

Mapping Complex Marine Environments with Autonomous Surface Craft

Jacques C. Leedeckerken, Maurice F. Fallon and John J. Leonard

Abstract This paper presents a novel marine mapping system using an Autonomous Surface Craft (ASC). The platform includes an extensive sensor suite for mapping environments both above and below the water surface. A relatively small hull size and shallow draft permits operation in cluttered and shallow environments. We address the Simultaneous Mapping and Localization (SLAM) problem for concurrent mapping above and below the water in large scale marine environments. Our key algorithmic contributions include: (1) methods to account for degradation of GPS in close proximity to bridges or foliage canopies and (2) scalable systems for management of large volumes of sensor data to allow for consistent online mapping under limited physical memory. Experimental results are presented to demonstrate the approach for mapping selected structures along the Charles River in Boston.

1 Introduction

Recent research in 3D mapping with autonomous robots has shown promising results for terrestrial and aerial vehicles operating in a wide range of environments [18, 3, 29, 20]. Research in the marine domain has primarily focused on bathymetric mapping and mine counter-measures with AUVs and ROVs [12, 13, 14, 24]. Typical AUV and ROV operations are performed in oceans with sufficient depth for launching the vehicle from a ship. Instead, we focus on shallow environments such as rivers and harbors. Related underwater mapping work in partially structured marine environments is demonstrated in [22, 23].

Our work presents a novel system for mapping **both above and below** the water surface using a surface craft equipped with an imaging sonar for subsurface perception as well as LIDAR, camera, and radar sensors for perception above the surface. By simultaneous sensing both above and below the water, we can improve the ac-

Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139.
jckerken,mfallon,jleonard@mit.edu

curacy of underwater maps using precise constraints derived from above surface perception and localization. To the authors' knowledge, our platform is unique in its ability to capture terrestrial and subsurface environments **simultaneously**.

Accurate mapping in marine and riverine environments in the presence of objects and buildings which block GPS reception, such as bridges and foliage, remains an open challenge. However it is possible to carry out sensor measurements of these same objects. We propose to use these sensor measurements to relatively localize the vehicle during GPS-denial, and hence to correct the resultant map produced by our system.

This problem fits within the broader context of Robust SLAM – creating mapping systems that can deal with realistic cluttered real-world situations including erroneous constraints, GPS dropouts, and sparse environmental structure as well as the observation of moving objects. Navigation and mapping systems for land vehicles experience similar issues when operating near tall buildings (the 'urban canyon' effect) or when traveling through tunnels.

1.1 Mapping in Sparse Environments

One approach to such corrupt GPS readings could be to simply ignore spurious position estimates. This is indeed possible for mapping dense environments — provided that sufficient environmental structure is continuously visible to the sensor before, during and after GPS dropout. However the marine environment is typically characterized by large open spaces in which little or no above-water environmental structure is within sensor range. In fact, the corruption of GPS measurements occurs specifically when the vehicle approaches these structures of interest. Hence, a key problem is dealing with the transitions that occur from periods of good GPS (but with little structure visible) to poor or no GPS (when close to structures of interest).

The accuracy of underwater mapping with AUVs is limited by the necessity to model in the absence of ground truth and an inability to cross-validate using a different sensing regime. For example, EKF-derived uncertainty can provide a measure of how well observations compare to predictions given such a model — with an underlying assumption that the model is sufficient.

However when carrying out a survey mission of an unknown structure or territory a prior map will be unavailable. The resultant map is inevitably be a projection of the observations, given only the estimated vehicle trajectory.

Unlike AUVs with sustained and complete GPS deprivation during operations, an autonomous surface craft will suffer from GPS deprivation for much shorter durations (although non-linear GPS distortion due to the structures is an alternative issue). The vehicle trajectory estimate can be improved using a combination of GPS and surface range sensing. Finally projections of surface laser sensing may be qualitatively assessed (with even small errors leaving visible artifacts), whereas direct validation of subsurface acoustic sensing is difficult.



(a)



(b)



(c)

Fig. 1 The ASC sensor suite shown in Figure 1(a) includes two Hokuyo LIDARs on a 1Hz rotating shaft (top center), three SICK LIDARs (center), an Oxford Technical Systems integrated MU and GPS (antennas fore and aft), a Blueview MB2250 imaging sonar (below hull), and an RDI DVL (under hull, not visible). Target environments are shown in Figure 1(b) and Figure 1(c).

