

Stochastic Planning for ASV Navigation Using Satellite Images

Yizhou Huang¹, Hamza Dugmag², Timothy D. Barfoot², and Florian Shkurti¹

Abstract—Autonomous surface vessels (ASV) represent a promising technology to automate water-quality monitoring of lakes. In this work, we use satellite images as a coarse map and plan sampling routes for the robot. However, inconsistency between the satellite images and the actual lake, as well as environmental disturbances such as wind, aquatic vegetation, and changing water levels can make it difficult for robots to visit places suggested by the prior map. This paper presents a robust route-planning algorithm that minimizes the expected total travel distance given these environmental disturbances, which induce uncertainties in the map. We verify the efficacy of our algorithm in simulations of over a thousand Canadian lakes and demonstrate an application of our algorithm in a 3.7 km-long real-world robot experiment on a lake in Northern Ontario, Canada.

I. INTRODUCTION

Autonomous Surface Vessels (ASVs) have seen increasing attention as a technology to monitor rivers, lakes, coasts, and oceans in recent years [1]–[8]. A fundamental challenge to the wide adoption of ASVs is the ability to navigate safely and autonomously in uncertain environments, especially for long durations. For example, a robot may precompute a plan that visits certain target areas on a map and execute it online. However, disturbances such as strong winds, waves, aquatic plants, unseen obstacles, and even simply changing visual appearances in a water environment are challenging for ASV navigation. Many potential low-level failures in robot perception and control systems may also undermine the overall success of the mission.

Our goal is to use an ASV to monitor lake environments and collect water samples for scientists. To make the system robust, we propose to identify uncertainties (e.g., obstacles, wind-prone areas) that could block waterways, incorporate them into the map, and generate stochastic policies that can adapt online. One planning framework that is suitable for modelling uncertain paths is the Canadian Traveller Problem (CTP) [9]. The most significant feature in a CTP graph is the stochastic edge, which has a probability of being blocked. The state of any stochastic edge can be disambiguated by visiting the edge. Once the state has been classified as traversable or not, it remains the same.

In this paper, we propose a navigation framework — the Partial Covering Canadian Traveller Problem (PCCTP) — to solve a route-planning problem in an uncertain environment. The framework uses a stochastic graph derived from coarse satellite images to plan an adaptive policy that visits all




Fig. 1. Our Clearpath Heron ASV for water-quality monitoring during a field test. The ASV has onboard sensors (GPS, IMU, underwater scanning sonar, stereo camera) and an Nvidia Jetson to process sensor measurements.

reachable target locations. Stochasticity in the graph represents possible events where a water passage between two points is blocked due to changing water levels, strong wind, and other unmapped obstacles.

Our proposed solution computes the optimal policy offline with a best-first tree-search algorithm. The benefits of our proposed approach are twofold: 1) Computing an optimal policy can reduce the total distance to visit all targets and return, 2) The execution of our stochastic policy is robust even in the presence of environmental disturbances and partial robot system failure. We evaluate our solution method on a random set of Canadian lakes and demonstrate its robustness on a robot system that we built for field tests.

II. RELATED WORKS

Remote sensing is a popular technique to build maps and monitor changes in water bodies around the world because of its efficiency [10], [11]. The JRC Global Surface Water dataset [12] maps changes in water coverage from 1984 to 2015 at a 30 m by 30 m resolution, produced using Landsat satellite imagery. Since water has a lower reflectance in the infrared channel, an effective method is to calculate water indices, such as NDWI [13] or MNDWI [14], from two or more optical bands (e.g., green and near-infrared). However, extracting water data using a threshold in water indices can be non-trivial due to variations introduced by clouds, seasonal changes, and sensor-related issues. To address this, [15] and [16] have developed techniques to select water-extraction thresholds adaptively. Our approach aggregates water indices from historical satellite images to estimate probabilities of water coverage (see Sec. III-C). Overall, we argue that it is beneficial to build stochastic models of surface water bodies due to their dynamic nature and imperfect

¹Department of Computer Science, University of Toronto, 40 St George St, Toronto, Canada {phuang, florian}@cs.toronto.edu

²University of Toronto Institute for Aerospace Studies, 4925 Dufferin St, Toronto, Canada hamza.dugmag@mail.utoronto.ca, tim.barfoot@utoronto.ca




Fig. 2. A toy example graph shown on a satellite image. The planned paths between nodes are simplified for ease of understanding. The number beside each edge of the top-right graph is the path length in km, and the number in brackets is the blocking probability, which is computed using the probability of water coverage in each pixel (represented by its shade of blue) on the path. Note that traversable and ambiguous edges are the state before any action.

knowledge derived from satellite images.

