2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
October 25-29, 2020, Las Vegas, NV, USA (Virtual)

Completeness Seeking Probabilistic Coverage Estimation using Uncertain State Estimates

Aditya Mahajan[†] and Stephen Rock[†‡]
{mahajan1, rock}@stanford.edu

Abstract— **This paper develops a coverage-centric adaptive path planner to visually survey a planar environment. This is achieved by modifying an existing path planning architecture to use a novel coverage estimation approach called convolved coverage estimation (CCE). The planner maximizes the probability of terrain coverage and exploits terrain features for loop closure to keep path uncertainty in check. The developed algorithm considers multi-dimensional uncertainty, operates in real-time, and does not require external correction methods like GPS. These characteristics are validated in high-fidelity simulation and flight tests on an unmanned aerial vehicle (UAV).**

I. Introduction

A. Motivation

This work is motivated by the need to conduct complete visual surveys of the sea floor for oceanographic science. Conceptually, the goal of the developed planner is to guarantee a minimum probability of coverage over the entire map with as few measurements as possible. The intent is to reduce the risk of missed coverage by developing a path planner that is coverage-conscious.

Ensuring a high-quality visual survey requires path planning that accounts for uncertainties in vehicle navigation, maintains sufficient sensor overlap to avoid missed coverage, and effectively uses terrain information for loop closure.

Current approaches to benthic surveying typically involve flying predetermined lawnmower paths with a high overlap between adjacent swaths [1], [2]. This low crosstrack spacing serves to minimize the risk of missed coverage by increasing image overlap. Many navigation methods exploit repeated observations of terrain by using visual-inertial simultaneous localization and mapping (SLAM) to adhere to the nominal trajectory. In many cases, such systems utilize expensive inertial sensors and prior knowledge of terrain information to maintain a trajectory estimate with low uncertainty.

As seen in the literature, complete coverage can be successfully achieved when operating with low navigation uncertainty, or in feature-rich terrain. However, in many situations high quality navigation estimates are difficult and an additional aid is needed.

B. Path Planner

The adaptive path planner described in this paper aims to automate the survey process and reduce the probability of missed coverage. This is done by estimating the probability of coverage online and directly using this estimate to plan the vehicle trajectory.

The developed planner has similarities to next best view (NBV) algorithms used in literature whereby the next optimal waypoint is chosen from a list of candidates. Optimality is based on several criteria including gaps in coverage, distance traveled, and whether a detour to informative terrain for loop closure would serve to improve the coverage estimate. More details are given in Section III.

II. Background

Coverage estimation and coverage path planning (CPP) are widely studied problems in the literature [3]. Missions focused on CPP have many underwater applications such as inspection of vessels [4]–[6], terrain surface reconstruction [7], and mine hunting [8], [9]. The underwater environment presents engineers with the added challenge of state estimation in the absence of GPS. This has led to the long and rich history of using vision aided underwater SLAM [10].

Research in underwater CPP has focused on improving navigation performance and minimizing localization uncertainty. Path planners have exploited visual information to induce loop closure and adhere to nominal boustrophedon [11] trajectories with great success. Kim [12] develops such a planner that minimizes uncertainty along a nominal path without explicitly tracking coverage. The result is a vehicle trajectory that closely follows the nominal path with deviations designed to improve navigation performance.

Others have improved CPP navigation performance by using a prior saliency map [13], following walls [14] or contours [7], or using a fleet of multiple robots [15]. Other works have exploited the ability to revisit informative terrain features [16] or intermittently surface to acquire GPS signals [9] to reduce uncertainty. The common thread for much of this research has been to improve coverage indirectly by attaining accurate navigation performance while following a path planned ahead of time.

A few works have performed CPP by addressing probabilistic coverage directly, which is a measure of coverage when the vehicle pose is non-deterministic. This idea is examined by Das *et al.* [17] where the authors note that in the presence of uncertainty, the path followed by a robot can deviate from the plan and result in an inaccurate coverage estimate. They show that performance can be improved by using an algorithm that accounts for path uncertainty and produces a feedback policy accordingly.