In the following section we describe the layout and unique sensing capabilities of our ASC as well as our proposed solution to mapping with repeated loss and reacquisition of GPS. In Section 4, we provide our experimental results, focusing on several structures in the Charles River near MIT including both the Harvard and Longfellow Bridges. Finally, in Section 5 we conclude the paper with a discussion of future research directions that our work enables.

1.2 Platform

Our approach to mapping marine environments begins with the construction of an ASC platform with novel sensing capabilities and sufficient computational power. Our platform is based on the hull of a SCOUT robotic kayak [8] that has been modified to provide the sensing capabilities required for our application. This small commercial kayak of approximately 2.8m length and 0.76m width. The vehicle is actuated with a single servoed thruster capable of speeds near 2m/s in calm water. A standard 12V marine battery provides power for 2.5 hours under full sensor and thruster load. The vehicle is pictured in Figure 1.

For positioning and localization, the vehicle has an Oxford Technical Systems integrated IMU and GPS and well as an RDI Doppler Velocity Logger (DVL) which estimates rate of motion through the water. Backup sensors include a Garmin 18x-5Hz GPS receiver and a Ocean Server 5000 3-axis magnetic compass. The terrestrial sensors include three SICK LMS291 scanning LIDARs, a low-cost webcam as well as two 30m range Hokuyo LIDAR mounted on a shaft rotating at approximately 1Hz. Two of the SICK LIDARs are mounted to scan vertically as shown in 1(a) — so as to permit 3D reconstruction as the vehicle passes along side features of interest.

Finally for underwater sensing, the vehicle has a Blueview MB2250 imaging sonar which has a 10m range and a 45 degree field-of-view. The sonar was mounted at a 45 degree pitch downwards which allows it to observe a vertical slice of the environment in front of the vehicle and downwards. The computational core of the vehicle consists of two motherboards carrying a quad-core 2.83 GHz CPU and a low-power dual-core 1.6 GHz CPU.

2 Technical Approach

We adopt a sub-mapping approach to the SLAM problem using a graph of local maps similar to the work in [5, 25, 19]. Local map estimation is achieved with Smoothing and Mapping (SAM), which is detailed in [9, 16]. While the SAM approach optimizes a graph of constraints from observations, the generation of quality constraints is subject to implementation. A key contribution of this work is the generation of the constraints from 3D lidar data to properly localize after GPS loss and reacquisition. Figure 3 shows the effects of repeated GPS loss and reacquisition on mapping error of a bridge. Using the lidar, we may generate constraints to localize the vehicle and reduce map error in the SAM framework.

We propose the following method for localization from lidar point clouds. Each local map will contain an octree of voxels, a 3D analogue of occupancy grids [28]. A similar approach to this was taken in [11] for mapping underwater tunnels. Using submaps generated over short durations, we construct an octree of the lidar scans from the three static LIDAR at a fine resolution (2.0cm). Then we conduct a match against the local map using Iterative Closest Point (ICP) [30] on the voxels using




Fig. 2 Left: GPS distortions without complete signal loss occurring while near a bridge. Trajectories are false-colored with the number of reported satellites available. Right: Vehicle trajectory computed by dead reckoning (green) shown in comparison to the estimated trajectory (blue). Together the GPS hazards and dead reckoning error represent the primary challenges of this research.

the IMU data to provide a prior estimate of the candidate octree to the local map. We use the point-to-point ICP variant over point-to-surface due to the complexity of surface normal estimation for skeletal steel, railings, and masts common in bridges and marinas.

The importance of such a registration method is demonstrated in Figure 2(a) and Figure 3. Figure 2(a) shows the GPS position estimates of the vehicle making two passes along side a portion of the Harvard Bridge. The outward southern path was a few metres further from the bridge than the northerly return path, otherwise the two paths were of similar smoothness. However the GPS position estimates of the return path shows pronounced degradation which is demonstrated by the discontinuities such as at (-300m, 475m). These discontinuities correspond to a small number of visible satellites. Figure 3 shows the effect that such behavior can have on the resultant map of the structure. The continuous observation of a sufficient GPS constellation can yield a reasonably consistent map (red). This compares with the map produced by winding through the bridge, which accumulates significant error over the course of the survey (yellow-green).

Before explaining how these two scenarios can be properly classified (in section 2.2), we will first briefly explain how the sonar data can be formed into a representation comparable to the LIDAR point returns commonly used in terrestrial mapping.




