

AquaVis: A Perception-Aware Autonomous Navigation Framework for Underwater Vehicles

Marios Xanthidis,¹ Michail Kalaitzakis,² Nare Karapetyan,¹ James Johnson,¹
Nikolaos Vitzilaios,² Jason M. O’Kane,¹ and Ioannis Rekleitis¹

Abstract—Visual monitoring operations underwater require both observing the objects of interest in close-proximity, and tracking the few feature-rich areas necessary for state estimation. This paper introduces the first navigation framework, called AquaVis, that produces on-line visibility-aware motion plans that enable Autonomous Underwater Vehicles (AUVs) to track multiple visual objectives with an arbitrary camera configuration in real-time. Using the proposed pipeline, AUVs can efficiently move in 3D, reach their goals while avoiding obstacles safely, and maximizing the visibility of multiple objectives along the path within a specified proximity. The method is sufficiently fast to be executed in real-time and is suitable for single or multiple camera configurations. Experimental results show the significant improvement on tracking multiple automatically-extracted points of interest, with low computational overhead and fast re-planning times.

Accompanying short video: <https://youtu.be/JKO.bbrIZyU>

I. INTRODUCTION

Autonomous underwater monitoring and navigation can be very hard for a variety of reasons. For example, the robot must move safely, avoiding obstacles and staying at depth. Planning and executing such motions can be particularly challenging for an AUV moving in three dimensions, with complex dynamics that have not been adequately modeled. Our previous work introduced AquaNav [1], which robustly solved these problems for very challenging environments, in simulation, in-pool, and open water conditions. Additionally, as a planning problem, visual monitoring of unknown underwater 3D environments in real-time is very challenging, due to the dimensionality of the problem, and the constraints introduced by the limited cameras’ range and field of view.

Furthermore, though visual data is generally utilized for state estimation, underwater environments tend to produce very noisy images due to lack of color saturation, insufficient illumination, and color attenuation. Moreover, good visual features are often concentrated on few nearby objects; while much of the visible terrain has few features. As a result, state-of-the-art methods fail to provide robust state estimation for the robot [2], [3], although previous work has addressed this problem by providing a very capable SLAM framework

¹ M. Xanthidis, N. Karapetyan, A. Johnson, J. M. O’Kane, and I. Rekleitis are with the Department of Computer Science and Engineering, University of South Carolina, Columbia, SC, USA. [mariosx,nare,jvj1]@email.sc.edu, [jokane,yiannisr]@cse.sc.edu

²M. Kalaitzakis and N. Vitzilaios are with the Department of Mechanical Engineering, University of South Carolina, Columbia, SC, USA. michailk@email.sc.edu, vitzilaios@sc.edu

This work was made possible through the generous support of National Science Foundation grants (NSF 1659514, 1849291, 1943205, 2024741).




Fig. 1. An environment with obstacles (grey) and feature-rich visual objectives indicated with stars. (a) AquaNav, considers only avoiding obstacles and minimizing the path length. (b) AquaVis, the method introduced here, navigates the robot safely by avoiding obstacles, while at the same time observing nearby visual objectives.

called SVIn [4], [5], under the assumption that an adequate number of high-quality features are visible throughout the path. However, even the most capable vision-based SLAM systems have trouble tracking the state when facing featureless homogeneous surfaces, turbidity, or the open (blue) water conditions that dominate the underwater domain.

For robust underwater navigation it is highly important to combine perception and motion planning, in order to avoid the obstacles, but also keep feature-rich objects in the cameras’ field of view. Bringing perception and motion planning closer not only assists state-estimation, but also produces trajectories that track and monitor points of interest, such as fish, corals, and structures. Such behavior is preferred for exploration and monitoring strategies that should collect diverse and meaningful-to-humans information [6], [7]. To this end, we propose a novel framework called AquaVis, whose objective is to generate motions that enable the robot not only to move efficiently and avoid obstacles safely, but also to observe areas of interest, that could be extracted automatically. The difference is illustrated in Figure 1.

This is achieved —utilizing the flexibility of the AquaNav framework [1]— by introducing two novel cost functions in the optimization process during planning, to direct the robot to observe specific points of interest while avoiding

the obstacles and respecting the kinematics of the robot. An analysis of the produced trajectories demonstrate object tracking with the desired proximity and safe navigation around obstacles.

The specific contributions of this paper are the following:

- 1) A novel and robust framework, called AquaVis, for autonomous 3D navigation for the Aqua2 robot [8], surpassing AquaNav’s capabilities by improving the perception capabilities.
- 2) A novel formulation of perception-aware navigation for mobile robots with an arbitrary number of cameras, tracking multiple visual objectives, moving in 3D. We also show how visual objectives could be extracted automatically, from perceived point-clouds, to assist state estimation.

II. RELATED WORK

The problem of actively planning and executing motions to improve state estimation performance, also known as Active SLAM, was first introduced by Feder *et al.* [9] in late 90’s, in the context of exploration. Stachniss and Burgard [10] provided a method that improved localization using SLAM, by attempting loop-closing.

Makarenko *et al.* [11] employed a laser, extracted landmarks that were used with an Extended Kalman Filter, and proposed a method that could be parameterized to trade-off exploring new areas with uncertainty. Martinez *et al.* [12] reduced pose and map estimates with Gaussian Processes. The work of Rekleitis introduced an exploration versus exploitation framework to reduce uncertainty for a single robot by visiting previously mapped areas for single [13] and multi-robot systems [14]. Zhang *et al.* [15], [16] employed hybrid metric-topological maps to reduce uncertainty. All these early works considered only the 2D case.

