

# Multi-Robot Exploration of Underwater Structures<sup>\*</sup>

Marios Xanthidis<sup>\*</sup> Bharat Joshi<sup>\*\*</sup> Jason M. O’Kane<sup>\*\*</sup>  
Ioannis Rekleitis<sup>\*\*</sup>

<sup>\*</sup> SINTEF Ocean, Trondheim 7010, Norway  
(e-mail: [marios.xanthidis@sintef.no](mailto:marios.xanthidis@sintef.no)).

<sup>\*\*</sup> Department of Computer Science & Engineering,  
University of South Carolina, Columbia, SC 29208 USA  
(e-mail: [bjoshi@email.sc.edu](mailto:bjoshi@email.sc.edu), [{jokane,yiannistr}@cse.sc.edu](mailto:{jokane,yiannistr}@cse.sc.edu)).

**Abstract:** This paper discusses a novel approach for the exploration of an underwater structure. A team of robots splits into two roles: certain robots approach the structure collecting detailed information (proximal observers) while the rest (distal observers) keep a distance providing an overview of the mission and assist in the localization of the proximal observers via a Cooperative Localization framework. Proximal observers utilize a novel robust switching model-based/visual-inertial odometry to overcome vision-based localization failures. Exploration strategies for the proximal and the distal observer are discussed.

Copyright © 2022 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

**Keywords:** Underwater Robotics, Mapping, Localization, Visual-Inertial Odometry, Cooperative Localization.

## 1. INTRODUCTION

Mapping and inspection of underwater structures is essential for a variety of fields such as marine archaeology, resource utilization, environmental monitoring, marine infrastructure maintenance, and public safety. Especially climate change, which potentially threatens many marine infrastructure and ecosystems, is expected to emphasize the need for scalable solutions for mapping and inspection in many of these domains. Currently, such underwater operations are performed by human operators, either as divers or Remotely Operated Vehicle (ROV) operators. Aside of the constant health risks and limitations associated with commercial divers, utilizing human operators requires significant logistics where scalability remains a major challenge. On the other hand, Autonomous Underwater Vehicles (AUVs) could assist such operations and not be bounded by the above-mentioned limitations, scaling better and being able to support the increased needs for such operation in the future.

The two major challenges for effective and safe operation of AUVs in the underwater domain are robust state estimation and motion planning with limited information. Regarding state estimation, previous work has shown the unique challenges that the underwater domain imposes on state-of-the-art SLAM (Joshi et al., 2019; Quattrini Li et al., 2016) approaches, and SLAM techniques such as SVIn2 (Rahman et al., 2018, 2019) improved the estimation accuracy. While minor state estimation errors might not affect significantly underwater operations locally, accumulated minor errors, especially while mapping



Fig. 1. Aqua2 AUV navigating towards the Stavronikita shipwreck, Barbados.

large structures, can significantly degrade the quality of the data collected from the target underwater structure, resulting in distorted maps. A novel approach is outlined in this paper utilizing a switching estimator employing an accurate VIO estimator (Rahman et al., 2019) and a model based estimator (Meger et al., 2015). Especially during navigation around underwater structures, when the AUV reaches the boundaries of the structure the cameras see only “blue water” and visual tracking is lost. The proposed switching estimator utilizes the model-based estimator to track the robot’s pose until VIO re-establishes tracking.

Regarding motion planning, underwater navigation poses significant challenges due to the noisy and limited sensing in the underwater domain, the external forces due to currents, the presence of many dynamic obstacles, and the often complex hydrodynamics experienced by many AUVs. Recent prior work has provided a robust vision-based real time 3D underwater navigation package called AquaNav (Xanthidis et al., 2020) that enabled an AUV

<sup>\*</sup> The authors would like to acknowledge the generous support of the National Science Foundation grants (NSF 2024741, 1943205) and the Research Council of Norway (ResiFarm, NO-327292)

with complex dynamics to navigate efficiently towards a desired goal and avoid obstacles safely in very challenging environments. Moreover, an extension of AquaNav called AquaVis (Xanthidis et al., 2021) was able to move the robot safely while it was also maximizing visibility of certain features in large structures such as shipwrecks, showing strong potential for mapping larger underwater structures.