Fig. 3 GPS signal loss and reacquisition leading to inconsistent localization and mapping. The winding traversals under the bridge (top, blue-green) accumulate error from GPS loss compared to the subsequent traversal along the side of the bridge (bottom, red) without GPS loss, yielding more consistent data but with incomplete coverage. Our SLAM algorithm corrects for the errors shown in the former. Note that the bridge length is over 600 meters, hence the figure is segmented as the arrows indicate.

2.1 Sonar Processing

Each ping of the imaging sonar provides an image covering the 45 degree field of view within the sensor's operating range of 1.0 to 10.0 meters. Pixel values represent intensities of the acoustic sonar returns. See Figure 4 for an example image. Processing of this image begins with an implementation of the Canny Edge Detector [7]. A one-dimensional range scan is created taking the range pixel with maximum intensity value preceded with a positive gradient for each bearing angle. Ranges less than 1m or having intensity values below a threshold are ignored.

The range data is projected into a voxel tree (as mentioned in the preceding section) and reduced to a digital elevation model by estimating a single depth value for occupied cells in the column at a given (x,y) coordinate. Projection into the voxel tree maps higher intensities to greater probability of occupancy. We also process the grid to remove extraneous midwater noise by removing cells at 2D grid coordinates (x,y) having multiple occupied cells z cells. When occupied neighbor (x,y) cells exist, only the z cells within a threshold distance of neighbors persist. A second pass retains the cell most consistent with the nearest neighbors (within two meters). These steps remove a significant proportion of the outliers from the bathymetry map. However due to the small sensor field of view and resultant sparse coverage, insufficient neighbors exist at boundaries of the coverage area.




Fig. 4 Example sonar ping image and corresponding range data. Image on left is color mapped for visibility, and image on right overlays extracted range data over the raw gray scale intensity image. Note the sonar is constructed of two physical staves, which causes the discontinuity.

Having explained how the sensor data is represented and efficiently stored, in the following section we will discuss motion modelling to identify GPS loss and recovery.

2.2 Motion Modeling

Dead reckoning is the integration of proprioceptive vehicle velocities so as to create a running vehicle position estimate. We fuse dead reckoning with measurements of GPS position, IMU rates, and DVL velocities within a Kalman Filter framework. Under conditions with poor GPS satellite visibility, GPS measurements become unreliable and biased. We address these issues with an adaptive error model. Using the device-reported GPS satellite count and dilution of precision (DOP) measurements as well as custom metrics from the lidar sensors we form an error model which gives one of two modes: S_0 and S_1 :

- S_0 : encompasses normal operations without obstructions to GPS. We assume diagonal covariances for GPS measurements with dilution of precision scaling in addition to the IMU rates, and DVL velocities.
- S_1 : occurs when GPS measurements are unreliable and covariances are scaled so as to have negligible effect (unless entirely rejected due to having an insuffi-

cient number of satellites visible). In this mode the IMU accelerometer and DVL velocities dominate the motion estimate.

In practice mode S_1 persisted for short durations during underpasses of the bridge trusses (approximately 10-20s). The proper double integration of IMU accelerations and the low frequency DVL velocities (0.5Hz) can result in reasonably slow accumulation of error during these periods.

The framework used to determine the motion model state is a linear classifier with a feature vector space \mathcal{F} for model discrimination where elements of $v \in \mathcal{F}$ are as follows:

- The GPS satellite count
- The GPS reported dilution of precision (DOP)
- The IMU reported velocity error
- The inverse mean of overhead lidar returns for the two vertical LIDAR taken together (i.e. within 30 degrees of vertical)
- The inverse mean of lateral lidar returns for the two LIDAR - taken separately for the left and right sides (within 15 degrees of horizontal).

Using several manually selected sequences of experimental data where the vehicle entered a region of GPS degradation, we manually selected a time boundary to discriminate between accurate GPS measurements prior to the boundary and inaccurate GPS measurements after the boundary. Using an optimization routine to determine the weighting vector which minimizes the classification error, we can then use the inner product of the feature vector with the weight vector to determine a classification as to when it would be appropriate to use the second vehicle error model

$$S(v) = \begin{cases} S_0 & v^T w \leq 0 \\ S_1 & v^T w > 0 \end{cases} \quad (1)$$

An example application of the method for real data is shown in Figure 6. Figure 6(a) results from using only the standard model S_0 . GPS distortions and the dependence of IMU velocities with GPS measurements internal to the IMU/GPS device's filter result in a path showing significant distortion of sensor data. Incorporating the second model shows significant improvement in Figure 6(b), yet minor error is visible during model transition as seen in the lower left portion of the figure.

