Active Online Calibration of Multiple Sensors for Autonomous Surface Vessels

Hordur K. Heidarsson and Gaurav S. Sukhatme

Robotic Embedded Systems Laboratory, University of Southern California
Los Angeles, CA 90089, United States
{heidarss,gaurav}@usc.edu
http://robotics.usc.edu/resl

Abstract. We present an approach to actively calibrate the exteroceptive sensors of an Autonomous Surface Vessel (ASV) autonomously. The approach consists of locating suitable calibration sites in the environment from aerial imagery, navigating to them, gathering calibration data and estimating the required parameters from data. We have conducted experiments using an ASV in a lake to validate our approach.

Keywords: Autonomous Surface Vessel, Calibration

1 Introduction

Autonomous Surface Vessels (ASVs) have the potential to become extremely useful tools for variety of tasks, such as environmental surveying, monitoring and mapping. However, at present, they are not ready to be operated at all times in shared environments such as marinas or lakes with boat traffic without human supervision. For practical long-term and cost-effective operation, 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 a variety of weather conditions (e.g. sun, 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. While most obstacles have expressions both below and above the water surface, there are some that have only one or the other. It is therefore important that we are able to sense in both those domains. In order to do so we employ a laser and a camera above the surface and a sonar below.

For smaller ASVs, it can be useful to have a modular configuration, where sensors and equipment can be reconfigured easily, even in the field, to better suit the task at hand and as well as for easier transportation. By moving sensors around we have introduced another problem: now we might not know the transforms between the various sensors which are needed to be able to fuse data from them. This can be measured by hand, but it can be time consuming to

do so accurately and especially so in the field. Furthermore, utilizing traditional data driven calibration methods using ready-made targets is non-trivial due to the different nature of some of the sensors, e.g. a laser and a camera need an above water target and a sonar needs an underwater target of a different type, which could be hard to deploy. This leads us to develop an active data driven method to obtain the calibration between the sensors from certain features we find in our environment.

2 ASV Description and Problem Statement

The ASV we consider in this paper is equipped with the following sensors: a forward-down facing monocular camera, a planar laser range finder that is mounted facing forward (with its measuring plane level with the water surface plane) and a mechanically scanning pencil-beam underwater sonar mounted facing forward (just below the surface of the water). Additionally, the ASV is equipped with a GPS receiver and an IMU for pose estimation. A diagram of the ASV can be seen in *Fig. 1*.

The monocular camera is utilized as a range sensor to points on the water surface plane. This is done by segmenting the water from non-water in the camera output[1,17] and using the transform between the camera and water surface plane to project the camera image on to the plane. This means that all three sensors are sensing in different parallel planes: the laser slightly above the surface of the water, the camera directly on the surface of the water and the sonar just below the surface of the water. This enables us to detect a wide variety of obstacles, whether they have an above- or below- the surface expression or both.

The sonar we use is a single pencil beam mechanically scanning sonar. Roughly speaking it functions similar to a laser range finder, but since there's only a single beam it can only take one distance measurement at a time. The width of the sonar beam is also not as narrow as a laser (our unit has a beam width of 1.8°). The data received from the sonar is the echo strength binned evenly over the set range of the unit. This means that the sonar can detect more than one separate obstacle at a time if the obstacle doesn't cover the full width of the sonar beam. This can also increase noise due to reflections and other sources. The speed of sound in water varies significantly with temperature, requiring sonar data to be corrected (scaling) based on the water temperature. Our sonar does not do temperature compensation internally, requiring us either to scale the data based on a temperature reading, or, as we have done, estimate this scaling factor during our calibration.

We assign a reference frame to each of the sensors. The camera frame $\{C\}$, with origin at the optical center of the camera and with z-axis aligned with the optical axis of the lens, the laser frame $\{L\}$, with origin at the center of the laser range finder and x-axis straight forward in the sensing plane, the sonar frame $\{S\}$, with origin at the center of the sonar with the x-axis straight forward in the sensing plane and the IMU and GPS frame $\{I\}$, with its origin at the center of the IMU body. Additionally, we define a fixed water surface frame $\{W\}$, that

has the XY-plane on the surface of the water and origin fixed with respect to the ASV (for simplicity we define it to share its z-axis with $\{L\}$ so in 2D they can be used interchangeably).