The other significant pillar of building an ASV navigation system is route planning. First formulated in the 1930s, the Travelling Salesperson Problem (TSP) [17] studies how to find the shortest path in a graph that visits every node once and returns to the starting node. Modern TSP solvers such as the *Google OR-tools* [18] can produce high-quality approximate solutions for graphs with about 20 nodes in a fraction of a second. In many cases, the problem graphs are built from real-world road networks and the edges are assumed to be always traversable. In CTP [9], however, edges can be blocked with some probability. The goal is to compute a policy that has the shortest expected path to travel from a start node to a single goal node. CTP can also be formulated as a Markov Decision Process [19] and solved optimally with dynamic programming [20] or heuristic search [21]. The robotics community has also studied ways in which the CTP framework can be best used in path planning [22], [23]. Our problem setting, PCCTP, lies at the intersection of TSP and CTP, where the goal is to visit a partial set of nodes on a graph with stochastic edges. A similar formulation, known as the Covering Canadian Traveller Problem (CCTP) [24], presents a heuristic, online algorithm named Cyclic Routing (CR) to visit every node in a complete n -node graph with at most $n - 2$ stochastic edges. A key distinction between CCTP and our setting is that CCTP assumes all nodes are reachable, whereas the robot may give up on unreachable nodes located behind an untraversable edge in PCCTP.

In recent years, more ASV systems and algorithms for making autonomous decisions to monitor environments have been built. Schiaretti et al. [25] classify the autonomy level for ASVs into 10 levels based on control systems, decision-making, and exception handling. Many works consider the

mechanical, electrical, and control subsystems of their ASV designs [3], [4], [7]. Dash et al. [6] validated the use and accuracy of deploying ASVs for water-quality modelling by comparing the data collected from ASVs against independent sensors. Two examples of vertically integrated autonomous water-quality monitoring systems using ASVs are presented in [26] and [5]. In contrast, our main contribution is a robust mission-planning framework that is complementary to existing designs of ASV systems. Finally, informative path planning is another orthogonal area where the robot relies on a probabilistic model to identify targets that maximize information gain; [27] reviews this topic.

III. METHODOLOGY

A. The Problem Formulation

We are interested in planning on a graph representation of a lake where parts of the water are stochastic (i.e., traversability is uncertain). Constructing such a graph using all pixels of satellite images is impractical since images are very high-dimensional. Thus, we extend previous works from CTP [9], [23], [24] and distill satellite images into a high-level graph G where some stochastic edges e may be untraversable with probability p . The state of a stochastic edge can be disambiguated only when the robot traverses the edge in question. The robot begins at the starting node s and is tasked to visit all reachable targets J specified by the user (e.g., scientists) before returning to the starting node. If some target nodes are unreachable because some stochastic edges block them from the starting node, the robot may give up on these sampling targets. Hence we call this problem PCCTP. The state of the robot is defined as a collection of the following: a list of target nodes that it has visited, the current node it is at, and its knowledge about the stochastic edges. A policy sets the next node to visit given the current state of the robot. The objective is to find the optimal policy π^* that minimizes the expected cost to cover all reachable targets. Formally, we define the following terms:

- $G = (V, E)$ is an undirected graph.
- $c : E \rightarrow \mathbb{R}_{\geq 0}$ is the cost function for an edge, which is the length of the shortest waterway between two points.
- $p : E \rightarrow [0, 1]$ is the blocking probability function.
- k is the number of stochastic edges.
- $s \in V$ is the start and return node.
- $J \subseteq V$ is the subset of target nodes to visit.
- $I = \{A, T, U\}^k$ is an information vector that represents the robot's knowledge of the status of all k stochastic edges. A, T, and U stand for ambiguous, traversable, and untraversable, respectively.
- $S \subseteq J$ is the subset of target nodes that have been visited.
- a is the current node the robot is at.
- $x = (a, S, I)$ is the state of the robot.
- π^* is the optimal policy that minimizes the cost $\mathbb{E}[\phi(\pi)]$, where ϕ is cost functional of the policy π .