3 Map Representation

The representation of the mapped environments using occupancy grids has been a classical approach in the robotics community. The pioneers of occupancy grid maps were Elfes [10] and Moravec [17]. The basic principle of an occupancy grid approach is to represent the environment as a field of binary random variables in a grid of cells. The random variable for each cell represents whether the location of

the cell is occupied or not. These occupancy grid cells are often given referred to as pixels in 2D maps and voxels in 3D maps.

In occupancy based mapping, the environment is partitioned in cells m_k such that the map m is the set $\{m_k\}$, and a function f that maps the cell index k to a physical location $l_k \in \mathbb{R}^d$. This allows the overall map probability to be formed as product of individual cells

$$p(m|z_{1:t}, x_{1:t}) = \prod_k p(m_k|x_{1:t}, z_{1:t}) \quad (2)$$

In practice the log-odds representation, shown in Equation 3 and 4, of each voxel's probability is used for numerical stability and simplified updates.

$$l_{k,t} = \log \frac{p_{k,t}}{1 - p_{k,t}} \quad (3)$$

$$p_{k,t} = 1 - \frac{1}{1 + \exp l_{k,t}} \quad (4)$$

3.1 Data Structures

A challenge in mapping large three dimensional environments is to develop a memory efficient representation. While feature-based representations reduce sensor data to a low dimensional feature parameter space, feature-based maps tend to under-represent the true structure by rejecting data not consistent with a limited alphabet of prior feature models. Conversely, metric grid maps work well with unstructured environments or when sufficient feature models are unavailable, but dense metric maps require large amounts of memory; thus, resolution and map scale compete under limited memory.

So as to exploit sparsity in natural environments, metric maps can be formed using efficient representations such as sparse voxel trees (an analogue of the two-dimensional quadtree). We employ an octree data structure of this form. While binary space partitioning methods, such as kd-trees provide an alternative representation, octrees offer several practical advantages for our application.

The key advantage of sparse voxel octrees comes from expansion of nodes on demand — rather than maintaining a dense grid of all addressable cells regardless of whether entire subtrees of cells are identical.

Each submap maintains an octree of this type. This limits the map size to a given resolution. A new submap is required when measurements approach the boundary of the current submap. In addition to the capacity rule for submap generation, we introduce a rule for starting a new submap at model transitions as described in section 2.2. In effect small submaps are created for each traversal under bridges, which simplifies later alignment into the global frame.

3.2 Map Registration

Our approach to solving the problem of registration, or alignment, of point sets uses the Iterative Closest Point (ICP) family of algorithms [4, 30]. The two major components of the ICP registration problem are: (1) correspondence and (2) minimization of an objective function. The correspondence problem finds matches between data points from two data sets. Given a model set $M = \{m_j\}$ and data set $D = \{d_k\}$ and an approximate prior transform $T_0 = (R_0, t_0)$, correspondences between the model set and data set are found using nearest neighbor euclidean distance. A region of interest is defined based upon the expected overlap of the maps given the prior. Given the correspondences, updates to the transform are computed as in [15]. In cases where maps have insufficient overlap or the resulting transform seems erroneous by distance relative to the prior, the computed transform is rejected.

4 Experiments

We present experimental results in the vicinity of several structures spanning or along the Charles River between Boston and Cambridge, MA. The Harvard and Longfellow Bridges, three sailing pavilions and a yacht club provide structures of interest, having both extensive superstructure and subsurface foundations.

The experiments covered an operating area of approximately 1.6km by 1.1km area with seven missions ranging from 30min to 90min with the longest trajectory being nearly 7km.

For continuity and clarity of the visual structure, we present mesh reconstructions of the resulting maps. Figure 7 shows a reconstructed map of the operating region. Due to memory constraints for visualization, the overview map is not shown in full detail. Instead we also present selected portions of the global map at full resolution. Figure 8 shows a detailed reconstruction of the Harvard Bridge. The superstructure coloring represents different submaps aligned into the global frame, with the bathymetric data is shown as a single color for clarity. A georeferenced 3D model of the Harvard Bridge piers (shown in Figure 5 and derived from historical blueprints from the Library of Congress[1]) is presented for quality assessment. Detailed views of a sailing pavilion are shown in Figure 9. Part of our ongoing work is generating a quantitative error metric of the maps relative to the model of the Harvard Bridge piers.