Our problem is then as follows. With an ASV in the described sensor configuration and deployed in an unmapped environment, we want to estimate the planar rigid-body transform with added scaling between the laser and sonar, $\theta_{LS} = (a_S, R_{LS}, t_{LS})$ such that $X_L = a_S R_{LS} X_S + t_{LS}$ where X_L are points in the laser reference frame, $\{L\}$, corresponding to the points, X_S, in the sonar reference frame, $\{S\}$. We also want to estimate the homography transform, H_{WC}, between the (intrinsically calibrated) camera image plane and the water surface plane, such that $X_W = H_{CW}^{-1} p_C$ where X_W are points in the water surface reference frame, $\{W\}$, corresponding to the points p_C in the intrinsically calibrated camera image in homogeneous coordinates.

We start without any previous data from the sensors and assume that the intrinsic model parameters of the camera are known (or can be estimated with a different approach). We also assume that the sonar and laser are mounted with level pitch and roll on the ASV since otherwise their utility would be limited.

3 Related Work

There is a good body of work in cross-calibration of some types of sensors, especially involving cameras and lasers. One approach is to manually use a checkerboard camera calibration target and detect it simultaneously with both camera and laser [9]. Other approaches calibrate from outdoor scenes and find suitable features for calibration [18,12], but use a 3D laser.

Intrinsic calibration of a camera can be done using a calibration target [19] or using online self-calibration [5]. Calibration between camera and IMU can also be done online from data [10]. Sonar calibration is usually done using deployed spherical targets [4,6], but focuses on the more low-level problem of estimating the intrinsics of the sensor rather than its mounting location as we desire.

Work on sensing for small ASVs includes[11,2]. Sonar sensing has mainly been done for AUVs, but some of the techniques [16] can be applied on ASVs as well.

4 Technical Approach

In this section we will begin by discussing the requirements for performing the calibration between the different sensors which will serve to motivate how we have approached the problem. We will follow that by outlining our process, step by step and finally provide a detailed explanation of each of the steps.

4.1 Calibration of Sensors in Different Domains

As previously described, the sensors we are considering are not sensing the exact same picture, but each in a different parallel plane. This has a great effect on

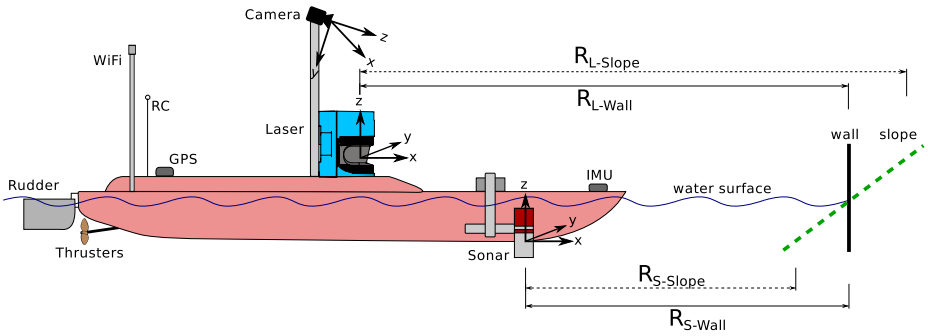

(a) Side view.

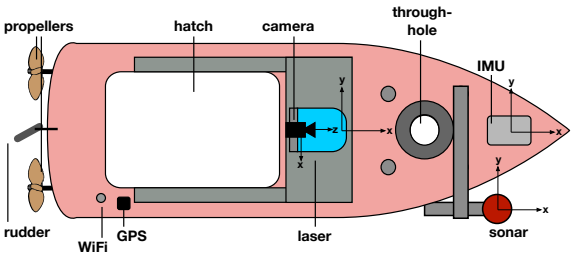

(b) Top view.

(c) In operation.

Fig. 1: The USC ASV. a) shows where the various sensors are mounted and a cross section of the sensing setup. The black vertical line and the angled dashed line to the left represent a vertical obstacle and a slanted one respectively. $R_{L-\{Wall,Slope\}}$ represent the detected ranges by the laser to the wall and sloped obstacle. $R_{S-\{Wall,Slope\}}$ similarly shows the detected ranges by the sonar. If the transform between laser and sonar is known, the detected ranges (converted to a common reference frame) are equal for the vertical obstacle, but differ for the sloped one, b) shows a top view, c) shows a photo of the ASV in operation.

what kind of data we can use to calibrate from. Looking at *Fig. 1a* one can see that when the different sensors look at a target that is not vertical, the different sensors will give conflicting range measurements to said target. Furthermore, since each sensor is in a different plane (and the transformation between them is currently not known), we won't even know that there is an inconsistency.