[†] Department of Aeronautics & Astronautics, Stanford University, Stanford, CA 94305, United States
[†‡]Monterey Bay Aquarium Research Institute, Moss Landing, CA 95039, United States

Probabilistic coverage is also addressed in more recent work [8], [9] where a side-scan sonar is used to calculate mine detection performance on each cell in a map. This is achieved by applying a transformation for continuous random variables [18]. The authors leverage the nature of the sensor (modeled by a closed form, differentiable, and one-to-one invertible function) to derive a probability density function (PDF) for mine detection. Due to the nature of the sensor, the PDF is calculated perpendicular to the vehicle trajectory only.

The approach presented in this paper is to directly address probability of coverage much like [8], [9], but in a fashion that allows real-time coverage estimates incorporating multidimensional uncertainty. However, the algorithm developed here calculates an expectation of coverage probability at each point on the map [19], not a complete PDF.

The adaptive path planner in this paper is structurally similar to the work by Kim [12] in that deviations to exploit loop closure on informative terrain are considered if this minimizes the cost function. This work is different in that the cost function directly incorporates an estimate of coverage probability. The change in cost function also dictates a change to the planner output. This adaptive planner outputs the location of the next optimal waypoint in addition to any intermediate deviations for loop closure. It does not strictly adhere to a nominal path (*e.g.* a lawnmower).

III. Technical Details

A. Adaptive Path Planner

The adaptive planner assumes that a history of vehicle poses and associated uncertainties is available. This information can be provided by several SLAM algorithms. The planner also depends on the ability to estimate coverage given vehicle poses and uncertainties. The coverage estimation process is based on convolved coverage estimation (CCE) [19] and is explained in Section III-B. The planner executes after every measurement and terminates when the estimated coverage is considered complete. The pseudocode for the adaptive planner is presented in Algorithm 1.

Conceptually, the adaptive planner chooses a target waypoint location from a list of candidates to maximize coverage. The candidate pool is generated based on the current estimate of vehicle position. Intermediate waypoints are placed based on observed terrain features if this increases the probability of coverage at the target. These choices are made to maximize the score S_i given in (1).

$$S_i = (c_i - \epsilon\ d_i)\ (1 - \mathbf{1}(g_i(y_i, \delta))) \tag{1}$$

where:

- i : candidate waypoint at location (x_i, y_i).
- c_i : area of new cells covered. See Section III-B for details.
- ϵ : small positive value.
- d_i : distance traveled to reach (x_i, y_i), inclusive of intermediate waypoints.
- δ : distance parameter for tuning the size of gaps that are to be ignored.
- g_i : size of coverage gap; area of cells not covered between y_i and $y_i - \delta$.

Maximizing the score results in the following emergent behavior:

- The distance traveled is only considered if multiple candidates give equal coverage. This allows the vehicle to deviate form the nominal trajectory to close a gap or close a loop at a terrain feature with little regard to distance traveled. The coverage/distance balance can be adjusted by varying the value of ϵ.
- The g_i function prevents the vehicle from deviating to cover a gap that is farther than δ. The gap will eventually be covered, but the algorithm determines this to be a lower priority.
- Ignoring small gaps in the short term has the added effect of giving gaps the opportunity to implicitly close in the event of a loop closure.
- If the terrain information is abundant and well distributed, the resulting trajectory is similar to a traditional lawnmower.

Algorithm 1: Adaptive path planner

Data: Alongtrack increment: Δx; map boundary; GraphSLAM output: $\vec{\mu},\ \Sigma$
Result: Next waypoint: (x_i, y_i^*); intermediate waypoint: (x_{int}, y_{int})