The contribution of this paper is a novel framework where a team of robots collaborates to map large underwater structures such as shipwrecks in order to improve map quality and reduce uncertainty at post-processing stage. The main idea, that utilizes the previously introduced techniques, is to separate the robots into two teams, the *proximal observers* that greedily attempt to map the target structure in close proximity emphasizing the collection of detailed measurements, and the *distal observers* that will maintain a global view of the situation by observing the proximal observers along with its relative position to the general structure of the shipwreck. A deep learning framework (Joshi et al., 2020) for estimating the relative pose between the two kinds of robots is also outlined.

Specifically, this paper presents technologies towards accomplishing the proposed framework, such as the cooperative localization module, the robust underwater state estimation method, along with the safe and efficient underwater navigation and active perception techniques for the two different behaviors developed.

## 2. RELATED WORK

Mapping shipwrecks is a problem that has been addressed using a variety of methodologies all around the globe, with the two most famous wreck explorations of the Antikythera (Williams et al., 2016) and the Titanic (Eustice et al., 2006) shipwrecks. Photogrammetry of manually collected images resulted in mosaics (Demesticha et al., 2014), or from an ROV, (Nornes et al., 2015). Comprehensive overviews of relevant robotic technologies is provided in works such as Menna et al. (2018), and the Arrows EU project (Allotta et al., 2015).

Kurazume et al. (1994); Kurazume and Hirose (2000) first conceptualized the concept of utilizing mutual sensor reading to localize two robots, and named it as cooperative positioning system, though in the literature, it is most often called Cooperative Localization (CL) (Rekleitis et al., 1998). Dieudonné et al. (2010) proved the NP-hardness for an arbitrary set of sensors, and theoretical analysis factors affecting error growth were studied in Roumeliotis and Rekleitis (2004). Additional research has examined the performance (Mourikis and Roumeliotis, 2006), the consistency (Huang et al., 2009), the sensing modalities (Zhou and Roumeliotis, 2012), proposed decentralized solutions (Leung et al., 2010) and introduced the integration with inertial sensors (Martinelli and Renzaglia, 2017).

Another approach to mapping underwater is to cast it as a coverage problem. Frolov et al. (2014) presented a motion planning approach for map uncertainty reduction by returning to regions with high uncertainty, while Chaves et al. (2016) proposed the use of loop closures for reducing

uncertainty. Recent work by Karapetyan et al. (2021) employed deep learning (DL) for vision-based navigation to cover underwater structures, such as shipwrecks, without dependence to state estimation. Other groups have recently showcased the potential of the Aqua2 (Sattar et al., 2008) AUV by presenting autonomous underwater navigation. DL-based approach was provided by Manderson et al. (2018) for obstacle avoidance by training using the decisions of a human pilot. Hong et al. (2021) used deep learning to classify objects into static and dynamic in conjunction with a potential field-based planner. Perception-aware underwater navigation by Manderson et al. (2020) provided an extension to their early work. Similarly to Manderson et al. (2018), this DL approach utilizes data produced by human operators driving the robot, which learns to avoid collisions with obstacles and to stay close to corals.

## 3. OVERVIEW

The key goal of the mapping framework is the ability to map the target structure in high resolution while at the same time maintaining an accurate overall situational awareness at all times. In order to achieve this goal the distal observer AUV(s) monitor the pose of the proximal observers together with the overall structure. The proximal observer explores near the structure and produces a dense 3D map. Key to the success of the mapping is to maintain an accurate estimate of the proximal observer's pose even when vision fails.

### 3.1 Cooperative Localization

Firstly, cooperative localization (Kurazume and Hirose, 2000; Rekleitis et al., 1998) needs to be solved for estimating the relative pose of Proximal Observer from Distal Observer. Previous work (Joshi et al., 2020), developed a deep learning framework to obtain the relative pose between the two AUVs from a single image and produced promising results for cooperative localization.