However, when looking at a vertical target that extends below the surface of the water, the sensors are measuring a common range, so this can be used for calibration. In the one dimensional case, this would be enough, but once we are in 2D, we have a set of lines that need to be aligned, but they could just slide along each other. Additionally, like previously mentioned, we are not able to tell if the target is slanted or not based on our sensor data. For these reasons, to have any hope of doing data association between the different sensors, we require our calibration target to have a corner and be vertical.

Since we want to calibrate from an existing feature in our environment, we would greatly benefit from some prior information on where to find those. To that effect we utilize overhead imagery to find potential calibration targets. With the overhead imagery we run into the same problem as before, it is not trivial to know if a potential target is vertical or sloped. However, most man made structures in water that have corners and appear straight edged (the main exception being ramps) tend to be vertical in the water. We will therefore build our approach on the assumption that a potential target from an overhead image will have a corner formed by straight edges, with the angle between them being close to 90°.

When calibrating the different sensors, we do not require to do all at once. It is sufficient to calibrate any two to each other at any time, but the additional sensor can provide additional redundancy against errors in case we are not looking at a suitable target.

4.2 Our Approach

Our calibration process is comprised of the following steps.

1. Collect overhead imagery.
2. Locate potential calibration targets.
3. Navigate to calibration site.
4. Collect calibration data.
5. Calibration.

Each step will now be described in detail.

Collect Overhead Imagery We start by gathering overhead image data. For up to date imagery an UAV with a downward facing camera can be used with simple auto generated way-point control to explore and gather overhead imagery. As an alternative, if available, existing aerial imagery can be used, such as from online mapping sources (at the risk of the imagery being outdated). We apply our previous approach [8] to detect water from the overhead imagery. We will refer to the resulting map as the obstacle map. This will facilitate knowing where the ASV can operate and provide input to our next step of locating potential sites for calibration.

Locate Potential Calibration Sites From the overhead imagery we want to locate potential calibration sites. To this end, we detect lines in the aerial imagery using a standard approach consisting of detecting edges in the image and then applying the probabilistic Hough transform [13]. Using the additional information about water/non-water from the previous step, we reject lines that appear either in mid-water (most likely ripples or noise) or lines that are on shore away from the water, leaving us with ones that close to the edge of the water. Now we cluster lines together based on location, where each cluster is a potential calibration site. Next we check all pairs of lines (first extending them by a small amount) for intersections to find corners and reject any that are outside some envelope of acceptable angles ($90° \pm n°$). Using this information allows us to assign a score based on number of corners, whether they are convex or concave, and length of the lines.

With the potential calibration sites found, we prioritize them by weighing the shortest path distance (estimated by planning a path on the obstacle map) from the ASV start location against the feature score. The site with the best combined score will be visited first. If there are many potential sites in the area, there might a benefit in finding the shortest expected path that leads to a suitable calibration target being found.

Navigate to Calibration Site Once we have picked a calibration site we need to navigate to it. In order to do so safely, we perform reactive obstacle avoidance along our planned path, based on a simple approach [7] using the sonar and laser individually. In case the site has multiple corners, we start with the closest one. We navigate to a point at a given distance out from the center line of the corner (see Fig. 2).

Collect Calibration Data Once in position, we keep station while performing a slow rotation to sweep the corner with our sensors to try to get good data points for the calibration. The need to keep relatively stationary comes from the fact that the mechanically scanning sonar has a single beam that sweeps back and forth in a arc. The scanning speed is limited by the speed of sound, so compared to the laser or camera, the sonar has a relatively slow update rate. Capturing a single frame (i.e. a full sweep from one side to the other), can take from 2 to 10 s, depending on the desired sector width and angular resolution. When a full frame has been captured, every measurement of it needs to be transformed to a common reference frame using pose estimates from the on-board pose estimate.

During the data collection phase, we need to verify that we are looking at a suitable calibration target before attempting to perform the calibration. Using the laser we can easily detect line and corner features with a fast update rate and low processing cost using a variety of algorithms [15]. After filtering out stray returns we use the probabilistic Hough transform on the laser data to extract lines. Next we merge close-by lines with similar angles and then elongate the remaining lines and check for intersections to locate corners.