More recent studies have expanded the problem from 2D to 3D, with the main platform considered being quadrotors, although a few studies utilizing manipulators also exist [17]. The work of Forster *et al.* [18] provided a method to minimize uncertainty of a dense 3D reconstruction, but it was based on a direct method that has weak performance underwater, and mostly fly-over motions were performed without robust obstacle avoidance. Penin [19] introduced a framework for producing trajectories taking into account the field of view limitations of the camera, but it was restricted to tracking only 4 points in close proximity to each other, no obstacles were considered, and no real-time performance.

The work of Spica *et al.* [20] combined visual servoing with Structure from Motion, but their primary focus was mapping and their method did not consider obstacles and operations in cluttered environments. Constante *et al.* [21] proposed a photometric method to drive the robot close to regions with rich texture, but as with Forster *et al.* [18], direct methods do not perform well underwater and the motions were constrained to fly-overs and near-hovering.

Sheckells *et al.* [22] provided an optimal technique for visual servoing with no obstacle avoidance and only one visual objective was considered for the duration of the

trajectory. Additionally, the work of Nageli *et al.* [23], [24] focuses on visual-objective tracking, rather than achieving a navigation goal with robust localization, and the potential field method applied for obstacle avoidance could result in a local minimum in cluttered environments.

Other related studies considered only one visual objective [25]–[29] or did not consider obstacle avoidance [21], [25], [30]–[32]. It is worth noting that Spasojevic *et al.* [31] were indeed able to track a set of landmarks, but with the constraint that they should always be tracked; in this work the robot needs to choose which objective(s) should be tracked from each position along the path.

Greef *et al.* [33] utilized a camera mounted on a gimbal. However, since that method was based on a teach and repeat approach, it is not applicable for unexplored environments. Given the camera configuration and the kinematics of the Aqua2 [8], neither the method of Zhou *et al.* [34], nor of Murali *et al.* [30] which allow lateral motions and free on-the-spot yaw rotations are suitable. Some techniques [26], [29] consider only one target, but they resulted in low-level controllers. A very recent result by Zhang and Scaramuzza [35] proposed a new topological model for map representation that could be used for guaranteeing uncertainty reduction in the entire map, but a computationally expensive offline computation on a known map is needed before planning, limiting the scope of online applications.

In the underwater domain, in the context of coverage, Frolov *et al.* [36] proposed a motion planning framework for reducing map uncertainty by revisiting areas of high uncertainty, while Chaves *et al.* [37] utilized loop closures to reduce uncertainty. Work by Karapetyan *et al.* [38] used vision based navigation to perform coverage of a shipwreck with no state estimation.

Recent work from different groups has emphasized the potential of the Aqua2 [8] platform by providing effective real-time underwater navigation methods. Manderson *et al.* [39] provided a deep learning-based approach for collision avoidance by training upon the decisions of a human operator, Hong *et al.* [40] utilized deep learning for classifying obstacles to static and dynamic on top of a potential field-based planner for obstacle avoidance, while Xanthidis *et al.* [1] produced a very capable model-based navigation framework using path-optimization.

Finally, on the front of perception-aware underwater navigation, Manderson *et al.* [41] provided an extension to their previous work. Similarly to [39], this deep-learning technique was based on fitting on data collected by a human operator controlling the robot. The robot was taught to stay close to corals, and avoid collisions with corals and rocks. Despite the effectiveness of this technique, the proposed solution is (a) unable to fully exploit the kinematic abilities of the robotic platform the way AquaNav does, because it does not consider roll motions and is limited to human intuition, (b) is naturally constrained to navigate only in similar environments (coral reefs), and (c) the motion commands follow a very reactive behavior and a short decision window that was compensated for by following predefined local goals.

On the other hand, AquaVis produces locally near-optimal motions for avoiding the obstacles, with no reliance on a potentially error-prone human training process. It also produces efficient trajectories for safe navigation in cluttered environments, similar to AquaNav. More importantly, since it operates on point-clouds, localization could be maintained with any kind of structures with rich texture, without the limitations dictated by a training dataset. Moreover, it is able to incorporate third-party object recognition modules for monitoring objects of interest, without the need of the time and resource intensive training on the motion planning module.

III. OVERVIEW

The goal of AquaVis is the safe navigation of an underwater robot, such as the Aqua2, moving freely through a cluttered three-dimensional environment. Such navigation should be accomplished while maintaining visibility of sparse visual objectives along its path. Specifically, the paths executed by the robot should reach a specified goal while keeping the path length short, avoid obstacles, and maximize the number of states along that path from which at least one visual objective is visible.

The proposed navigation framework, operating in unknown environments, selects from among the observed areas the ones that satisfy the visibility objectives and guides the trajectory towards the most appropriate ones. For example, in the presence of feature rich clusters, the trajectory is morphed to keep these clusters in the field of view. Barring any prior or sensed information about the environment, the main driver is motion towards the destination while minimizing the distance travelled. Although our primary focus lies on underwater robots, there are no explicit assumptions or limitations introduced that prohibit applications on other platforms and domains.