As motion capture systems used to obtain ground truth poses are difficult to set up underwater, we simulated a robot swimming by projecting the robot's 3D model over underwater images using Unreal Engine 4. Unfortunately, the produced images differed from the actual underwater images because of poor visibility and color loss underwater. We utilized CycleGAN (Zhu et al., 2017) to generate synthetic images that look like actual underwater footage in effect bridging the gap between simulation and actual underwater images. Then, we train the Convolutional Neural Network on the synthetic dataset and then evaluate on actual underwater footage.

We augmented the YOLOv3 (Redmon and Farhadi, 2018) network by passing the output of the standard backbone encoder to the object detection stream and the pose regression stream. The pose regression stream predicts the 2D projections of 8 corners of the robot's 3D model and the confidence score for each 2D keypoint prediction. The network architecture can be viewed as dividing an image into grids, and each grid votes for 2D keypoint projections and bounding box predictions. To focus on image areas belonging to the robot, predictions pertaining to grid cells

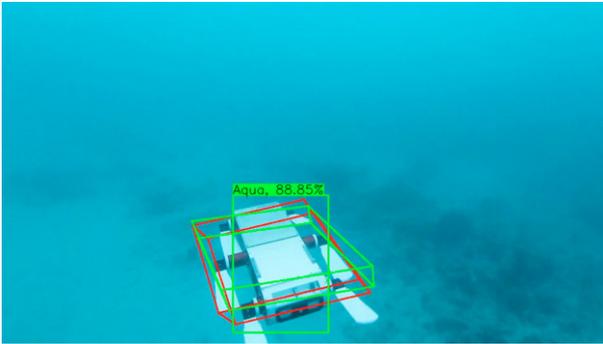


Fig. 2. An Aqua2 AUV detected from another Aqua2 while traveling over a coral reef, Barbados.

that fall outside the object detection box were discarded. For each 3D keypoint, 2D keypoint predictions with high confidence scores were selected to produce a set of 2D-to-3D correspondences. Then, RANSAC-based PnP (Lepetit et al., 2009) is used for robust relative pose estimation.

The deep learning framework has been evaluated in different types of environments – pool, ocean – and with multiple cameras, including several GoPro cameras, and an Aqua2 robot. The tests demonstrated its robustness to camera intrinsics changes, variations in underwater environment, and color calibration, Fig. 2. In addition, compared with state-of-the-art algorithms (Koreitem et al., 2018), the framework performed better in terms of translation and orientation accuracy on the pool dataset.

The discussed pose estimation network sometimes fails to predict the accurate pose when there is inaccurate detection of the bounding box. Thus, the robot’s correct relative pose with respect to the observer can only be obtained in case of accurate object detection and then interpolated for failed frames. In the future, we aim to fuse information from pressure-based depth sensor and IMU to detect disparities in the estimated pose across multiple frames and improve robustness and accuracy.

### 3.2 Robust State Estimation Underwater

Vision provides rich semantic information and through place recognition results in loop closures. Unfortunately, as demonstrated in recent work on comparing numerous open-source packages of visual and visual/inertial state estimation (Quattrini Li et al., 2016; Joshi et al., 2019), in an underwater environment there are frequent failures due to a variety of reasons. In contrast to above water scenarios, GPS based localization is impossible. In addition to the traditional difficulties of vision based localization, the underwater environment is prone to rapid changes in lighting conditions, limited visibility, low contrast, and color saturation with depth.

One important characteristic of vision underwater is that the range of the camera is limited to a few meters. Especially during operations around structures, AUVs often move in a way that sets the camera facing past the structure into open water. Even if there are obstacles more often than not they are not visible after a few meters. In contrast, aerial vehicles even if they look past a building, they would see the next building over or the ground. Figure 3 presents the view of the camera as an Aqua2



Fig. 3. Image acquired by an Aqua2 AUV during navigation over the Stavronikita shipwreck in Barbados. Features are detected only at the lower part of the image.