coverage $\longleftarrow CCE(\vec{\mu},\ \Sigma)$
if *coverage is complete* **then**
    return $\emptyset$
end
$x_i \longleftarrow x_{i-1} + \Delta x$, increment alongtrack coordinate (decrement if x_i outside map boundary)
Generate list of N candidates $y_i^{(1:N)}$ in map boundary
for *Each candidate* n
    Simulate direct travel to $(x_i, y_i^{(n)})$
    Calculate vehicle uncertainty, $\Sigma_i^{(n)}$
    if $\mathrm{tr}(\Sigma_i^{(n)}) >$ *threshold* **then**
        for *Each terrain feature,* j
            Simulate travel to $(x_i, y_i^{(n)})$ via (x_j, y_j)
            Calculate vehicle uncertainty, $\Sigma_j^{(n)}$
        end
        $j^* \longleftarrow \mathrm{argmin}_j\ \mathrm{tr}(\Sigma_j^{(n)})$
        $(x_{int}, y_{int}) \longleftarrow (x_{j^*}, y_{j^*})$
        $\Sigma_i^{(n)} \longleftarrow \Sigma_{j^*}^{(n)}$
    else
        $(x_{int}, y_{int}) \longleftarrow \emptyset$
    end
    Get CCE from visiting $(x_i, y_i^{(n)})$ via (x_{int}, y_{int})
    Calculate score, $S_i^{(n)}$;
end
$n^* \longleftarrow \mathrm{argmax}_n\ S_i^{(n)}$
$y_i^* \longleftarrow y_i^{(n^*)}$
return (x_i, y_i^*); (x_{int}, y_{int})

As seen on Lines 6-7 of Algorithm 1, the adaptive path planner chooses one from a list of N generated candidate waypoints after every measurement. To keep processing time to a minimum, the candidates $1:N$ all have have the same along-track coordinate x_i, but vary in terms of their crosstrack coordinate y_i. An example of this approach is presented in Figure 1. This illustrates a sample of candidate waypoints (*A*, *B*, and *C*) in red, located at various crosstrack coordinates on the map.

The planner simulates travel to each candidate and propagates vehicle position uncertainty with EKF–SLAM. If position uncertainty (measured as the trace of the covariance) can be reduced by traveling via a previously observed terrain feature, an intermediate waypoint is placed at the terrain feature.

The trace of the covariance is used as a scalar measure of uncertainty as it preserves more of the overall size [20]. The determinant of the covariance is more common but this measure collapses to zero even if a single eigenvalue of the matrix goes to zero. The trace does not suffer from this drawback.

Once the vehicle uncertainty at each candidate is determined, the probability of coverage from traveling to each candidate is calculated using CCE on Line 21. The sensor footprints (shown in Figure 1 as dashed boxes) are determined by truncating this probability with a user-defined threshold explained in Section III-B. This is used to calculate the score on Line 22.

The optimal waypoint is chosen to maximize the score, and the planner outputs that waypoint's location. For the illustrated example, waypoint *B* is chosen. This is because *C* doesn't result in any new coverage and *A* results in too large of a coverage gap in the $-y$ direction.

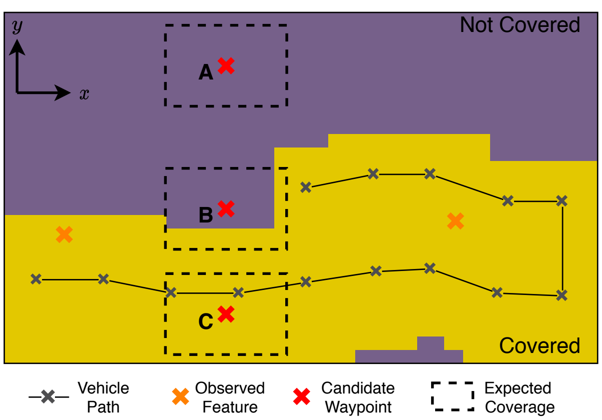

Fig. 1: The adaptive path planner evaluates candidate waypoints based on: expected coverage, vehicle pose uncertainty, distance traveled, and gaps in coverage.

Overall, the adaptive planner greedily maximizes coverage but adapts to vehicle uncertainty in real-time. When vehicle uncertainty is low, overlap is reduced and more new terrain is covered. When uncertainty is high, overlap is increased, either by eroding the expected coverage using CCE (see Section III-B) or by revisiting informative terrain.

B. Coverage Estimation

A key step in the planning process is estimating the probability of coverage. This estimate is calculated by convolving the state estimate with the sensor footprint. CCE has the effect of eroding the sensor footprint in proportion to the direction of vehicle pose uncertainty. The extent of erosion is controlled by a user defined threshold that represents a trade-off between false positive and false negative coverage. This work is presented in [19]. Overall, CCE results in fewer false positives when compared to only using the best estimate of the vehicle pose, particularly at instances of high uncertainty.