B. Solving PCCTP with AO*

We extend the AO* search algorithm [21] used in CTP to find exact solutions to our problem. AO* is a heuristic,

edges of the removed target. The heuristic is the cost of the best plan that covers all definitively reachable targets and disambiguates one of the critical stochastic edges. For example, consider computing the heuristic at starting node 0 in Fig. 2. The optimistic plan is to visit node 2, disambiguate the edge (2, 3), and return to node 0. If the robot reaches node 3, the heuristic would then be the remaining cost to visit node 1 and return to the start. This heuristic is always admissible because the path to disambiguate a critical edge is always a subset of the eventual policy. We can compute this by constructing an equivalent generalized travelling salesman problem [30] and solve it with any optimal TSP solver.

Fig. 3 shows the result of applying PCCTP-AO* to the example problem in Fig. 2. Note that the AO* algorithm stops expanding as soon as the lower bound of the cost of the right branch exceeds that of the left branch. This guarantees the left branch has the lower cost and thus, is optimal.

C. Estimating Stochastic Graphs From Satellite Imagery

We will now explain our procedure to estimate the high-level stochastic graph from satellite images.

Water Masking Our goal in the first step is to build a water mask of a water area across a specific period of time (e.g., 30 days). We use the *Sentinel-2* Level 2A dataset [31], which has provided images at 10 m by 10 m resolution since 2017. Each geographical location is revisited every five days by a satellite. We then select all satellite images in the target spatiotemporal window and filter out the cloudy images using the provided cloud masks. For each image, we calculate the Normalized Difference Water Index (NDWI) [13] for every pixel using green and near-infrared bands. However, the distribution of NDWI values varies significantly across different images over time. Thus, we separate water from land in each image and aggregate the indices over time. We then fit a bimodal Gaussian Mixture Model on the histogram of NDWIs to separate water pixels from non-water ones for each image. We average all water masks over time to calculate the probabilistic water mask at the target spatiotemporal window. Each pixel on the final mask represents the probability of water coverage on this 10 m by 10 m area. If a pixel may or may not have water, then we call it a stochastic pixel. Finally, we identify the boundary of all deterministic water pixels.

Stochastic Edge Detection: Pinch Point We can now identify those stochastic water paths (i.e., narrow straits, pinch points [22]) that are useful for navigation. A pinch point (e.g., Fig. 2) is defined as a sequence of stochastic water pixels that connect two parts of topologically far (or distinct) but metrically close water areas. Essentially, this edge is a shortcut connecting two points on the water boundary that are otherwise far away or disconnected. To find all such edges, we iterate over all boundary pixels, test each shortest stochastic water path to nearby boundary pixels, and include those stochastic paths that are shortcuts. The blocking probability of a stochastic edge is one minus the minimum water probability along the path. Since this process will produce many similar stochastic edges around the same

narrow passage, we run DBSCAN [32] and only choose the shortest stochastic edge within each cluster.

Stochastic Edge Detection: Windy Edges The second type of stochastic edges are those with strong wind. In practice, when an ASV travels on a path far away from the shore, there is a higher chance of running into a strong headwind or wave, making the path difficult to traverse. We define an edge to be a windy edge if it is 200 m away from the water boundary at some point, and assign a small probability for the event where the edge is blocked by wind.

Path Generation The next step is to construct the geotagged path and calculate all edge costs in the high-level graph. The nodes in the high-level graph are composed of all sampling targets, endpoints of stochastic edges, and the starting node. We run A* [33] on the deterministic water pixels to calculate the shortest path between every pair of nodes except for the stochastic edges found in the previous step. Since the path generated by A* connects neighbouring pixels, we smooth them by downsampling. Then, we can discard any unnecessary stochastic edges if they do not reduce the distance between a pair of nodes. Finally, we check if each deterministic edge is a windy edge, and obtain the high-level graph used in PCCTP.

IV. SIMULATIONS

A. Testing Dataset

We evaluate our route-planning framework on Canadian lakes selected from the *CanVec Series* Ontario dataset [34]. Published by *Natural Resources Canada*, this dataset contains geospatial data of over 1.1 million water bodies in Ontario. Considering a practical mission length, lakes are filtered such that their bounding boxes are 1-10 km by 1-10 km. Then, water masks of the resulting 5190 lakes are generated using *Sentinel-2* imagery across 30 days in June 2018-2022 [31]. We then detect any pinch points on the water masks and randomly sample five different sets of target nodes on each lake, each with a different number of targets. The starting locations are sampled near the shore.