AUV explored over the bow of the Stavronikita shipwreck in Barbados. While the railing of the deck is visible at the bottom, most of the image is dominated by blue water, the sunlight penetrating from the surface makes the top part much brighter. Brisk features are marked on the bottom part of the image.

A robust switching estimator is utilized to ensure continuous pose estimates despite the loss of visual features. When the AUV is facing the structure, the SVIn2 (Rahman et al., 2019) VIO package is used. A health monitor checks the quality of the visual data based on the number and quality of features detected. When vision becomes unreliable, a model based estimator, termed primitive estimator (PE), is utilized. PE uses the water-depth, IMU, and the motion commands to estimate the motion of the AUV. When the PE starts, it is initialized at the last good pose produced by SVIn2. When the VIO recovers, the estimator switching back to SVIn2 which is initialized in the corresponding pose of the PE. Loop closure is applied to the combined trajectory ensuring that when an area is revisited the correction propagates back through both SVIn2 and the PE portions of the trajectory.

### 3.3 Proximal Observer Exploration

Different exploration and coverage patterns can be utilized to guide the proximal observer. Regardless the motion pattern, the proximal observer has to avoid obstacles and to keep the target structure in the field of view. The AquaVis framework, (Xanthidis et al., 2021), is utilized to produce an efficient trajectory selecting poses preferentially for keeping the prominent features of the structure in the camera’s view.

AquaVis is an active perception framework that extracts visual objectives online, moves towards a target position, and maximizes visibility of nearby visual objectives from a desired distance. For the purposes of this work, visual objectives were extracted by utilizing the centroids of areas with high density of features. This was achieved by clustering the 3D point-cloud, provided by state estimation, with DBSCAN (Ester et al., 1996) to drive the robot towards the shipwreck. In the future third-party object detection methods could be utilized to provide a more cognitive behavior. With the visual objectives as input, AquaVis

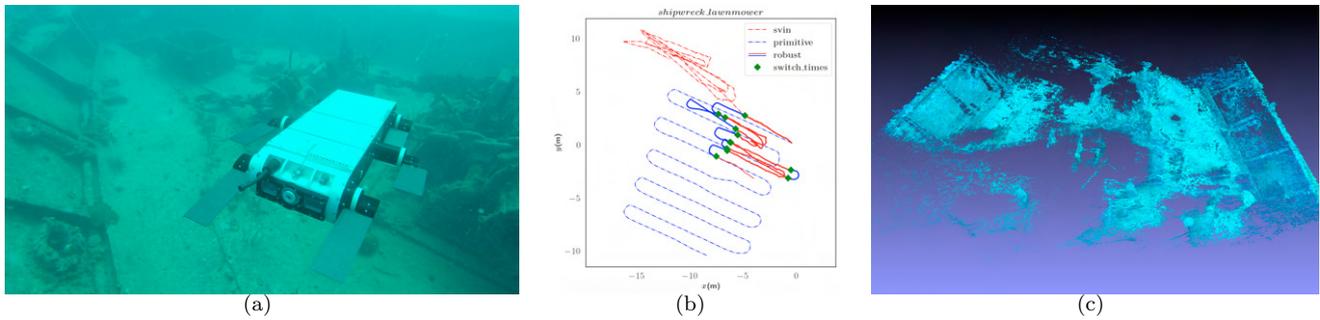


Fig. 4. A robust switching estimator utilized for an Aqua2 AUV navigating over the Stavronikita shipwreck (a). The Trajectory according to the primitive estimator; using the same algorithm as the robot controller PE believes the AUV performed a near perfect lawnmower pattern (blue dashed line); the trajectory according to SVIn2 (Rahman et al., 2019); due to tracking loss the VIO is way off the actual wreck (red dashed line); the trajectory from the switching estimator which utilized the robust parts of VIO (in red) switching to PE when tracking was lost (in blue) (b). Dense reconstruction utilizing the global optimization package COLMAP (c).

provides the waypoints to be traversed by the robot from a path-optimization process that combines constraints for minimizing path length, guarantying clearance, encouraging the robot to observe the extracted visual objectives, and enforcing the robot's kinematics, within a modified version of Trajopt (Schulman et al., 2014) for mobile robots moving in 3D.