The probabilistic estimate from CCE is thresholded with a user-defined value. The area of cells that exceeds this threshold is c_i.

IV. Results

This section shows the results of a coverage mission with a UAV in a high-fidelity simulated environment as well as flight tests. For simulated missions, the adaptive path planner was implemented on a simulated UAV flying over flat terrain. The UAV model was a generic quadrotor available in the open-source *Gazebo* environment (see Figure 2a).

Flight experiments were conducted in the Boeing Flight & Autonomy Laboratory at Stanford University. This facility is equipped with an *OptiTrack* motion capture system, capable of determining vehicle position and orientation to sub-millimeter accuracy. Data from *OptiTrack* was used as ground truth for evaluating results. The UAV chosen for these tests was an *Intel® Aero RTF Drone*, equipped with a downward facing camera. Figure 2b shows a sample flight test in progress. Similar to *Gazebo*, fiducial tags (with no prior location knowledge) were used as analogs for terrain information.

In both environments, the vehicle was initialized in a corner and commanded to achieve a minimum 90% probability of coverage. Artificial noise was added to all three position states as well as heading to simulate error due to dead-reckoning. This included a constant, unknown drift rate for the x and y position states. Pose estimation was done using graph-based simultaneous localization and mapping (GraphSLAM) [21]. This algorithm generates motion constraints between odometry and terrain features to solve the full SLAM problem. AprilTags [22] fiducial markers were placed on the ground plane to serve as terrain information. Fiducial tags were used for convenience to enable feature correspondence that is outside the scope of this work.

A. Simulated Mission - Predetermined Lawnmower

The extent of the artificially added noise is illustrated in Figure 3 where the vehicle was commanded to fly a predetermined lawnmower path, followed by a diagonal pass to induce loop closure. This figure shows the path that was estimated and the path that was flown. The vehicle trajectory exhibits significant overlap, but not over informative regions. This results in incomplete coverage.

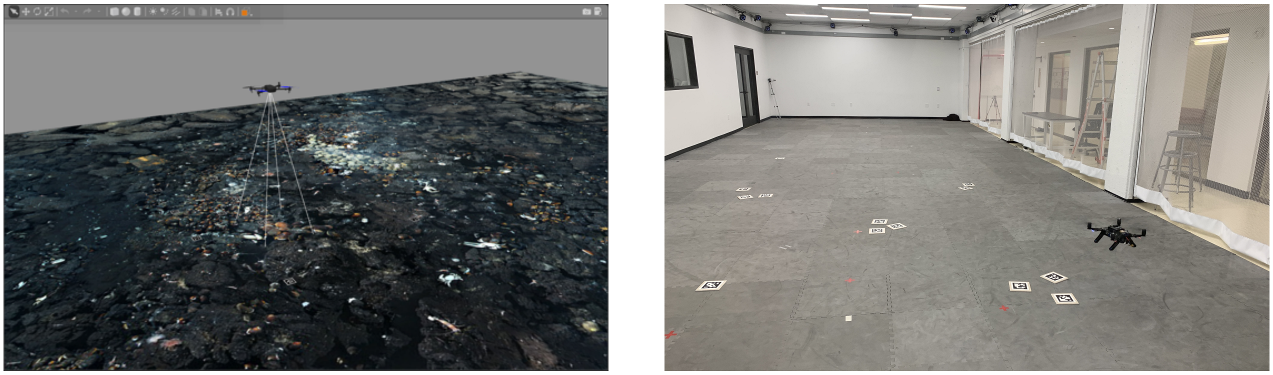

(a) Screenshot of UAV simulation in *Gazebo*

(b) Image of UAV during flight

Fig. 2: Results were obtained using a UAV in simulation as well as flight experiments