Furthermore, we generate the high-level graphs as well as windy edges from the water mask. Graphs with no stochastic edges are removed as well as any instances with more than nine stochastic edges due to long run times. In the end, we evaluate our algorithm on 2217 graph instances, which come from 1052 unique lakes.

B. Baselines

The simplest baseline is an online greedy algorithm that always goes to the nearest unvisited target node assuming all ambiguous edges are traversable. For a graph with k stochastic edges, we simulate all 2^k possible worlds, each with a different traversability permutation, and evaluate our greedy actor on each one. The greedy actor recomputes a plan at every step and queries the simulator if it encounters a stochastic edge to disambiguate it. Also, it checks the reachability of every target node upon discovering an untraversable edge and gives up on any unreachable targets.

A more sophisticated baseline is the optimistic TSP algorithm. Instead of always going to the nearest target node, it

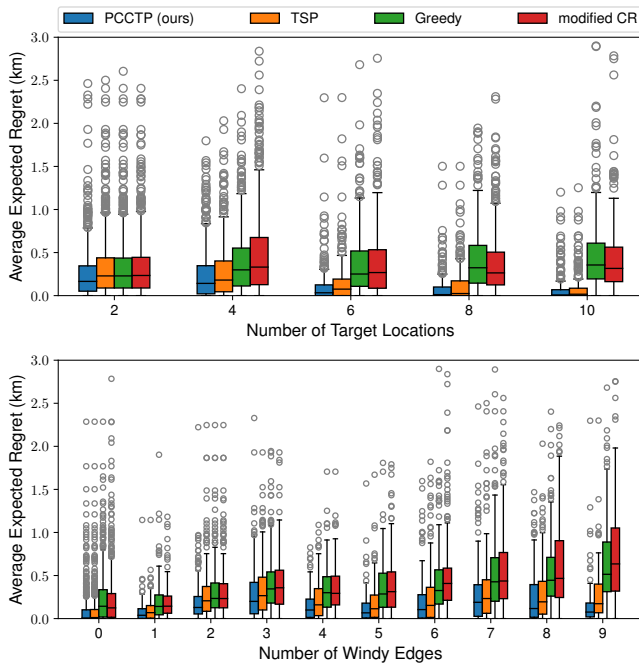


Fig. 4. Average expected regret of PCCTP and baselines. A total of 2217 graphs are grouped according to the number of target nodes and windy edges. When there are no windy edges, all stochastic edges are pinch points.

computes the optimal tour using dynamic programming to visit all remaining targets assuming all ambiguous edges are traversable. Similar to the greedy actor, TSP recomputes a tour at every step and may change its plan after it encounters an untraversable edge. The expected cost is also computed via a weighted sum on all 2^k possible worlds. In contrast to PCCTP, both TSP and greedy require onboard computation to update their optimistic plans.

Lastly, we modify the CR algorithm, originally a method for CCTP [24], to solve PCCTP. CR precomputes a cyclic sequence to visit all target nodes using the Christofides algorithm [35] and tries to visit all target nodes in multiple cycles while disambiguating stochastic edges. If a target node turns out to be unreachable, we allow CR to skip this node in its traversal sequence.

C. Results

Fig. 4 compares our algorithm against all baselines. To measure the performance across various graphs of different sizes, we use the average expected regret over all graphs. The expected regret of a policy π for one graph G is defined as

$$\mathbb{E}_w[\text{Regret}(\pi)] = \sum_w [p(w)(\phi(\pi, w) - \phi(\pi^P, w))],$$

where π^P is a privileged planner with knowledge of the states of all stochastic edges, ϕ is the cost functional, and w is a possible world of the graph. PCCTP precomputes the optimal policy in about 50 seconds on average in our evaluation, and there is no additional cost online. Compared to the strongest baseline (TSP), our algorithm saves the robot about 1%(50m) of travel distance on average and 15%(1.8km) in the extreme case. Although the improvement is marginal on average, our planner can still be beneficial in edge cases

(e.g., high blocking probability, long stochastic edges). The performance of PCCTP may be further enhanced if the estimated blocking probabilities of the stochastic edges are refined based on historical data.