Figure 6 shows the vehicle traversing under a bridge while suffering GPS degradation, and projected range data displayed for the trajectory. The dead reckoned path using the naïve measurement error model in Figure 6(a) is shown relative to the optimized trajectory in Figure 6(b) using the improved vehicle model and range registration. However, small errors remain in the optimized result as seen in the lower left of Figure 6(b) suggesting improvements in registration in future work.





Fig. 5 Model of Harvard Bridge from Library of Congress — used for validating map quality.


5 Conclusions and Future Work

This paper has described a new system and approach for robotic mapping and navigation in complex marine environments. We constructed a novel platform designed for operation in marine environments and equipped for perception both above and below the water surface. We have presented a mapping framework to accurately capture the three dimensional environment on a large scale with dense sensor sampling. The proposed method addresses the technical difficulties of GPS deprivation on navigational and mapping accuracy in our submap SLAM approach and presents experimental results in which surface sensing validates the accuracy of subsurface mapping.

Three key goals for future work include (1) incorporation of autonomous path planning, (2) improving the robustness of local map matching, and (3) increasing the efficiency of the online map representation. Integrating mapping with autonomous path planning will be necessary to achieve complete surveys of a region of interest while performing collision avoidance. This would address the problem of coverage gaps due to the limited sonar footprint of our existing system, which was remotely piloted. In our results, ICP registration of maps for portions of bridges had noticeable error in some instances. Symmetry in the longitudinal direction of the bridge could be a cause of some of the error. A promising direction for improved error models is to apply machine learning techniques to estimate hidden states for parameters such as the time-variant GPS covariances, extending upon the work in [2]. Approaches to improve efficiency include the use of multiple resolutions [27] or voxel content distribution attributes, as seen in [6, 26, 21].



(a)



(b)

Fig. 6 Area C of Figure 7: A comparison of projected range data using the naïve trajectory (a) to the smoothed trajectory for a traversal under a large bridge. Laser color mapped by height (green to red), sonar colormapped by intensity (tan colors), and trajectory shown as a single yellow line.




Fig. 7 Reconstructed map data for a portion of the Charles River near MIT. The Harvard Bridge is on the left and the MIT Sailing Pavilion is at the top of the figure.




Fig. 8 Area A of Figure 7: A bridge reconstruction with scattered coloring for individual submaps. A prior model of the bridge piers is shown for comparison. Note the base of the piers are submerged, but sonar coverage was insufficient to map them.

Acknowledgements

This work has been supported by the Office of Naval Research under Grants N00014-06-10043, N00014-05-10244 and N00014-07-11102 and by the MIT Sea Grant College Program under grant 2007-R/RCM-20.




Fig. 9 Area B of Figure 7 : Two views of the reconstructed map data for the MIT Sailing Pavilion: upper: lidar data only; lower: lidar and sonar data.

References

1. Harvard bridge, spanning charles river at massachusetts avenue, boston, suffolk county, ma. <http://www.loc.gov/pictures/item/MA1293/>
2. Abbeel, P., Coates, A., Montemerlo, M., Ng, A.Y., Thrun, S.: Discriminative training of kalman filters. In: Proceedings of Robotics: Science and Systems. Cambridge, USA (2005)
3. Bachrach, A., He, R., Roy, N.: Autonomous flight in unstructured and unknown indoor environments. In: European Micro Aerial Vehicle Conference (2009)
4. Besl, P., McKay, N.: A method for registration of 3-d shapes. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **14**(2), 239–256 (1992)
5. Bosse, M., Newman, P., Leonard, J., Teller, S.: Simultaneous localization and map building in large-scale cyclic environments using the Atlas framework. *Intl. J. of Robotics Research* **23**(12), 1113–1139 (2004)
6. Bosse, M., Zlot, R.: Continuous 3D scan-matching with a spinning 2D laser. In: IEEE Intl. Conf. on Robotics and Automation (ICRA). Kobe, Japan (2009)
7. Canny, J.: A computational approach to edge detection. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **8**(6), 679–698 (1986). DOI 10.1109/TPAMI.1986.4767851
8. Curcio, J., Leonard, J., Patrikalakis, A.: SCOUT - a low cost autonomous surface platform for research in cooperative autonomy. In: Proceedings of the IEEE/MTS OCEANS Conference and Exhibition. Washington DC (2005)
9. Dellaert, F., Kaess, M.: Square Root SAM: Simultaneous localization and mapping via square root information smoothing. *Intl. J. of Robotics Research* **25**(12), 1181–1203 (2006)
10. Elfes, A.: Integration of sonar and stereo range data using a grid-based representation. In: Proc. IEEE Int. Conf. Robotics and Automation (1988)
11. Fairfield, N., Kantor, A.G., D., W.: Real-time SLAM with octree evidence grids for exploration in underwater tunnels. *J. of Field Robotics* (2007)
12. Fairfield, N., Wettergreen, D.: Active localization on the ocean floor with multibeam sonar. In: Proceedings of the IEEE/MTS OCEANS Conference and Exhibition, pp. 1–10 (2008). DOI 10.1109/OCEANS.2008.5151853