### 3.4 Distal Observer Oversight

Visual servoing of the proximal observer is a fundamental strategy for the distal observer. Although, the presented technique does not take into consideration the visibility of the explored structure yet, future work will address this problem soon. To achieve the desired behavior the distal observer has to extract in real-time the position of the proximal observer, along with its future positions within an acceptable horizon. Please note that either minimal underwater communication between the robots, or a third-party tracking and motion prediction technique could satisfy such assumption. The proposed methodology for enabling cooperative localization presented in section 3.1 is exploring possibilities towards the latter direction.

Similarly to the proximal observer, this is also an instance of active perception, thus we proposed to utilize Aqua-Vis (Xanthis et al., 2021) too, with few fundamental changes. Firstly, the visual objectives are not extracted automatically by the point cloud, but they are communicated by the proximal observer, who shares its planned trajectory with the distal observer. Secondly, instead of using a collection of nearby visual objectives for each state during optimization, only a single one corresponding to a future position of the proximal observer was picked. Assuming known constant speeds for both the proximal and the distal observers, we pick only the visual objective that corresponds to the same time frame for each future state of the distal observer.

## 4. EXPERIMENTAL RESULTS

During an early deployment of the Aqua2 AUV over the bow of the Stavronikita shipwreck, Barbados, the AUV performed a fixed boustrophedon (Choset and Pignon, 1998) coverage pattern; see Fig. 4(a) for the AUV traveling

over the wreck. The open loop controller used a model based estimator fusing motion commands, IMU, and water depth sensor data resulting in a evenly spaced lawnmower patten; see the dashed-blue line in Fig. 4(b). The estimates of SVIn2 (Rahman et al., 2019) drifted off when the AUV was facing blue water; see the red dashed-line in Fig. 4(b). However, current pushed the AUV off the trajectory, utilizing a switching estimator resulted in accurate pose estimates; see the continues line, the color signifying which estimator was used, in Fig. 4(b). The collected images were then processed using the global optimization package COLMAP (Schonberger and Frahm, 2016) the resulting reconstruction can be seen in Fig. 4(c). Finally, Fig. 5 shows a instance of a simulated distal and proximal observer mapping a shipwreck utilizing the proposed technique in Gazebo (Koenig and Howard, 2004). For simulation purposes, odometry was provided by the simulator, while cameras were represented by simulated lidar sensors with the same field of view. The proximal observer (red) moves closer to the wreck while the distal observer (brown) follows keeping the proximal observer in view.

## 5. CONCLUSIONS

This paper presented a novel approach for the mapping and monitoring of underwater structures. The key concept is the utilization of two behaviours for AUVs, close inspection of the structure, enabling high definition observations, and distal observations that maintain an overall picture of the situation, including the rough shape of the structure and the pose of the proximal AUV with respect to the structure. In pursuit of this goal, a deep learning approach for estimating the relative pose between the two robots was presented together with a robust switching estimator using model based and VIO pose estimation. Motion strategies for the two AUVs are also discussed.

Extensions of this work include the distal observer being also the illumination carrier in environment without natural illumination, i.e. caves; see Fig. 6 where a diver with a VIO setup and a light plays the role of the distal observer. While traveling to and from the structure, the two vehicles will convoy (Shkurti et al., 2017) with the distal observer keeping the proximal in its field of view. For