A. AquaNav Overview

To achieve the desired behavior of AquaVis, we utilized the robust navigation architecture of AquaNav [1], and extended its core planning module, utilizing Trajopt, to use the location of visual features as a constraint, Figure 2. In short, AquaNav is a waypoint navigation system, capable of real-time replanning and execution of trajectories with a guaranteed clearance, to ensure safety in the challenging and unpredictable underwater conditions. The AquaNav system is robust enough to enable real-time replanning, efficient and safe navigation in unknown environments, and is tested in real open-water conditions. Trajopt is the primary path-optimization planner that ensures the above guarantees, and is assisted by a sampling-based warm-starting method, to overcome local minimum challenges. This sort of path-optimization based approach not only generates high quality solutions rapidly, but also offers adequate flexibility, enabling modifications in the form of novel constraints and cost functions. For the complete description of the AquaNav framework please refer to Xanthidis *et al.* [1].

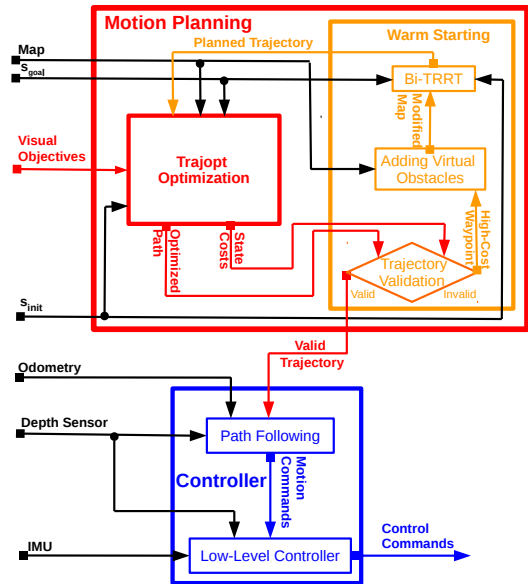


Fig. 2. System architecture of AquaVis, which is based on AquaNav. AquaVis alters the core planning component by incorporating visual objectives, shown with red, while modules for warm-starting, shown with orange, and path following, shown with blue, are kept the same.




Fig. 3. Example of the visibility formulation used in Equation 1. The visibility manifold F_s for 2 cameras mounted on the robot is shown in light green. Visual objectives v_1 , v_2 , and v_3 are indicated with stars. Only v_2 is visible because it is inside F_s , while v_1 and v_3 are not observable from the robot's current state s .

B. AquaVis Objective

The robot's state s describes its position and orientation in some fixed coordinate frame. The robot is equipped with one or more cameras, such that from state s , a region $F_s \subseteq \mathbb{R}^3$ is visible from at least one of the cameras. Let V denote a finite set of visual objectives. For a given state s , each visual objective $v \in V$ may be visible (i.e. $v \in F_s$) or not ($v \notin F_s$); see Figure 3 where visual objective v_2 is visible by the front camera. Additionally, let the continuous path of the robot be approximated by a sequence of consecutive states s_1, \dots, s_n . Note that as n increases, we can approximate the robot's continuous path with arbitrary precision. We quantify the path's success in maintaining visibility of the visual objects via the following function:

$$M(s_1, \dots, s_n) = \frac{|\{s_i \mid F_{s_i} \cap V \neq \emptyset\}|}{n} \quad (1)$$

This function provides the fraction of the states in the path that observe at least one objective. It reaches 1 if all the states are able to observe at least one visual objective, or 0 if no visual objectives were observed during traversing the entire path. Thus, the objective of AquaVis is to minimize

the path length, avoid obstacles, and maximize Equation 1.

IV. PROPOSED APPROACH

This section describes the enhancements of AquaVis upon the AquaNav pipeline. These enhancements consist of ways to automatically extract the visual objectives, modify the planning process to accommodate them, and ensure the satisfaction of the kinematic constraints.

A. Extracting Visual Objectives

With respect to the AquaVis pipeline (Figure 2), the visual objectives are considered as an input in the form of a list of 3D points. These visual objectives are either user-defined, or automatically extracted online. For example, known methods that detect corals [42]–[44], or other Aqua2 robots [45] and extract the 3D positions of those features could be employed for application specific purposes, such as environmental monitoring or multi-robot exploration.

Visual objectives could be extracted automatically to assist underwater state estimation, by utilizing the output of most SLAM techniques. In particular, AquaNav employs the robot to navigate through an unknown environment, using a state estimation package, such as SVin2 [5], that outputs both the odometry and a representation of the sensed environment as a 3D point cloud. Thus, the raw point-cloud could be processed to extract visual objectives with high density of features, then these visual objectives could assist the odometry as landmarks. Such an approach is a necessity in the underwater domain, which is notoriously challenging for vision-based state estimation [2], [3], in part because the quality of the features is often low, and their spatial distribution uneven, with most features concentrated in only a few places.

We propose extracting visual objectives from a point-cloud by treating the problem as density-based clustering. DBSCAN [46] is applied on the 3D point cloud to detect clusters with high density and then the centroids of these clusters are chosen as the visual objectives. DBSCAN has a minimal number of parameters: the minimum number of samples per cluster and the minimum proximity. The operator decides the quality of the objectives to be tracked, in terms of both number and density of good features [47].

In each planning cycle, the above preprocessing step produces the visual objectives used during planning. Though, not keeping past information of previously detected clusters, could result to a highly sub-optimal reactive behavior. Thus, the set V contains a maximum of m computed visual objectives, in order to ensure real-time planning, and to avoid excessive computation from an ever increasing number of visual objectives. Initially, the visual objectives are added to the list until $|V| = m$. Then, any new measurement replaces the closest one if they are in close proximity by updating the center of the cluster, or in any other case, the oldest one to favor locality and computational efficiency.