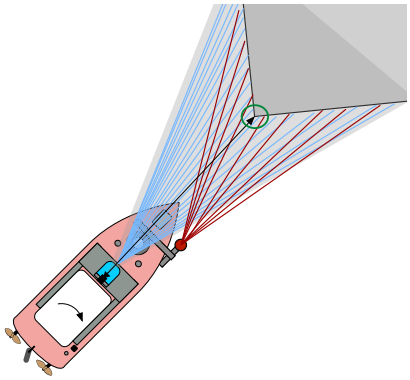

Fig. 2: ASV gathering calibration data.

Fig. 3: An example of a camera view during calibration. Below is a result from an edge extraction with multiple unwanted edges from water ripples.

Once a suitable laser frame is found, we check the corresponding sonar frame using a very similar process: filter away stray returns and extract lines to find corners, but since the sonar measurements are much noisier and less accurate than a laser, we use a variation of the Hough transform designed for underwater sonar [16] for improved accuracy. If corners are found in the laser and sonar images, we check if they are within a certain distance from each other (a rough upper bound on the translation between them) and add them to our calibration set. Example of the laser and sonar data can be seen in *Fig.* 5.

For the camera we use the fact that straight lines are preserved under projective transformation and attempt to find the corner(s) in the camera image. However, this is not as simple as for the overhead case or for the laser or sonar. Edge detection on the camera image generally results in many edges on the surface of the water in addition to those along our intended edges (refer to *Fig.* 3). These edges can be hard to filter out based on a single image. To remedy this, we use a technique based on [17] of employing optical flow [3] to do tracking over a few sequential frames. In our case, we place our trackers uniformly on the edges found in the image. Now we can compute the entropy of the tracked trajectories and also compute the dissimilarity between neighboring trajectories. Trackers with trajectories of high entropy or those highly dissimilar to their neighbors are very likely trying to track water and can be discarded along with their edges. We can now apply the Hough transform to the remaining edges in order to find lines which we then elongate and look for intersections. If we find a corner in the image which has no lines below it (*i.e.* larger y-coordinate), we save its coordinates along with the closest corner from the laser image as a pair

for the camera calibration. Since we are tracking the corner we can get series of point correspondences through the sweeping maneuver. For the case of stereo cameras, this process is easier since we should be able to find the corner from the disparity information fairly easily and then calibrate in the same way from a set of point correspondences.

Calibration Once we have collected enough calibration data we can start estimating the calibration parameters. For the laser and sonar this is done by registering the two point sets to each other. For this we utilize the Coherent Point Drift (CPD) algorithm [14], a probabilistic registration algorithm capable of solving for scale as well as the rigid transform, with a slight modification where we force correspondence probabilities, $P(i|s_j)$, to zero for point pairs that should not be matched (*i.e.* points in separate frames). This allows us to jointly estimate the transform on several frames of data at once, resulting in a more robust estimation.

For the camera calibration we utilize the RANSAC version of findHomography in OpenCV, with which we supply our point correspondences between the image coordinates and water frame coordinates.

5 Experiments and Results

In this section we will show the results for our two main pieces of the calibration process, namely the locating of potential sites, and then the calibration itself.

5.1 Experimental Setup

For our experiments we have used the USC ASV (*Fig. 1*), which is approximately 1.8 m long and 0.8 m wide. It is actuated with two electric thrusters and a rudder, and is capable of speeds up to 1.6 m/s. The sensor package consists of a Microstrain 3DM-GX3 IMU, a U-Blox GPS unit, a Pt. Grey Grasshopper 3 USB3 camera, a SICK LMS-200 laser range finder and an Imagenex 881L sonar. On-board control and computation is done with a Core i7 computer running ROS.

Our field experiments have been conducted in Puddingstone Reservoir, which is an artificial lake in San Dimas, CA, USA.

5.2 Locating Calibration Sites

Using aerial imagery from Google Maps since we were unable to fly our UAV at the test site, we have performed the second step of our approach off-line where we locate potential calibration sites. We have identified two potential calibration sites in the area of operation and are able to detect both of them from the overhead imagery, see *Fig.4*. The process resulted in several line edges along the shoreline of the lake, but they were rejected due the lack of suitable