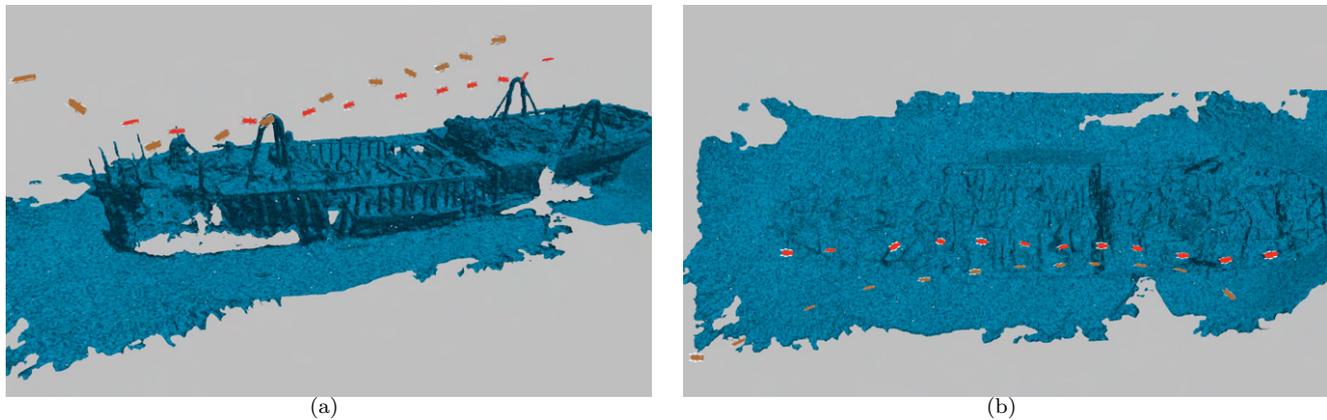


Fig. 5. Side and top view of a pair of robots exploring over the USS YP-389 shipwreck (NOAA Monitor National Marine Sanctuary, 1972).

reduced visibility environments, equipping the proximal observer with a light pattern will facilitate the relative pose estimate.



Fig. 6. Aqua2 AUV navigating inside the Ginnie Ballroom, FL. A diver plays the role of the distal observer with a GoPro camera (Joshi et al., 2022) and a light illuminating the AUV and the surroundings.

An additional advantage of the proposed approach is the capability of the distal observer to record the behaviour of the proximal observer over time, providing valuable insights on the exploration process. Deployments over shipwrecks and around aquaculture farms will provide historical records and information about maintenance and repair needs.

## REFERENCES

- Allotta, B., Costanzi, R., Ridolfi, A., Colombo, C., Bellavia, F., Fanfani, M., Pazzaglia, F., Salvetti, O., Moroni, D., Pascali, M.A., et al. (2015). The arrows project: adapting and developing robotics technologies for underwater archaeology. *IFAC-PapersOnLine*, 48(2).
- Chaves, S.M., Kim, A., Galceran, E., and Eustice, R.M. (2016). Opportunistic sampling-based active visual slam for underwater inspection. *Autonomous Robots*, 40(7).
- Choset, H. and Pignon, P. (1998). Coverage path planning: The boustrophedon cellular decomposition. In *Field and service robotics*, 203–209. Springer.
- Demesticha, S., Skarlatos, D., and Neophytou, A. (2014). The 4th-century bc shipwreck at mazotos, cyprus: new techniques and methodologies in the 3d mapping of shipwreck excavations. *Journal of Field Archaeology*, 39(2), 134–150.
- Dieudonné, Y., Labbani-Igbida, O., and Petit, F. (2010). Deterministic robot-network localization is hard. *IEEE Trans. Robot.*, 26(2), 331–339.
- Ester, M., Kriegel, H.P., Sander, J., Xu, X., et al. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. In *kdd*, volume 96.
- Eustice, R.M., Singh, H., Leonard, J.J., and Walter, M.R. (2006). Visually mapping the RMS Titanic: Conservative covariance estimates for SLAM information filters. *The Int. J. of Robotics Research*, 25(12), 1223–1242.
- Frolov, S., Garau, B., and Bellingham, J. (2014). Can we do better than the grid survey: Optimal synoptic surveys in presence of variable uncertainty and decorrelation scales. *Journal of Geophysical Research: Oceans*, 119(8).
- Hong, J., de Langis, K., Wyeth, C., Walaszek, C., and Sattar, J. (2021). Semantically-aware strategies for stereovisual robotic obstacle avoidance. In *2021 International Conference on Robotics and Automation (ICRA)*. IEEE.
- Huang, G.P., Trawny, N., Mourikis, A.I., and Roumeliotis, S.I. (2009). On the consistency of multi-robot cooperative localization. In *Robotics: Science and Systems*.
- Joshi, B., Modasshir, M., Manderson, T., Damron, H., Xanthidis, M., Quattrini Li, A., Rekleitis, I., and Dudek, G. (2020). DeepURL: Deep pose estimation framework for underwater relative localization. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1777–1784. Las Vegas, NV.
- Joshi, B., Rahman, S., Kalaitzakis, M., Cain, B., Johnson, J., Xanthidis, M., Karapetyan, N., Hernandez, A., Quattrini Li, A., Vitzilaios, N., and Rekleitis, I. (2019). Experimental Comparison of Open Source Visual-Inertial-Based State Estimation Algorithms in the Underwater Domain. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 7221–7227. Macau.
- Joshi, B., Xanthidis, M., Rahman, S., and Rekleitis, I. (2022). High definition, inexpensive, underwater mapping. In *IEEE International Conference on Robotics and Automation (ICRA)*, accepted. Philadelphia, PA, USA.
- Karapetyan, N., Johnson, J., and Rekleitis, I. (2021). Human diver-inspired visual navigation: Towards coverage path planning of shipwrecks. *Marine Technology Society Journal*, "Best of OCEANS 2020", 55(4), 24–32.

