Multi-Target Rendezvous Search

Malika Meghjani, Sandeep Manjanna and Gregory Dudek

***Abstract*— In this paper, we examine multi-target search, where one or more targets must be found by a moving robot. Given the target's initial probability distribution or the expected search region, we present an analysis of three search strategies - Global maxima search, Local maxima search, and Spiral search. We aim at minimizing the mean-time-to-find and maximizing the total probability of finding the target. This leads to two types of illustrative performance metrics: minimum time capture and guaranteed capture. We validate the search strategies with respect to these two performance metrics. In addition, we study the effect of different target distributions on the performance of the search strategies. We also consider the practical realization of the proposed algorithms for multi-target search. The search strategies are analytically evaluated, through simulations and illustrative deployments, in open-water with an Autonomous Surface Vehicle (ASV) and drifting sensor targets.**

I. INTRODUCTION

This paper addresses the problem of searching for one or more targets for which we either have an initial probability distribution describing their suspected initial location or sparser information such as their initial expected bounding region. We want a searching robot to find lost targets as effectively as possible, but we observe that in this context, being effective can have two distinct meanings: minimizing the time to detect some subset of the targets, and maximizing the likelihood of detecting all the targets. Our analysis shows that the strategies that minimize overall capture time may not be the same as those that guarantee the eventual capture of the targets. To put this in grotesquely concrete terms, if a set of people are adrift in the ocean, it may be more important to find at least one person rapidly than to find all of them too late to assure their survival. On the contrary, if the mission is to look for an important piece of evidence for an investigation, such as a black box from a plane crash, then guaranteed capture of the target becomes more important than the time taken to capture it.

More generally, search problems encompass a large ensemble of application domains and problem formulations, and have many connections to coverage. While search and rescue is the most familiar problem, many other problems such as environmental assessment and threat detection can also be characterized this way. When the search environment is open water, this task becomes even more challenging because of constant drift of the targets, and constraints on the communication range. Since our application focus is marine environments, we ignore the potential impact of obstacles in the environment. Some of the example scenarios for open water search problems are search for lost divers, floating debris from a plane crash, and passively floating scientific data collectors. Our proposed strategies can also be applied in aerial and open terrestrial domains, given the same initial conditions.

The authors are with the Centre for Intelligent Machines, McGill University, Montréal, Québec, Canada.
email:{malika, msandeep, dudek}@cim.mcgill.ca

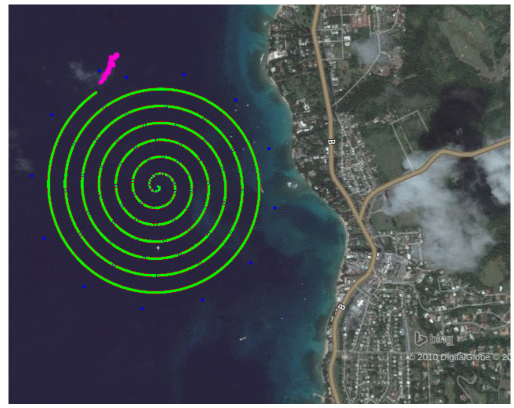

Fig. 1: Simulated search pattern on Microsoft Bing Maps ©. The spiral (in green) indicates the path followed by an ASV terminating at the target waypoints (in pink).

A search problem requires an expected area where the probability of finding the target is high. We refer to this region as the expected search region. Given the probability distribution of the target in a search region or the boundaries of the expected search region, we propose and analyze three search strategies: Global-maxima search, Heuristic local-maxima search and Spiral search. The global and heuristic local-maxima search strategies are dictated by the global and local peaks of the underlying target belief distributions, respectively. The spiral strategy is deterministic, which is useful in finding the targets when there is only a given expected search region with no prior distribution. As discussed below, spiral search has also received substantial prior consideration in the context of the abstract theory of searching. An instance of the spiral search pattern is shown in Fig. 1.

In the spiral search strategy, we consider two search patterns: outward spiral and inward spiral. The outward spiral pattern is a greedy strategy which initializes from the center of the search region and expands outwards to cover the entire region. In contrast, the inward spiral pattern first encapsulates the search area and then moves inwards towards the center. We hypothesize that the inward spiral pattern gives guaranteed search outcomes with a longer search time, whereas the outward spiral pattern, minimizes the search time but compromises on the success-rate.

We analyze the performance of the search strategies by comparing their mean-time-to-find the target in simulation and real field trials. In addition, we report the number of missed targets and propose a cost metric to penalize the search strategies for the missed targets. For field trials, we used an Autonomous Surface Vehicle (ASV, seen in Fig. 8) as the searcher and passively floating data collectors (Drifters, seen in Fig. 9) as the targets.

Contributions of this paper are: parametrization of single target and multi-target search problems; analysis of specific classes of search strategies in terms of performance parameters such as time to find the target, target search failure rate, and cost of the search; theoretical reasoning behind the trade-off between guaranteed capture speed and minimum time capture; and finally, partial validation of the search strategies in open water with a real robot combined with an analysis using our robot simulator.

II. RELATED WORK

Search strategies for robotics applications have often been based on models inspired by natural processes [1], geometric patterns [2] and complete coverage methods [3]. The most common class of search algorithms across these domains are those that generate a spiral pattern, used by animals for foraging food [1], and by rescue robots searching for lost targets [4] [5]. The most fundamental and abstract root problem is that of searching for a point on a line (going back and forth with a specific metric), whereas spiral search provides optimal worst-case performance whenever the target's probability distribution is Gaussian (or comes from a broader class or "realistic" distribution) [6]. In addition, using spiral patterns for global coverage task is also shown to be useful for multi-target search application [7].