B. Motion Planning Modifications

AquaVis modifies the path optimization element of the AquaNav framework, which is built upon the optimization-




Fig. 4. Top (a) and side (b) views of a state using the novel constraints during optimization. The blue square indicates a visual objective, and the red circle marks the next waypoint. Minimizing d_{obj} will result on the robot observing the objective, while minimizing d_{align} will result on the robot to be consistent with the kinematics assumed during path execution and planning.

based package Trajopt. A brief review of the original Trajopt formulation is discussed next.

1) *Original Trajopt formulation:* Trajopt attempts to minimize the function

$$f(S) = \min_S \sum_{i=1}^{n-1} \|s_{i+1} - s_i\|, \quad (2)$$

where $S = \langle s_0, s_1, \dots, s_n \rangle$ the sequence of n states of the robot considered during optimization. $f(S)$ is the sum of squared displacements, which minimizes path length.

Collision constraints are enforced for every state $s \in S$:

$$h(s) = \sum_{o \in O} |d_{safe} - sd(PC_s, o)| \quad (3)$$

where O is the set of obstacles, PC_s is the 3D geometry of the robot in state s , and sd represents the minimum Euclidean distance to separate two 3D convex objects. More details about sd appear in [48].

The above constraint, given successful convergence, guarantees that each waypoint on the path will maintain distance at least d_{safe} from the closest obstacle, but has no guarantees on the transitions between waypoints. To enforce continuous time safety, instead of Equation 3, the following function is applied for each pair s_{i-1}, s_i of consecutive states:

$$H(s_i, s_{i+1}) = \sum_{o \in O} \left| d_{safe} - sd(L_{s_{i-1}}^{s_i}, o) \right|, \quad (4)$$

in which $L_{s_{i-1}}^{s_i}$ is the convex hull of $PC_{s_{i-1}} \cup PC_{s_i}$.

Two additional cost functions are introduced. The first new cost function incentivizes the robot to view visual objectives, whereas the second one forces the path to self-correct and maintain the kinematic constraints to visit each waypoint assumed by the path follower. The functions are described below, and an example that outlines these novel cost functions is shown in Figure 4.

2) *Visibility Constraints:* The visibility constraint is intended to direct the robot to observe a known set of visual objectives. The core idea is to project a set of points F_s^\sim in front of the robot's cameras to approximate F_s and then attempt to minimize the distance d_{obj} between the closest visual objective to the closest projected point of F_s^\sim . By minimizing the above distance to zero, the robot is guaranteed to track at least one visual objective in state




Fig. 5. Different perspectives of the projected points of the F_s^{\sim} visibility set approximating the F_s visibility manifold corresponding to the front camera. s . Figure 5 shows an example of the above concept for the front camera of the robot.

Given a set of objectives V and a visibility set F_s^{\sim} for the state s , the general form of the proposed constraint applied to each state is:

$$\text{Vis}(s) = \min_{v \in V} \min_{f \in F_s^{\sim}} \|f - v\| \quad (5)$$

Trajopt was modified to utilize the above cost function to minimize the distance d_{obj} between the projected desired point and the nearest visual objective. So, given successful convergence, at least one visual objective will be visible at a desired direction and distance for each state.

It should be noted that there is an important trade-off between approximating F_s accurately and real-time performance. During the optimization process, for each state, the distance between each visual objective and each projected point f , $f \in F_s^{\sim}$ is calculated, resulting potentially to slow re-planning. Real-time performance requires selecting a small set F_s^{\sim} , thus further relaxing path optimality.

Acceptable performance that trades path optimality and smooth transitions for low computational cost, and fast re-planning, could even be achieved with a set F_s^{\sim} formed by a single point in front of the camera at a desired distance. Let T_w^s be the transformation from the world reference frame to the robot's local reference frame at state s , T_r^c is the transformation from the robot's local reference frame at state s to the camera c , C is the set of cameras, and d_{vis} is the desired distance of observing a visual objective, then:

$$F_s^{\sim} = \bigcup_{c \in C} \left\{ T_w^s T_r^c [d_{vis} \ 0 \ 0 \ 1]^T \right\} \quad (6)$$

This formulation guarantees that during the planned trajectory, multiple objectives could be observed from multiple cameras, contrary to the works discussed previously. Also with the simplification described in Equation 6, AquaVis achieves real-time behavior with no significant added delays to the AquaNav replanning baseline. Moreover, although path-optimality was sacrificed for real-time performance, in our experiments it is shown that efficient and smooth paths were produced thanks to the greedy nature of the formulation combined with fast replanning.

3) *Kinematic Constraints*: The only objective of AquaNav was to minimize the path-length in terms of both translation and rotation. So the produced plans could be

executed directly by a way-point follower, since unnecessary aggressive rotations are not expected. Similarly, the paths produced by using only the visibility constraints could be directly executed by a holonomic robot, enabling it to move and rotate in way to track the necessary visual objectives. However, the Aqua2 vehicle is not a holonomic robot, thus by using only the visual constraints, many motions requiring lateral translation could not be executed by the path follower. To ensure that the resulting trajectory for observing the visibility objectives satisfies the kinematic constraints of the vehicle a second constraint is introduced. The path follower module in the AquaNav pipeline accepts the 3D coordinates that need to be reached by the robot, along with a constant desired roll orientation during the motion. Ideally, the robot would move along the straight line segments connecting successive waypoints. Thus, the robot after achieving the waypoint p_{s_i} should maintain an orientation pointing directly to the next waypoint $p_{s_{i+1}}$, with p_{s_i} indicating the 3D coordinates of state s_i .