- Koenig, N. and Howard, A. (2004). Design and use paradigms for gazebo, an open-source multi-robot simulator. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*(IEEE Cat. No. 04CH37566), volume 3, 2149–2154. IEEE.
- Koreitem, K., Li, J., Karp, I., Manderson, T., Shkurti, F., and Dudek, G. (2018). Synthetically trained 3d visual tracker of underwater vehicles. In *MTS/IEEE OCEANS*. Charleston, SC, USA.
- Kurazume, R. and Hirose, S. (2000). An experimental study of a cooperative positioning system. *Autonomous Robots*, 8, 43–52.
- Kurazume, R., Nagata, S., and Hirose, S. (1994). Cooperative positioning with multiple robots. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 2, 1250–1257.
- Lepetit, V., Moreno-Noguer, F., and Fua, P. (2009). EPnP: An accurate O(n) solution to the PnP problem. *Int. Journal of Comp. Vision*.
- Leung, K.Y.K., Barfoot, T.D., and Liu, H.H.T. (2010). Decentralized localization of sparsely-communicating robot networks: A centralized-equivalent approach. *IEEE Trans. Robot.*, 26, 62–77.
- Manderson, T., Higuera, J.C.G., Cheng, R., and Dudek, G. (2018). Vision-based autonomous underwater swimming in dense coral for combined collision avoidance and target selection. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1885–1891.
- Manderson, T., Higuera, J.C.G., Wapnick, S., Tremblay, J.F., Shkurti, F., Meger, D., and Dudek, G. (2020). Vision-based goal-conditioned policies for underwater navigation in the presence of obstacles. *Robotics: Science and Systems XVI*.
- Martinelli, A. and Renzaglia, A. (2017). Cooperative visual-inertial sensor fusion: Fundamental equations. In *Int. Symp. Multi-Robot and Multi-Agent Systems*.
- Meger, D., Higuera, J.C.G., Xu, A., Giguere, P., and Dudek, G. (2015). Learning legged swimming gaits from experience. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2332–2338. IEEE.
- Menna, F., Agrafiotis, P., and Georgopoulos, A. (2018). State of the art and applications in archaeological underwater 3d recording and mapping. *Journal of Cultural Heritage*, 33, 231–248.
- Mourikis, A. and Roumeliotis, S. (2006). Performance analysis of multirobot cooperative localization. *IEEE Trans. Robot.*, 22(4), 666–681.
- NOAA Monitor National Marine Sanctuary (1972). ex-uss yp-389. <https://monitor.noaa.gov/shipwrecks/yp-389.html>. Accessed 2022-04-21.
- Nornes, S.M., Ludvigsen, M., Ødegard, Ø., and Sørensen, A.J. (2015). Underwater photogrammetric mapping of an intact standing steel wreck with rovs. *IFAC-PapersOnLine*, 48(2), 206–211.
- Quattrini Li, A., Coskun, A., Doherty, S.M., Ghasemlou, S., Jagtap, A.S., Modasshir, M., Rahman, S., Singh, A., Xanthidis, M., O’Kane, J.M., and Rekleitis, I. (2016). Experimental comparison of open source vision based state estimation algorithms. In *International Symposium of Experimental Robotics (ISER)*. Tokyo, Japan.