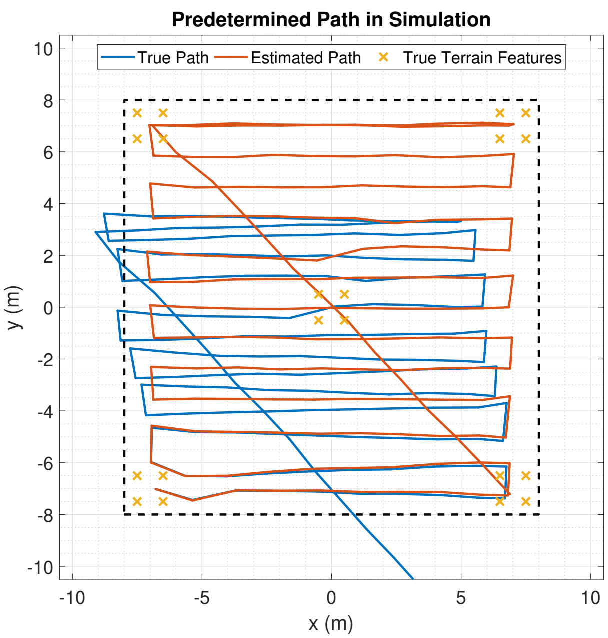

Fig. 3: UAV trajectory when following a predetermined path with artificially corrupted inertial measurements. The red estimated path follows a lawnmower that is expected to cover the map. The true trajectory, shown in blue, exhibits significant drift and fails to adequately cover the workspace.

B. Simulated Mission - Adaptive Path Planner

The simulated mission was carried out with the aim of covering a square area measuring 16 meters on a side. The UAV was first commanded to image one edge of the map, then the adaptive planner was activated to complete the survey autonomously.

For this mission, the UAV proceeded to cover the map using the adaptive planner for optimizing waypoint location and CCE for coverage estimation. The UAV made a few visits to informative terrain (see Figure 4a) and those were spread between the features in the middle of the map (moderately certain) to those in the lower left corner (highly certain). Repeatedly visited areas can also be seen in Figure 4b, where the bright yellow region in the center indicates more frequent revisits to informative terrain. Note that there is no red (false positive coverage) in this figure, indicating complete coverage.

In another experiment, the simulated UAV was flown over highly informative terrain (see Figure 5). The resulting survey was a lawnmower-like path with very little overlap. No specific areas were visited with more frequency as the terrain features were evenly distributed throughout the workspace. Complete coverage was achieved in this case as well.

C. Flight Test - Adaptive Path Planner

The flight test was carried out in a smaller area (measuring 7 m $\times$ 5 m) due to space constraints. The approach taken to conduct this mission was identical to that for the simulated terrain.

The path flown has similarities to the simulated mission. It can be seen in Figure 6a that the UAV makes several trips to the feature-rich region in the center of the map. Note that true terrain feature locations were unavailable so the feature locations estimated using GraphSLAM are plotted instead.

The coverage resulting from this mission was 99.9% complete with only 0.12% of the area falsely identified as being covered. This slight gap in coverage is visible in Figure 6b and this highlights the fact that CCE does not guarantee complete coverage. The results of these tests are displayed in Table I.

V. Conclusion

This paper develops a coverage-centric algorithm that directly addresses probability of coverage when conducting visual surveys. CCE ties uncertainties in the vehicle path estimate to uncertainty in the coverage estimate and the adaptive path planner maximizes coverage by optimizing the location of the next waypoint. This process takes aspects such as vehicle pose uncertainty, distance traveled, and gaps in coverage into consideration.

Although this approach does not provide completeness guarantees, the algorithm maximizes the probability of coverage by revisiting informative terrain only when necessary.