V. REAL ROBOT EXPERIMENT

A. Robot System

Our ASV platform, shown in Fig. 1, is a customized *Clearpath Heron* ASV with a GPS, IMU, *ZED* stereo camera, *Ping360* scanning sonar, and an *Nvidia Jetson* for processing sensor inputs. A separate *Micro-Start XP-1* battery powers the *Jetson*, stereo camera, and sonar for approximately one hour, while two NiMH cells power the motors and remaining electronics for about two hours. We did not integrate water samplers for this mission but we can add them in the future. Our system also includes a custom web interface that displays the status of the robot and the planner. We have a remote controller used for manual mode if needed for safety.

B. Autonomy Stack for Policy Execution

Our philosophy is to design a robust autonomy framework where the success of policy execution is less dependent on lower-level components such as perception and local planning. At a high level, the planner precomputes the navigation policies from satellite images given user-selected sampling locations. A local planner then follows the global path during a mission. When the robot is disambiguating a stochastic edge, the policy executor will decide the traversability of the edge based on a timer. If the robot finishes travelling the stochastic edge within the time limit, the edge is marked as traversable. Otherwise, the edge is untraversable. The executor then branches into different policy cases depending on the disambiguation result. Using a timer allows us to account for issues we cannot directly sense (e.g., heavy headwind), and does not require a perfect terrain assessment to identify a stochastic edge successfully.

The local planner uses the Dynamic Window Algorithm (DWA) [36] to plan a desired velocity that tracks the global plan and avoids obstacles. The robot's odometry and global pose are estimated with an extended Kalman filter that combines GPS and IMU data. The local planner maintains a 2D rolling-window cost map that tracks obstacles in the local, 30 m by 30 m surroundings of our robot. We use a point-cloud-based method [37] to detect and track obstacles visible in the field of view of the stereo camera. Obstacles are classified as clusters of points above the water plane and tracked using a temporal consistency filter. The tracked obstacles are added or discarded only after consecutive detections for two seconds and used to update the cost map. The sonar is only used for data collection and not for navigation in this work.

C. Testing Site

We tested our algorithm at Nine Mile Lake in McDougall, Ontario, Canada, with five sampling locations of our choice. An overview of the mission is shown in Fig. 5. To reach the target at the bottom-left of the satellite image, the ASV has to disambiguate three consecutive stochastic edges with several large rocks. There are two more stochastic edges,