The cost function utilized by AquaNav for aligning the robot properly during planning is similar to the cost function described for the visibility constraints: A point is projected in front of the robot at a specific distance and the distance d_{align} of this point to the next waypoint is minimized. More precisely, let $S = [s_1, s_2, \dots, s_{n-1}, s_n]$ be the trajectory to be optimized, the cost function applied for each state s_i is

$$A(s_i) = \left\| T_w^{s_i} \begin{bmatrix} av(S) - \epsilon \\ 0 \\ 1 \end{bmatrix} - \begin{bmatrix} p_{s_{i+1}}^T \\ 1 \end{bmatrix} \right\|, \quad (7)$$

where p_{s_i} is the coordinates of state s_i in the world frame, and $av(S)$ is the average length of the S trajectory given by:

$$av(S) = \frac{\sum_{i=2}^n \left\| p_{s_i}^T - p_{s_{i-1}}^T \right\|}{n-1} \quad (8)$$

The first element of the first vector, similar to Equation 6, is the distance the point will be projected, in this case forward. The distance is calculated as the average distance between two consecutive states reduced by a small positive value ϵ . So the projected point is adjusted automatically to be the same for every state to encourage consistency, while the ϵ factor is used to encourage shorter path lengths.

V. EXPERIMENTAL RESULTS

The performance of AquaVis was validated in simulation. Its flexibility in controlling the trade-off between path-length, tracking visual objectives, and satisfying the way-point navigation kinematics was explored. In our experiments the desired clearance for obstacle avoidance, similarly to [1], was set to 0.6 m, the desired visibility distance d_{vis} to 1.0 m, and the linear velocity of the robot to 0.4 m/s.

Additionally, the original camera configuration of Aqua2 — our target system — was used. Aqua2 leverages a forward-looking stereo camera system for state estimation. That system has a field of view 90° vertically and 120° horizontally. The forward-looking camera is tilted downwards by 40° to further assist state estimation.

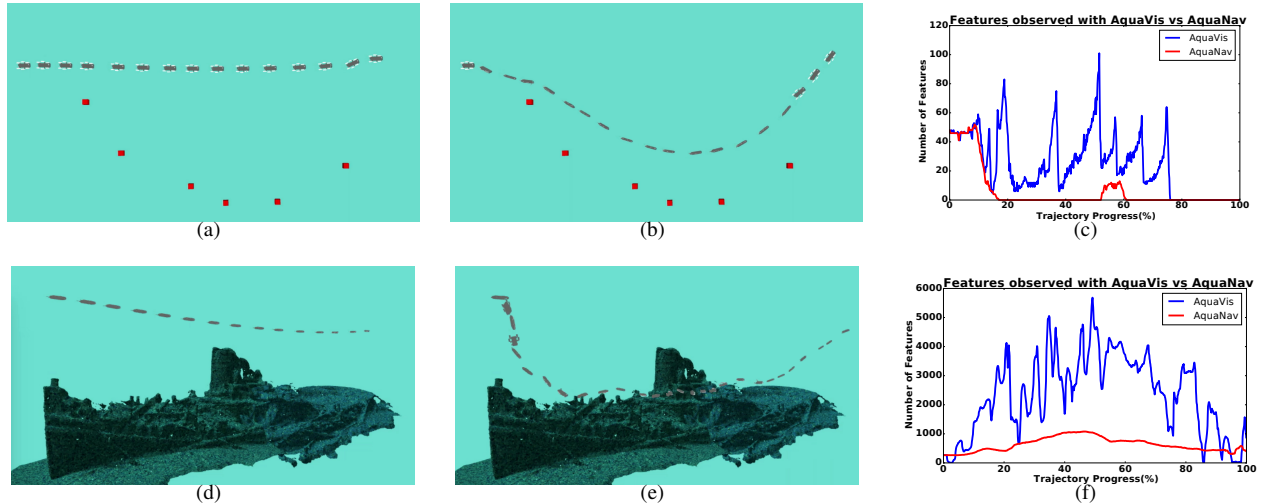


Fig. 6. The results for the Pilars environment are shown on the top, and for the Shipwreck at the bottom row. The trajectories produced by AquaNav are shown in (a) and (d), and for AquaVis at (b) and (e). The features observed for both methods are shown in (c) and (f).

A. Simulations

We simulate the detected features from the stereo VIO with a lidar sensor that returns 3D clusters of features in select obstacles. In the real system, fewer features will be detected but that does not affect the planning process negatively, instead the process of extracting visual objectives is expected to be executed faster. The simulated lidar has the same field of view with the Aqua2 front cameras, a resolution of 100×75 , and range of 6 m, to represent the expected turbidity of the underwater domain. To extract visual objectives automatically using DBSCAN [46], the maximum distance between features was set to 0.2 m with a minimum number of 5 features per cluster. A maximum set of 15 visual objectives was maintained, with new visual objectives replacing the closest of the old ones that were in a distance less than 0.5 m, or the oldest in the set. AquaVis is tested online against AquaNav in 2 different environments, the Pilars, and the Shipwreck shown in Figure 6.