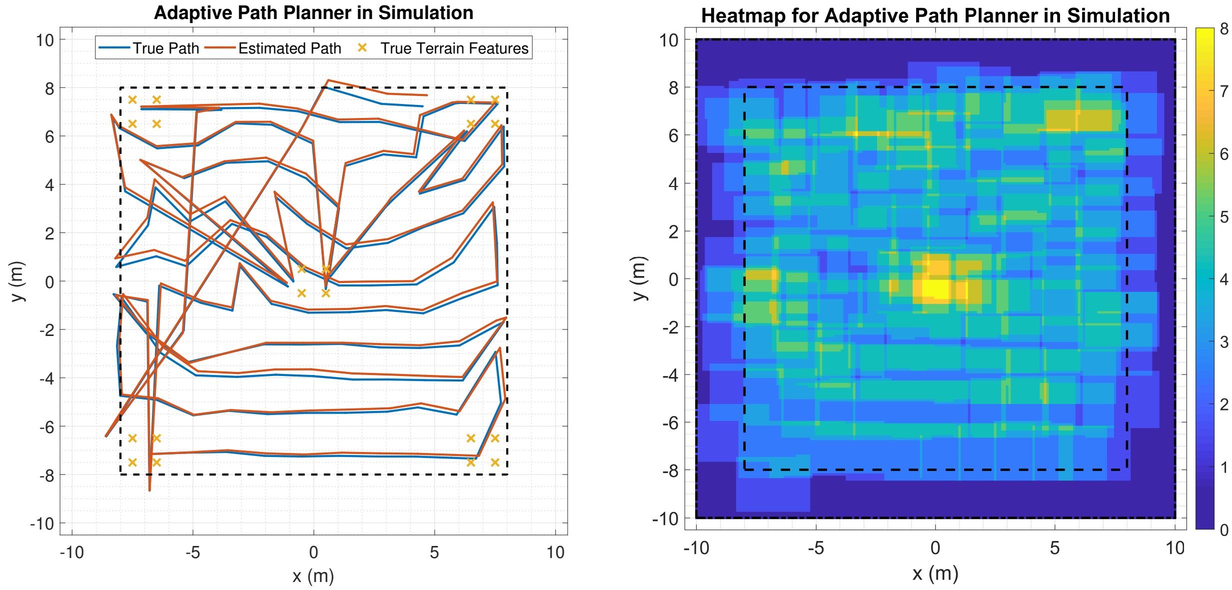

(a) UAV trajectory from adaptive path planning. The estimated trajectory (red) is very close to the true trajectory (blue). Yellow crosses represent true locations of terrain information.

(b) Heatmap indicating revisited areas for adaptive path planning in simulation. The absence of red areas indicates complete coverage.

Fig. 4: Simulation results from surveying an area with a UAV

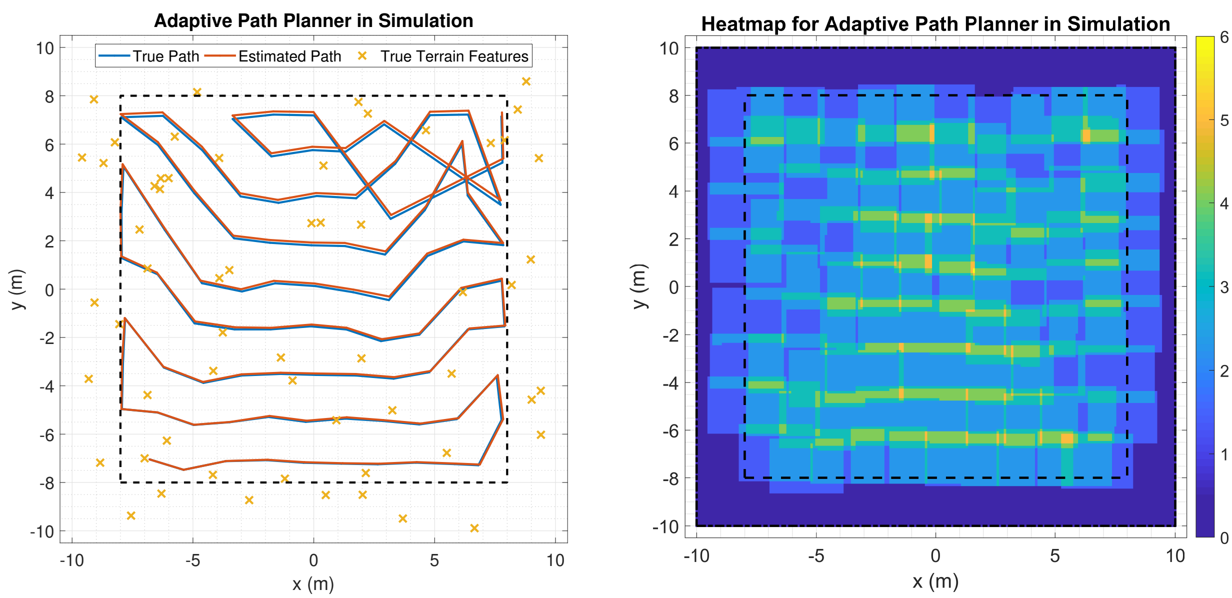

(a) UAV trajectory from adaptive path planning on feature-rich terrain. The ubiquity of terrain information makes revisits to specific areas unnecessary and the trajectory looks more like a typical lawnmower in this case.