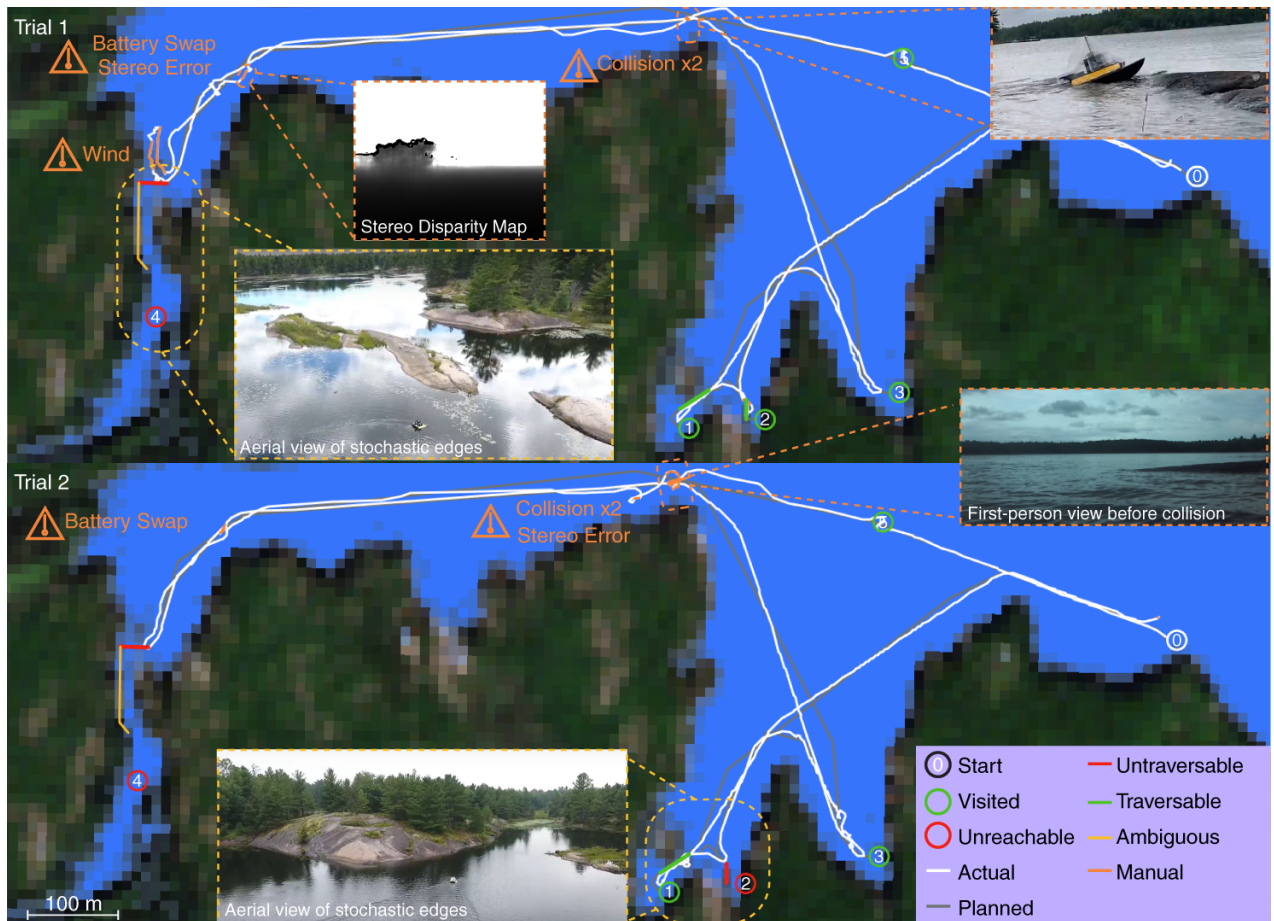


Fig. 5. The GPS path of our robot on satellite water masks for two runs of our field test. The total length traveled autonomously is about 3.7 km. The pictures with a dashed yellow border are drone aerial views of the stochastic edges. We have labeled all instances of manual interventions, including battery changes, failures to detect a rock, and a stereo matching failure in orange. Note that the labels untraversable, traversable, and ambiguous are the state of stochastic edges after disambiguation.

each leading to a target node at the lower bay of the lake. The aerial views of both sections are shown in Fig. 5. The large inconsistency between the satellite image and the local aerial view demonstrates the need for a robust algorithm. Our planner does exactly that by treating these narrow water passages as stochastic edges and plans accordingly.

D. Results

We conducted two trials: first on a rainy day and then on a sunny day. The high-level objective is to navigate the ASV to all reachable target nodes and return home. The robot successfully achieved that in both trials. Manual interventions and their causes are labeled along the path in Fig. 5. Our robot did not find a traversable path to target node 4 because the global plan derived from satellite images was blocked by large rocks (see left of Fig. 5) and the local planner failed to identify an alternative. We observed several other issues during the field tests: power, terrain assessment, and environmental disturbances. The *XP-1* battery must be swapped after about an hour of use. Our stereo matching algorithm struggled against sun glare and very calm water. A shallow rock was misclassified as part of the water plane (see. right of Fig. 5) and consistently resulted in collisions. Lily pads damaged one of our underwater plastic propeller

blades and very strong winds overpowered the motors at one point. Despite these partial system failures and disturbances, the ASV can still execute the full policy, visit other reachable target nodes, and return home. This can be credited to the robustness of our framework, particularly because our policy execution only depends on the robot's global pose and an independent timer.

VI. CONCLUSION

In this paper, we present a route-planning algorithm for ASV navigation using historical satellite images. We propose to model the environmental disturbances as well as inaccuracies of the satellite images by building a stochastic roadmap of the target waterway. The route-planning problem can then be solved offline optimally by a heuristic-based tree-search algorithm. The correctness and utility of our planner are verified in a simulation of over a thousand lakes against two online planning baselines. We tested our algorithm in the field and showed how our planner can be part of a robust ASV navigation system.

ACKNOWLEDGMENT

We thank the Natural Sciences and Engineering Research Council of Canada (NSERC) for partially supporting this work.