- Rahman, S., Quattrini Li, A., and Rekleitis, I. (2018). Sonar Visual Inertial SLAM of Underwater Structures. In *IEEE International Conference on Robotics and Automation*, 5190–5196. Brisbane, Australia.
- Rahman, S., Quattrini Li, A., and Rekleitis, I. (2019). An Underwater SLAM System using Sonar, Visual, Inertial, and Depth Sensor. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1861–1868. Macau, (IROS ICROS Best Application Paper Award. Finalist).
- Redmon, J. and Farhadi, A. (2018). YOLOv3: An incremental improvement.
- Rekleitis, I., Dudek, G., and Milios, E.E. (1998). On multiagent exploration. In *Vision Interface*, 455–461. Vancouver, BC, Canada.
- Roumeliotis, S.I. and Rekleitis, I. (2004). Propagation of uncertainty in cooperative multirobot localization: Analysis and experimental results. *Autonomous Robots*, 17(1), 41–54.
- Sattar, J., Dudek, G., Chiu, O., Rekleitis, I., Giguere, P., Mills, A., Plamondon, N., Prahacs, C., Girdhar, Y., Nahon, M., and Lobos, J.P. (2008). Enabling autonomous capabilities in underwater robotics. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 3628 – 3634. Nice, France.
- Schonberger, J.L. and Frahm, J.M. (2016). Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Schulman, J., Duan, Y., Ho, J., Lee, A., Awwal, I., Bradlow, H., Pan, J., Patil, S., Goldberg, K., and Abbeel, P. (2014). Motion planning with sequential convex optimization and convex collision checking. *The International Journal of Robotics Research*, 33(9).
- Shkurti, F., Chang, W.D., Henderson, P., Islam, M.J., Higuera, J.C.G., Li, J., Manderson, T., Xu, A., Dudek, G., and Sattar, J. (2017). Underwater multi-robot convoying using visual tracking by detection. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 4189–4196. IEEE.
- Williams, S.B., Pizarro, O., and Foley, B. (2016). Return to antikythera: Multi-session slam based auv mapping of a first century bc wreck site. In *Field and Service Robotics*, 45–59. Springer.
- Xanthidis, M., Kalaitzakis, M., Karapetyan, N., Johnson, J., Vitzilaios, N., O’Kane, J., and Rekleitis, I. (2021). Aquavis: A perception-aware autonomous navigation framework for underwater vehicles. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 5387–5394. Prague, Czech Republic.
- Xanthidis, M., Karapetyan, N., Damron, H., Rahman, S., Johnson, J., O’Connell, A., O’Kane, J., and Rekleitis, I. (2020). Navigation in the presence of obstacles for an agile autonomous underwater vehicle. In *IEEE International Conference on Robotics and Automation*.
- Zhou, X. and Roumeliotis, S. (2012). Determining 3D relative transformations for any combination of range and bearing measurements. *IEEE Trans. Robot.*, 29(2).
- Zhu, J.Y., Park, T., Isola, P., and Efros, A.A. (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. In *International Conference on Computer Vision*.