13. Folkesson, J., Leederkerken, J., Williams, R., Leonard, J.: Feature tracking for underwater navigation using sonar. In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. San Diego, CA (2007)
14. Folkesson, J., Leederkerken, J., Williams, R., Patrikalakis, A., Leonard, J.: A feature based navigation system for an autonomous underwater robot. In: *Field and Service Robotics (FSR)*, vol. 42, pp. 105–114 (2008)
15. Horn, B.K.P.: Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America* **4**, 629–642 (1987)
16. Kaess, M., Ranganathan, A., Dellaert, F.: iSAM: Incremental smoothing and mapping. *IEEE Trans. Robotics* **24**(6), 1365–1378 (2008)
17. Moravec, H.: Sensor fusion in certainty grids for mobile robots. In: *Sensor Devices and Systems for Robotics*, pp. 253–276. Springer-Verlag (1989). Nato ASI Series
18. Newman, P., Cole, D., Ho, K.: Outdoor SLAM using visual appearance and laser ranging. In: *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*, pp. 1180–1187. IEEE (2006)
19. Ni, K., Steedly, D., Dellaert, F.: Tectonic SAM: Exact, out-of-core, submap-based SLAM. In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pp. 1678–1685 (2007)
20. Nüchter, A., Lingemann, K., Hertzberg, J., Surmann, H.: 6D SLAM-3D mapping outdoor environments. *Journal of Field Robotics* **24**(8-9), 699–722 (2007)
21. Pauly, M., Gross, M., Kobbelt, L.: Efficient simplification of point-sampled surfaces. In: *Visualization, 2002. VIS 2002*. IEEE, pp. 163–170 (2002). DOI 10.1109/VISUAL.2002.1183771
22. Ribas, D., Ridao, P., Neira, J., Tardós, J.: SLAM using an imaging sonar for partially structured underwater environments. In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)* (2006)
23. Ribas, D., Ridao, P., Tardós, J., Neira, J.: Underwater SLAM in man-made structured environments. *Journal of Field Robotics* **25**(11-12), 898–921 (2008)
24. Roman, C., Singh, H.: Improved vehicle based multibeam bathymetry using sub-maps and SLAM. In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pp. 3662–3669 (2005)
25. Roman, C., Singh, H.: Consistency based error evaluation for deep sea bathymetric mapping with robotic vehicles. In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. Orlando, FL (2006)
26. Rusu, R.B., Blodow, N., Marton, Z.C., Beetz, M.: Aligning point cloud views using persistent feature histograms. In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. Nice, France (2008)
27. Ryde, J., Hu, H.: 3D mapping with multi-resolution occupied voxel lists. *Autonomous Robots* **28**(2) (2010)
28. Thrun, S.: Learning occupancy grids with forward sensor models. In: *Proceedings of the Conference on Intelligent Robots and Systems (IROS'2001)*. Hawaii (2001)
29. Thrun, S., Montemerlo, M., Dahlkamp, H., Stavens, D., Aron, A., Diebel, J., Fong, P., Gale, J., Halpenny, M., Hoffmann, G., Lau, K., Oakley, C., Palatucci, M., Pratt, V., Stang, P., Strohband, S., Dupont, C., Jendrossek, L.E., Koelen, C., Markey, C., Rummel, C., van Niekerk, J., Jensen, E., Alessandrini, P., Bradski, G.: Stanley: The robot that won the DARPA grand challenge. *Journal of Field Robotics: Special Issue on the DARPA Grand Challenge* **23**, 661–692 (2006)
30. Zhang, Z.: Iterative point matching for registration of free-form curves and surfaces. *Intl. Journal of Computer Vision* **13**(2), 119–152 (1994)