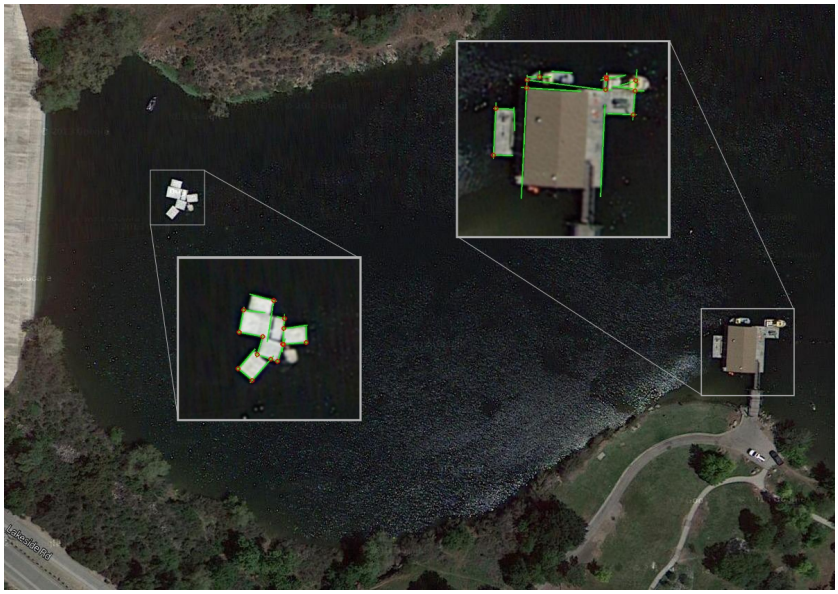

Fig. 4: Detected calibration sites in experimental area. Green lines indicate the detected lines and red circles indicate detected corners.

corners for the calibration. We also configured our process to use only convex corners as we did not want the ASV to navigate in to. This eliminated the dam on the left hand side of the image as a potential site.

Depending on where the ASV was launched the two calibration sites were ranked differently with the site in the middle of the lake being more valuable but the boat house site being closer to our launch spots.

5.3 Calibration

An example of the alignment between sonar and laser data can be seen in *Fig. 5*. This shows the alignment done on a single frame of data, which can work, but can suffer from scaling problem if there is only a single corner in the frame. To remedy that, we perform the alignment on multiple frames of data simultaneously, captured at different viewpoints so the scaling will stay fixed. *Fig.6* shows a camera frame from the calibration, where the detected lines and corner are overlaid on the image along with the tracking points.

Finally, *Fig. 7* shows the camera view with sonar and laser data overlaid using the estimated transforms.

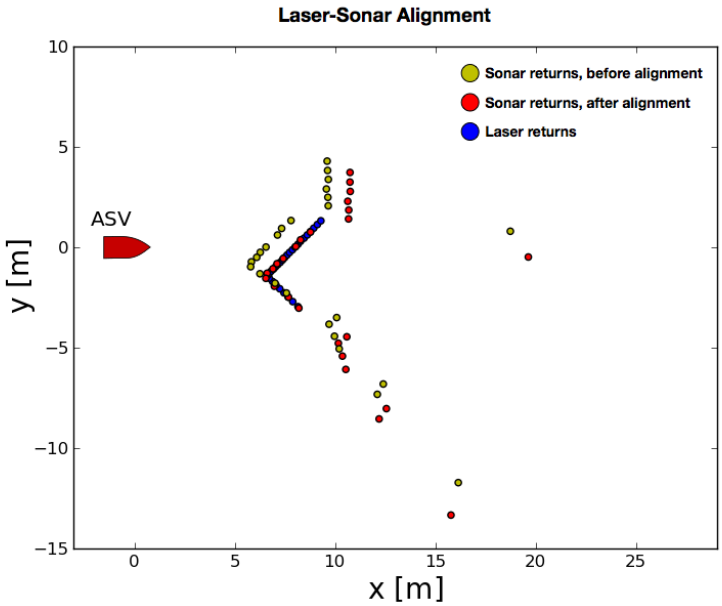

Fig. 5: Calibration between laser and sonar. The plot shows the alignment for a single frame from the calibration process. The blue dots are laser returns, the yellow are the uncalibrated sonar returns and the red are the calibrated sonar returns.

Fig. 6: Camera view during calibration. The floating platform in front of the ASV is our calibration target that is being sensed by the sonar and laser in Fig 5. The blue lines are straight lines detected in the image and the green circle indicates the corner that has been found. The blue diamonds are trackers that have been randomly placed on lines that have been found in previous frames.