REFERENCES

- [1] M. Dunbabin and L. Marques, "Robots for Environmental Monitoring: Significant Advancements and Applications," *IEEE Robot. Autom. Mag.*, vol. 19, no. 1, pp. 24–39, Mar. 2012.
- [2] A. Odetti, G. Bruzzone, M. Altosole, M. Viviani, and M. Caccia, "SWAMP, an Autonomous Surface Vehicle Expressly Designed for Extremely Shallow Waters," *Ocean Eng.*, vol. 216, p. 108205, Nov. 2020.
- [3] G. Ferri, A. Manzi, F. Fornai, F. Ciuchi, and C. Laschi, "The HydroNet ASV, a Small-Sized Autonomous Catamaran for Real-Time Monitoring of Water Quality: From Design to Missions at Sea," *IEEE J. Oceanic Eng.*, vol. 40, no. 3, pp. 710–726, July 2015.
- [4] D. Madeo, A. Pozzebon, C. Moccenni, and D. Bertoni, "A Low-Cost Unmanned Surface Vehicle for Pervasive Water Quality Monitoring," *IEEE Trans. Instrum. Meas.*, vol. 69, no. 4, pp. 1433–1444, Apr. 2020.
- [5] H. Cao, Z. Guo, S. Wang, H. Cheng, and C. Zhan, "Intelligent Wide-Area Water Quality Monitoring and Analysis System Exploiting Unmanned Surface Vehicles and Ensemble Learning," *Water*, vol. 12, no. 3, p. 681, Mar. 2020.
- [6] P. Dash, R. J. Moorhead, J. Herman, W. Beshah, M. S. Sankar, J. Moorhead, G. D. Chesser, W. Lowe, J. Simmerman, G. Turnage, A. P. Katkar, and J. P. Liles, "Evaluation of Water Quality Data Collected using a Novel Autonomous Surface Vessel," in *OCEANS 2021: San Diego – Porto*, Sept. 2021, pp. 1–10.
- [7] Y.-T. Ang, W.-K. Ng, Y.-W. Chong, J. Wan, S.-Y. Chee, and L. B. Firth, "An Autonomous Sailboat for Environment Monitoring," in *2022 Thirteenth International Conference on Ubiquitous and Future Networks (ICUFN)*, July 2022, pp. 242–246.
- [8] S. MahmoudZadeh, A. Abbasi, A. Yazdani, H. Wang, and Y. Liu, "Uninterrupted Path Planning System for Multi-USV Sampling Mission in a Cluttered Ocean Environment," *Ocean Eng.*, vol. 254, p. 111328, June 2022.
- [9] C. H. Papadimitriou and M. Yannakakis, "Shortest Paths Without a Map," *Theoretical Computer Science*, vol. 84, no. 1, pp. 127–150, 1991.
- [10] C. Huang, Y. Chen, S. Zhang, and J. Wu, "Detecting, Extracting, and Monitoring Surface Water From Space Using Optical Sensors: A Review," *Rev. Geophys.*, vol. 56, no. 2, pp. 333–360, June 2018.
- [11] X. Yang, S. Zhao, X. Qin, N. Zhao, and L. Liang, "Mapping of Urban Surface Water Bodies from Sentinel-2 MSI Imagery at 10 m Resolution via NDWI-Based Image Sharpening," *Remote Sensing*, vol. 9, no. 6, p. 596, June 2017.
- [12] J.-F. Pekel, A. Cottam, N. Gorelick, and A. S. Belward, "High-Resolution Mapping of Global Surface Water and its Long-Term Changes," *Nature*, vol. 540, no. 7633, pp. 418–422, Dec. 2016.
- [13] S. K. McFeeters, "The Use of the Normalized Difference Water Index (NDWI) in the Delineation of Open Water Features," *International Journal of Remote Sensing*, vol. 17, pp. 1425–1432, 1996.
- [14] H. Xu, "Modification of Normalised Difference Water Index (NDWI) to Enhance Open Water Features in Remotely Sensed Imagery," *Int. J. Remote Sens.*, vol. 27, no. 14, pp. 3025–3033, July 2006.
- [15] J. Li and Y. Sheng, "An Automated Scheme for Glacial Lake Dynamics mapping using Landsat Imagery and Digital Elevation Models: a Case Study in the Himalayas," *Int. J. Remote Sens.*, vol. 33, no. 16, pp. 5194–5213, Aug. 2012.
- [16] G. L. Feyisa, H. Meilby, R. Fensholt, and S. R. Proud, "Automated Water Extraction Index: A New Technique for Surface Water Mapping Using Landsat Imagery," *Remote Sens. Environ.*, vol. 140, pp. 23–35, Jan. 2014.