The Pilars Environment, shown in Figure 6(a-c), is intended to test AquaVis in an environment where feature rich areas are distributed sparsely in the environment; a top down view is presented. AquaNav, by optimizing path length, moves on a straight line, disregarding the features, which are essential for localization. AquaViz, in contrast, reaches the same goal while passing in proximity and observing the feature rich areas (red cubes). The plot in Figure 6(c) confirms our expectations: AquaNav cannot observe any features for the majority of the time, whereas AquaVis consistently tracked enough features. It is worth noting, that AquaVis introduced a 90° roll to bring the visual objectives of the pilars in the field of view. Moreover, AquaVis maintained tracking for the first 75% of the trajectory that visual objectives could be observed, and lost track at the last 25% where no visual objectives were present that could be observed with a forward looking camera.

Similarly, in the Shipwreck environment, shown in Figure 6-(d-f), AquaVis was able to observe consistently more




Fig. 7. Sensitivity Analysis environments. Here, obstacles are colored red and the visual objectives are denoted with blue squares

features than AquaNav, excluding ascent and descent, which is expected given the kinematics of Aqua2. Also the robot not only oriented itself to track most of the shipwreck but also created the desired proximity, indicating potential use for mapping purposes. On the other hand, AquaNav moved in a straight line, unaware of the feature rich areas, and tracked only a small portion of the features tracked by AquaVis.

B. Sensitivity Analysis

To evaluate the flexibility of AquaVis to control the trade-off between the path length, kinematics, and tracking the visual objectives, the system was tested on five simulated environments of gradually increasing complexity and difficulty. The environments considered are shown in Figure 7 plus one more environment with objectives as Figure 7(a) but without any obstacles. The initial setup for all environments has Aqua in a straight trajectory 12m long.

There are six different weights in the cost function of AquaVis. Two weights for the initial TrajOpt formulation controlling the translation and rotation change between states, an obstacle avoidance weight, and then the two weights introduced in AquaNav to ensure the traversability

of the trajectory and finally the visual objectives weight introduced in this work. Ten different weight sets were used to find their effect on the total path length and the tracking of the visual objectives.

For all environments and all sets of weights, the total trajectory length, the average distance from the nearest visual objective (d_{obj}) and average alignment distance (d_{align}) for each state s were measured. A linear regression of the form $K = Q^T W + b$ was used to quantify the effect of the weights on the parameters in question. Here Q is the vector of the coefficients and W is the vector that contains the base-10 logarithm of the weights. The results of the sensitivity analysis are shown in Table I for the coefficients of translation t_w , rotation r_w , clearance o_w , way-points uniformity d_w , alignment a_w , and visibility v_w .

TABLE I
LINEAR REGRESSION COEFFICIENTS

	t_w	r_w	o_w	d_w	a_w	v_w	b
Vis. Obj.	0.14	0.32	0.24	0.72	0.83	-1.22	-0.29
Traj. Len.	-0.50	-0.06	0.22	-0.57	-0.32	0.88	13.48
Align Dis.	-0.24	-0.03	-0.30	-0.15	-0.79	0.81	2.42


As expected, the only parameter that pushes the trajectory close to the visual objectives is the visual objective weight while the same parameter is the most significant in making the trajectory longer and worsening the alignment between the states. The trajectories shown in Figure 7(a),(c) were obtained using equal weights for all parameters and thus have smooth trajectories and are moving close to the visual objectives. On the other hand, for the trajectories in Figure 7(b) and (d), the weights for the visual objectives are increased to make the trajectory elongate and move towards the obstacles. In these cases, equally weighted parameters would cause AquaVis to miss all the visual objectives.

VI. CONCLUSION

The proposed framework will enable operations of the Aqua2 vehicle in a diverse set of environments. Environmental monitoring of coral reefs will be enhanced by guiding the robot towards corals with rich features instead of sand patches; see Figure 8a. Mapping underwater structures, such as shipwrecks, will benefit by ensuring the AUV operates in close proximity to the wreck and does not stray into open water where there are no features by which to localize; see Figure 8b. Finally, underwater caves—one of the most challenging environments for autonomous robots—present additional challenges due to the restricted lighting conditions [49]. AquaVis will guide the robot towards areas with enough light and texture to ensure safe operations; see Figure 8c.

REFERENCES


[1] M. Xanthidis, N. Karapetyan, H. Damron, S. Rahman, J. Johnson, A. O’Connell, J. M. O’Kane, and I. Rekleitis, “Navigation in the presence of obstacles for an agile autonomous underwater vehicle,” in *IEEE International Conference on Robotics and Automation*, 2020, pp. 892–899.



(a)



(b)



(c)

Fig. 8. The Aqua2 Autonomous Underwater Vehicle operating in different environments. (a) Surveying a coral reef, mapping the corals [44]. (b) Operating over the Stavronikita Shipwreck in Barbados. (c) Collecting data inside the Ballroom cavern in Ginnie Springs Florida.

[2] B. Joshi, S. Rahman, M. Kalaitzakis, B. Cain, J. Johnson, M. Xanthidis, N. Karapetyan, A. Hernandez, A. Quattrini Li, N. Vitzilios, and I. Rekleitis, “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)*, Macau, Nov. 2019, pp. 7221–7227.