(b) Heatmap indicating revisited areas for feature-rich terrain. The adaptive planner minimizes overlap in order to maximize coverage when surveying terrain with many features that are evenly distributed.

Fig. 5: Simulation results from surveying a feature-rich area with a UAV

TABLE I: Summary of results: visual surveys conducted using adaptive path planner in simulation and UAV flight test.

Environment	Terrain Condition	Mission Time (s)	Coverage (%)		Path Length (m)
			True	False Pos.	
Simulation	Poor features, localized in clusters	877	100	0	292
Simulation	Feature-rich, randomly & uniformly distributed	582	100	0	173
Flight test	Poor features, localized in clusters	312	99.9	0.12	72.7

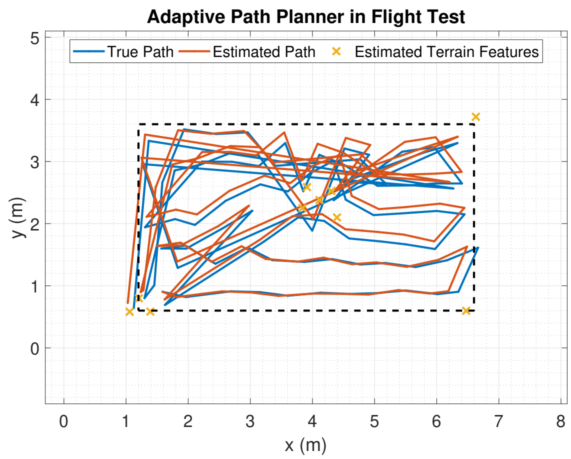

(a) UAV flight trajectory from adaptive path planning. The estimated trajectory (red) is very close to the true trajectory (blue) at times of high certainty (lower part of the survey). Yellow crosses represent estimated locations of terrain information.

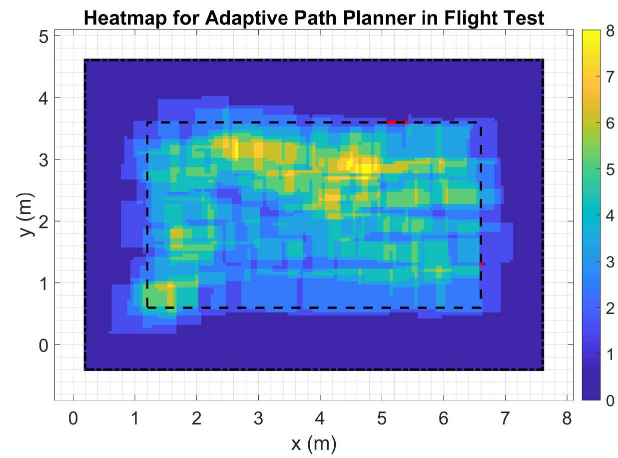

(b) Heatmap indicating revisited areas for adaptive path planning in flight experiments. The red area on the upper edge indicates missed coverage.

Fig. 6: Flight experiment results from surveying an area with a UAV

This is shown in experiments carried out in high-fidelity simulated environments and flight tests.

Overall, the developed algorithm allows the vehicle to revisit missed areas and exploit terrain information to avoid redeployment. This reduces the cost of the mission and the burden on operators.

ACKNOWLEDGMENT

This work has been supported in part by the Monterey Bay Aquarium Research Institute (MBARI).

REFERENCES

[1] K. Richmond and S. M. Rock, "An Operational Real-Time Large-Scale Visual Mosaicking and Navigation System," in *OCEANS 2006*, pp. 1–6, IEEE, 9 2006.

[2] K. Richmond, *Real-Time Visual Mosaicking and Navigation on the Seafloor*. PhD thesis, Stanford University, Stanford, CA 94305, 3 2009.