- [17] G. Laporte, "The Traveling Salesman Problem: An Overview of Exact and Approximate Algorithms," *European Journal of Operational Research*, vol. 59, pp. 231–247, 1992.
- [18] L. Perron and V. Furnon, "OR-Tools," Google. [Online]. Available: <https://developers.google.com/optimization/>
- [19] R. Bellman, "A Markovian Decision Process," *Journal of mathematics and mechanics*, pp. 679–684, 1957.
- [20] G. H. Polychronopoulos and others, "Stochastic Shortest Path Problems with Recourse," *Networks*.
- [21] V. Aksakalli, O. F. Sahin, and I. Ari, "An AO* Based Exact Algorithm for the Canadian Traveler Problem," *INFORMS Journal on Computing*, vol. 28, no. 1, pp. 96–111, 2016.
- [22] D. Ferguson, A. Stentz, and S. Thrun, "Planning with Pinch Points," Carnegie-Mellon Univ Pittsburgh PA Robotics Inst, Tech. Rep., Jan. 2004.
- [23] H. Guo and T. D. Barfoot, "The Robust Canadian Traveler Problem Applied to Robot Routing," in *2019 International Conference on Robotics and Automation (ICRA)*, May 2019, pp. 5523–5529.
- [24] C.-S. Liao and Y. Huang, "The Covering Canadian Traveller Problem," *Theoretical Computer Science*, vol. 530, pp. 80–88, 2014.
- [25] M. Schiavetti, L. Chen, and R. R. Negenborn, "Survey on Autonomous Surface Vessels: Part I - A New Detailed Definition of Autonomy Levels," in *Computational Logistics*. Springer International Publishing, 2017, pp. 219–233.
- [26] H.-C. Chang, Y.-L. Hsu, S.-S. Hung, G.-R. Ou, J.-R. Wu, and C. Hsu, "Autonomous Water Quality Monitoring and Water Surface Cleaning for Unmanned Surface Vehicle," *Sensors*, vol. 21, no. 4, Feb. 2021.
- [27] S. Bai, T. Shan, F. Chen, L. Liu, and B. Englot, "Information-Driven Path Planning," *Current Robotics Reports*, vol. 2, no. 2, pp. 177–188, Apr. 2021.
- [28] C. L. Chang and J. R. Slagle, "An Admissible and Optimal Algorithm for Searching AND/OR Graphs," *Artif. Intell.*, vol. 2, no. 2, pp. 117–128, Sept. 1971.
- [29] A. Martelli and U. Montanari, "Optimizing Decision Trees Through Heuristically Guided Search," *Commun. ACM*, vol. 21, pp. 1025–1039, 1978.
- [30] C. E. Noon and J. C. Bean, "An Efficient Transformation of the Generalized Traveling Salesman Problem," *INFOR Inf. Syst. Oper. Res.*, vol. 31, no. 1, pp. 39–44, Feb. 1993.
- [31] M. Drusch, U. Del Bello, S. Carlier, O. Colin, V. Fernandez, F. Gascon, B. Hoersch, C. Isola, P. Laberinti, P. Martimort, A. Meygret, F. Spoto, O. Sy, F. Marchese, and P. Bargellini, "Sentinel-2: ESA's Optical High-Resolution Mission for GMES Operational Services," *Remote Sensing of Environment*, vol. 120, pp. 25–36, 2012, the Sentinel Missions - New Opportunities for Science.
- [32] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise," in *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, ser. KDD'96. AAAI Press, 1996, p. 226–231.
- [33] P. E. Hart, N. J. Nilsson, and B. Raphael, "A Formal Basis for the Heuristic Determination of Minimum Cost Paths," *IEEE transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.
- [34] Natural Resources Canada, "Lakes, Rivers and Glaciers in Canada - CanVec Series - Hydrographic Features," 2019.
- [35] N. Christofides, "Worst-Case Analysis of a New Heuristic for the Travelling Salesman Problem," *Operations Research Forum*, vol. 3, 1976.
- [36] D. Fox, W. Burgard, and S. Thrun, "The Dynamic Window Approach to Collision Avoidance," *IEEE Robotics & Automation Magazine*, vol. 4, no. 1, pp. 23–33, 1997.
- [37] J. Muhovic, R. Mandeljc, B. Bovcon, M. Kristan, and J. Pers, "Obstacle Tracking for Unmanned Surface Vessels Using 3-D Point Cloud," *IEEE j. ocean. eng.*, vol. 45, no. 3, pp. 786–798, July 2020.