[3] A. Quattrini Li, A. Coskun, S. M. Doherty, S. Ghasemlou, A. S. Jagtap, M. Modasshir, S. Rahman, A. Singh, M. Xanthidis, J. M. O’Kane, and I. Rekleitis, “Experimental Comparison of open source Vision based State Estimation Algorithms,” in *International Symposium of Experimental Robotics (ISER)*, Tokyo, Japan, Mar. 2016.

[4] S. Rahman, A. Quattrini Li, and I. Rekleitis, “Sonar Visual Inertial SLAM of Underwater Structures,” in *IEEE International Conference on Robotics and Automation*, Brisbane, Australia, May 2018, pp. 5190–5196.

[5] —, “An Underwater SLAM System using Sonar, Visual, Inertial, and Depth Sensor,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Macau, (IROS ICROS Best Application Paper Award Finalist), Nov. 2019, pp. 1861–1868.

[6] Y. Girdhar, “Unsupervised semantic perception, summarization, and autonomous exploration for robots in unstructured environments,” Ph.D. dissertation, McGill University Libraries, 2015.

[7] E. Bourque and G. Dudek, “On the automated construction of image-based maps,” *Autonomous Robots*, vol. 8, no. 2, pp. 173–190, 2000.

[8] G. Dudek, P. Giguere, C. Prahacs, S. Saunderson, J. Sattar, L.-A. Torres-Mendez, M. Jenkin, A. German, A. Hogue, A. Ripsman, *et al.*, “Aqua: An amphibious autonomous robot,” *Computer*, vol. 40, no. 1, pp. 46–53, 2007.

[9] H. J. S. Feder, J. J. Leonard, and C. M. Smith, “Adaptive mobile robot navigation and mapping,” *The International Journal of Robotics Research*, vol. 18, no. 7, pp. 650–668, 1999.

[10] C. Stachniss, D. Hahnel, and W. Burgard, “Exploration with active loop-closing for fastslam,” in *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*(IEEE Cat. No. 04CH37566), vol. 2. IEEE, 2004, pp. 1505–1510.