[3] E. Galceran and M. Carreras, "A survey on coverage path planning for robotics," *Robotics and Autonomous Systems*, vol. 61, pp. 1258–1276, 12 2013.

[4] B. Englot and F. S. Hover, "Sampling-based Coverage Path Planning for Inspection of Complex Structures," in *Proceedings of the 22nd International Conference on Automated Planning and Scheduling*, (Atibaia, Sao Paulo, Brazil), pp. 29–37, AAAI Press, 2012.

[5] P. Ozog and R. M. Eustice, "Toward long-term, automated ship hull inspection with visual SLAM, explicit surface optimization, and generic graph-sparsification," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3832–3839, IEEE, 5 2014.

[6] A. Kim and R. M. Eustice, "Active visual SLAM for robotic area coverage: Theory and experiment," *The International Journal of Robotics Research*, vol. 34, pp. 457–475, 4 2015.

[7] E. Galceran, R. Campos, N. Palomeras, D. Ribas, M. Carreras, and P. Ridao, "Coverage Path Planning with Real-time Replanning and Surface Reconstruction for Inspection of Three-dimensional Underwater Structures using Autonomous Underwater Vehicles," *Journal of Field Robotics*, vol. 32, pp. 952–983, 10 2015.

[8] L. Paull, M. Seto, and H. Li, "Area coverage planning that accounts for pose uncertainty with an AUV seabed surveying application," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6592–6599, IEEE, 5 2014.

[9] N. Abreu, N. Cruz, and A. Matos, "Accounting for uncertainty in search operations using AUVs," in *2017 IEEE Underwater Technology (UT)*, pp. 1–8, IEEE, 2017.

[10] I. Mahon and S. B. Williams, "SLAM using natural features in an underwater environment," in *ICARCV 2004 8th Control, Automation, Robotics and Vision Conference, 2004.*, vol. 3, pp. 2076–2081, IEEE, 2004.

[11] H. Choset and P. Pignon, "Coverage Path Planning: The Boustrophedon Cellular Decomposition," in *Field and Service Robotics*, pp. 203–209, London: Springer London, 1998.

[12] A. Kim and R. M. Eustice, "Perception-driven navigation: Active visual SLAM for robotic area coverage," in *2013 IEEE International Conference on Robotics and Automation*, pp. 3196–3203, IEEE, 5 2013.

[13] E. Galceran, S. Nagappa, M. Carreras, P. Ridao, and A. Palomer, "Uncertainty-driven survey path planning for bathymetric mapping," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 6006–6012, IEEE, 11 2013.

[14] E. U. Acar and H. Choset, "Exploiting critical points to reduce positioning error for sensor-based navigation," in *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No.02CH37292)*, vol. 4, pp. 3831–3837, IEEE, 2002.

[15] S. Tully, G. Kantor, and H. Choset, "Leap-Frog Path Design for Multi-Robot Cooperative Localization," in *Field and Service Robotics* (A. Howard, K. Iagnemma, and A. Kelly, eds.), pp. 307–317, Berlin, Heidelberg: Springer Berlin Heidelberg, 2010.

[16] A. Kim, *Active Visual SLAM with Exploration for Autonomous Underwater Navigation*. PhD thesis, The University of Michigan, 2012.

[17] C. Das, A. Becker, and T. Bretl, "Probably approximately correct coverage for robots with uncertainty," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1160–1166, IEEE, 9 2011.

[18] A. M. Mood, *Introduction to the Theory of Statistics*. McGraw-Hill Book Company, 1950.

[19] A. Mahajan and S. M. Rock, "A Pilot Aid for Real-Time Vision-Based Coverage Estimation for Seabed Surveying Applications," in *OCEANS 2020 MTS/IEEE Singapore*, (in press), IEEE, 2020.

[20] R. Sim and N. Roy, "Global A-Optimal Robot Exploration in SLAM," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pp. 661–666, IEEE, 2005.

[21] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press, 2005.

[22] E. Olson, "AprilTag: A robust and flexible visual fiducial system," in *2011 IEEE International Conference on Robotics and Automation*, pp. 3400–3407, IEEE, 5 2011.