Fig. 7: Camera view with laser returns (red dots) and sonar returns (dark blue dots or lines) overlaid using transforms estimated by the calibration process. It is worth noting that since the floating platforms rise so low out of that water that the laser might not get a hit from the closest one. Also worth noting is that the overlaid sonar returns include the whole echo profile, not only the peak, which explains why they appear as lines.

6 Conclusions and Future Work

6.1 Conclusions

We have presented an approach to calibrate different sensors sensing in different planes on an ASV. Our method is based on finding physical corner features in the environment with the aid of overhead imagery, navigating to them and leveraging them to obtain a calibration between the different sensors. We have performed field experiments with our ASV in order to test and verify our approach.

6.2 Future Work

We plan to investigate the feasibility of getting a rough estimate of the slope of a target from the return data of the sonar. Since the sonar data contains information over the whole sensing range we have some additional information that we are currently mostly discarding. The spread and shape of the return could give some insight into the characteristics of the obstacle that generates the return. This would help improving the calibration process as well as general mapping and obstacle avoidance using the sonar and the other sensors which is our main focus. We are also working on applying self-supervised learning techniques for using one sensor to improve the sensor model of another. This allows us to use a sensor in which we have greater trust or that has a more established sensor model.

6.3 Acknowledgments

Special thanks to RESL members Jnaneshwar Das, Stephanie Kemna, Jörg Müller and Carl Oberg for their help with conducting field experiments and the staff at Frank G. Bonelli Regional Park for being helpful and allowing us to conduct our experiments in the lake.

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 Bandyophadyay, 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. Jean-Yves Bouguet. Pyramidal implementation of the affine lucas kanade feature tracker description of the algorithm. *Intel Corporation*, 2:3, 2001. 01395.
4. D. Chu, K.C. Baldwin, K.G. Foote, Yanchao Li, L.A. Mayer, and G.D. Melvin. Multibeam sonar calibration: target localization in azimuth. In *MTS/IEEE Conference and Exhibition OCEANS, 2001*, volume 4, pages 2506–2510 vol.4, 2001.
5. J. Civera, Diana R. Bueno, A.J. Davison, and J. M M Montiel. Camera self-calibration for sequential bayesian structure from motion. In *IEEE International Conference on Robotics and Automation, 2009. ICRA '09*, pages 403–408, May 2009.

6. Kenneth G. Foote, Dezhang Chu, Terence R. Hammar, Kenneth C. Baldwin, Larry A. Mayer, Lawrence C. Hufnagle Jr, and J. Michael Jech. Protocols for calibrating multibeam sonar. *The Journal of the Acoustical Society of America*, 117(4):2013–2027, April 2005. 00074.
7. 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.
8. 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.
9. Abdallah Kassir and Thierry Peynot. Reliable automatic camera-laser calibration. In *Australasian Conference on Robotics and Automation*, 2010.
10. 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.
11. Jacques C. Leedekerken, Maurice F. Fallon, and John J. Leonard. Mapping complex marine environments with autonomous surface craft. Delhi, India, December 2010. 00000.
12. Jesse Levinson and Sebastian Thrun. Automatic online calibration of cameras and lasers. In *Proceedings of Robotics: Science and Systems*, 2013.
13. J. Matas, C. Galambos, and J. Kittler. Robust detection of lines using the progressive probabilistic hough transform. *Computer Vision and Image Understanding*, 78(1):119–137, April 2000.
14. Andriy Myronenko and Xubo Song. Point set registration: Coherent point drift. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32(12):2262–2275, 2010.
15. V. Nguyen, A. Martinelli, N. Tomatis, and R. Siegwart. A comparison of line extraction algorithms using 2D laser rangefinder for indoor mobile robotics. In *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2005. (IROS 2005)*, pages 1929–1934, August 2005.
16. David Ribas, Pere Ridao, Juan Domingo Tardos, and Jose Neira. Underwater SLAM in man-made structured environments. *Journal of Field Robotics*, 25:898–921, November 2008. 00071 ACM ID: 1464494.
17. P. Santana, R. Mendonca, and J. Barata. Water detection with segmentation guided dynamic texture recognition. In *2012 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pages 1836–1841, 2012.
18. Davide Scaramuzza, Ahad Harati, and Roland Siegwart. Extrinsic self calibration of a camera and a 3d laser range finder from natural scenes. In *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*, pages 4164–4169. IEEE, 2007. 00097.
19. Zhengyou Zhang. A flexible new technique for camera calibration. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(11):1330–1334, 2000.