- [11] A. A. Makarenko, S. B. Williams, F. Bourgault, and H. F. Durrant-Whyte, "An experiment in integrated exploration," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 1. IEEE, 2002, pp. 534–539.
- [12] R. Martinez-Cantin, N. de Freitas, A. Doucet, and J. A. Castellanos, "Active policy learning for robot planning and exploration under uncertainty," in *Robotics: Science and Systems*, vol. 3, 2007, pp. 321–328.
- [13] I. Rekleitis, "Simultaneous Localization and Uncertainty Reduction on Maps (SLURM): Ear based exploration," in *2012 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2012, pp. 501–507.
- [14] —, "Multi-Robot simultaneous Localization and Uncertainty Reduction on Maps (MR-SLURM)," in *2013 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2013, pp. 1216–1221.
- [15] Q. Zhang, D. Whitney, F. Shkurti, and I. Rekleitis, "Ear-based exploration on hybrid metric/topological maps," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2014, pp. 3081–3088.
- [16] Q. Zhang, I. Rekleitis, and G. Dudek, "Uncertainty reduction via heuristic search planning on hybrid metric/topological map," in *2015 12th Conference on Computer and Robot Vision*. IEEE, 2015, pp. 222–229.
- [17] R. Lopez Padilla and R. Murrieta-Cid, "Maintaining visibility of a landmark using optimal sampling-based path planning," *Computación y Sistemas*, vol. 23, no. 4, 2019.
- [18] C. Forster, M. Pizzoli, and D. Scaramuzza, "Appearance-based active, monocular, dense reconstruction for micro aerial vehicles," in *Proceedings of Robotics: Science and Systems*, Berkeley, USA, July 2014.
- [19] B. Penin, R. Spica, P. R. Giordano, and F. Chaumette, "Vision-based minimum-time trajectory generation for a quadrotor uav," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 6199–6206.
- [20] R. Spica, P. Robuffo Giordano, and F. Chaumette, "Coupling active depth estimation and visual servoing via a large projection operator," *The International Journal of Robotics Research*, vol. 36, no. 11, pp. 1177–1194, 2017.
- [21] G. Costante, J. Delmerico, M. Werlberger, P. Valigi, and D. Scaramuzza, "Exploiting photometric information for planning under uncertainty," in *Robotics Research*. Springer, 2018, pp. 107–124.
- [22] M. Shekells, G. Garimella, and M. Kobilarov, "Optimal visual servoing for differentially flat underactuated systems," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Daejeon, South Korea, 2016, pp. 5541–5548.
- [23] T. Nägeli, J. Alonso-Mora, A. Domahidi, D. Rus, and O. Hilliges, "Real-time motion planning for aerial videography with dynamic obstacle avoidance and viewpoint optimization," *IEEE Robotics and Automation Letters*, vol. 2, no. 3, pp. 1696–1703, 2017.
- [24] T. Nägeli, L. Meier, A. Domahidi, J. Alonso-Mora, and O. Hilliges, "Real-time planning for automated multi-view drone cinematography," *ACM Transactions on Graphics (TOG)*, vol. 36, no. 4, pp. 1–10, 2017.
- [25] C. Potena, D. Nardi, and A. Pretto, "Effective target aware visual navigation for uavs," in *2017 European Conference on Mobile Robots (ECMR)*. IEEE, 2017, pp. 1–7.
- [26] D. Falanga, P. Foehn, P. Lu, and D. Scaramuzza, "Pampc: Perception-aware model predictive control for quadrotors," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1–8.
- [27] L. Yang, Z. Liu, X. Wang, and Y. Xu, "An optimized image-based visual servo control for fixed-wing unmanned aerial vehicle target tracking with fixed camera," *IEEE Access*, vol. 7, pp. 68 455–68 468, 2019.
- [28] C. Potena, D. Nardi, and A. Pretto, "Joint vision-based navigation, control and obstacle avoidance for uavs in dynamic environments," in *2019 European Conference on Mobile Robots (ECMR)*. IEEE, 2019, pp. 1–7.
- [29] K. Lee, J. Gibson, and E. A. Theodorou, "Aggressive perception-aware navigation using deep optical flow dynamics and pixelmpc," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1207–1214, 2020.
- [30] V. Murali, I. Spasojevic, W. Guerra, and S. Karaman, "Perception-aware trajectory generation for aggressive quadrotor flight using differential flatness," in *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 3936–3943.
- [31] I. Spasojevic, V. Murali, and S. Karaman, "Perception-aware time optimal path parameterization for quadrotors," *arXiv preprint arXiv:2005.13986*, 2020.
- [32] V. Indelman, L. Carlone, and F. Dellaert, "Planning in the continuous domain: A generalized belief space approach for autonomous navigation in unknown environments," *The International Journal of Robotics Research*, vol. 34, no. 7, pp. 849–882, 2015.
- [33] M. Greeff, T. D. Barfoot, and A. P. Schoellig, "A perception-aware flatness-based model predictive controller for fast vision-based multi-rotor flight," in *21st IFAC World Congress*, 2020.
- [34] B. Zhou, J. Pan, F. Gao, and S. Shen, "Raptor: Robust and perception-aware trajectory replanning for quadrotor fast flight," *arXiv preprint arXiv:2007.03465*, 2020.
- [35] Z. Zhang and D. Scaramuzza, "Fisher information field: an efficient and differentiable map for perception-aware planning," *arXiv preprint arXiv:2008.03324*, 2020.
- [36] S. Frolov, B. Garau, and J. Bellingham, "Can we do better than the grid survey: Optimal synoptic surveys in presence of variable uncertainty and decorrelation scales," *Journal of Geophysical Research: Oceans*, vol. 119, no. 8, pp. 5071–5090, 2014.
- [37] S. M. Chaves, A. Kim, E. Galceran, and R. M. Eustice, "Opportunistic sampling-based active visual slam for underwater inspection," *Autonomous Robots*, vol. 40, no. 7, pp. 1245–1265, 2016.
- [38] N. Karapetyan, J. Johnson, and I. Rekleitis, "Coverage path planning for mapping of underwater structures with an autonomous underwater vehicle," in *MTS/IEEE OCEANS - Singapore*, 2020.
- [39] T. Manderson, J. C. G. Higuera, R. Cheng, and G. Dudek, "Vision-based autonomous underwater swimming in dense coral for combined collision avoidance and target selection," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1885–1891.
- [40] J. Hong, K. de Langis, C. Wyeth, C. Walaszek, and J. Sattar, "Semantically-aware strategies for stereo-visual robotic obstacle avoidance," in *2021 International Conference on Robotics and Automation (ICRA)*. IEEE, 2021.
- [41] T. Manderson, J. C. G. Higuera, S. Wapnick, J.-F. Tremblay, F. Shkurti, D. Meger, and G. Dudek, "Vision-based goal-conditioned policies for underwater navigation in the presence of obstacles," *Robotics: Science and Systems XVI*, Jul 2020.
- [42] M. Modasshir, A. Quattrini Li, and I. Rekleitis, "Mdnet: Multi-patch dense network for coral classification," in *OCEANS 2018 MTS/IEEE Charleston*. IEEE, 2018, pp. 1–6.
- [43] M. Modasshir, S. Rahman, O. Youngquist, and I. Rekleitis, "Coral Identification and Counting with an Autonomous Underwater Vehicle," in *2018 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2018, pp. 524–529.
- [44] M. Modasshir, S. Rahman, and I. Rekleitis, "Autonomous 3D Semantic Mapping of Coral Reefs," in *12th Conference on Field and Service Robotics (FSR)*, Tokyo, Japan, 2019.
- [45] B. Joshi, M. Modasshir, T. Manderson, H. Damron, M. Xanthidis, A. Quattrini Li, I. Rekleitis, and G. Dudek, "DeepURL: Deep Pose Estimation Framework for Underwater Relative Localization," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Las Vegas, NV, USA, 2020, pp. 1777–1784.
- [46] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *KDD'96: Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, Portland, Oregon, 1996, pp. 226–231.
- [47] F. Shkurti, I. Rekleitis, and G. Dudek, "Feature tracking evaluation for pose estimation in underwater environments," in *Canadian Conference on Computer and Robot Vision (CRV)*, St. John, NF Canada, 2011, pp. 160–167.
- [48] J. Schulman, Y. Duan, J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, "Motion planning with sequential convex optimization and convex collision checking," *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1251–1270, 2014.
- [49] N. Weidner, S. Rahman, A. Quattrini Li, and I. Rekleitis, "Underwater cave mapping using stereo vision," in *IEEE International Conference on Robotics and Automation (ICRA)*, Singapore, May 2017, pp. 5709 – 5715.