The search strategies proposed by Bourgault *et al.* [8], Stone [9], and Furukawa *et al.* [10] represent the target's possible location using Bayesian statistics. Furukawa *et al.* [11] also summarize the mathematical basis of multi-robot, single-target search within a recursive Bayesian framework. Their goal is to unify search and tracking under a single objective function. This allowed them to retain the state estimation of the target when transitioning from tracking to search and vice-versa. Their optimal control strategy for a single searcher, single target problem, also resulted in a spiral search pattern. Since spiral pattern was proven to be dominant in the aforementioned references, we implemented it in this paper.

There are other searching strategies which update the search patterns based on their on-line observations [12]. A similar search problem is addressed by Saigol *et al.* [13]. In this work, the authors propose an automated planning algorithm to deal with uncertainty in searching for hydrothermal vents. They suggest an information lookahead and entropy change planners. The information lookahead is based on a POMDP formulation which was reported to be computationally expensive. Whereas, the entropy change maximization method requires a continuous feedback based on the probabilities of the observing chemical traces. In our work, we do not have such observations available to track and find the target.

Das *et al.* [14] have considered an application very similar to ours and achieve some powerful results. They use a floating drifter to coordinate exploration of a moving ocean patch by an Autonomous Underwater Vehicle (AUV). The floating target demarcates the center of the ocean patch, while the AUV moves in a box pattern around the perimeter, changing depth in a saw-tooth motion. Their work substantially differs from ours in that the target and AUV maintain near constant contact through the use of satellite communication.

In our previous work [15], we illustrated the track and search problem of a floating target using an AUV. The AUV performed spiral patterns using the sawtooth motion. In this work, we provide an analysis for guaranteed capture and minimum time capture along with a performance comparison of the deterministic spiral search strategy with two probabilistic strategies.

III. SEARCH STRATEGIES

We consider the problem of finding a drifting target with a mobile searcher, for example, a robot boat searching for a lost target in the sea. We propose three search strategies given, either an initial target probability distribution or a bounded search region. These search strategies include: Global-Maxima, Heuristic Local-Maxima and Spiral search. The former two strategies are useful when both boundary of the search region and the initial target probability distribution are known, whereas, the latter requires only the knowledge about the boundary of the search region. We summarize each of these strategies, below.

1) Global Maxima Search: Given an initial target probability distribution, the global-maxima search strategy aims at clearing parts of the search region with highest probabilities. The search region is discretized into grids and each grid-cell is assigned with a value equal to the integral of the probability under that grid-cell. Global-maxima search strategy generates a trajectory that visits the grid-cell with highest value and clears that grid-cell once visited, then continues to visit the next maxima until the target is found or the search region is covered. The prior target distribution is application dependent and can be chosen based on the expert knowledge. In our analysis, we observed that with an assumption of Gaussian prior, this strategy provides a trajectory with multiple overlapping segments as the global-maxima shifts across the search region. An example path of the searcher with Gaussian distribution as the prior belief is shown in Fig. 2(a).

2) Heuristic Local Maxima Search: The local-maxima strategy, sequentially clears the grid-cells with maximum values within a local maxima-search radius. The trajectories generated by local maxima tend to be less overlapping as seen in Fig. 2(b). This search strategy is dependent on the initial location of the robot and is prone to getting stuck in the local maxima. The maximum success rate with this strategy is observed to be extremely low. Hence, we introduced a heuristic to overcome this drawback. According

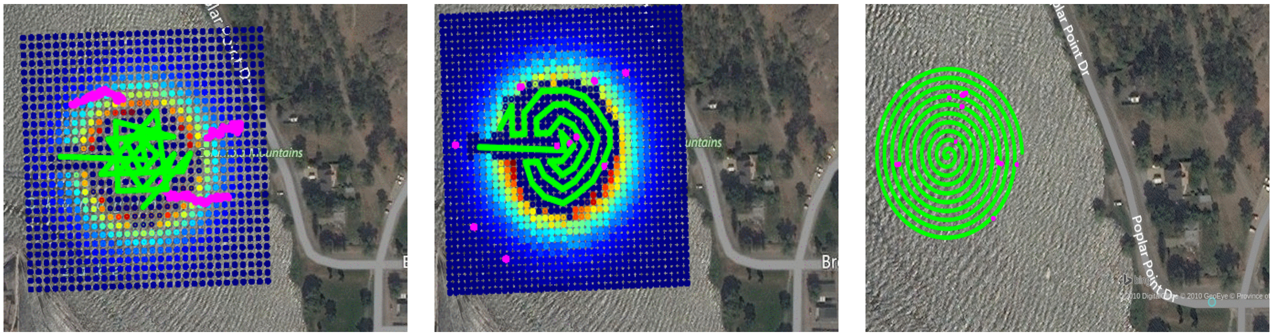

(a) Global Maxima (b) Heuristic Local Maxima (c) Spiral

Fig. 2: Simulated global, heuristic local and spiral search paths of the robot (green) and target waypoints (pink).

to this heuristic, when the searcher is stuck in a local-maxima, we iteratively increase the maxima-search radius, until the searcher recovers from the local maxima or the radius becomes equal to the radius of the entire search region.

3) Spiral Search: The spiral search is a deterministic strategy which is useful when only the boundaries of the search region are known. This strategy does not require the discretization of the search region. The spiral pattern (Fig. 2(c)) dictating the robot's position $(x_r(t), y_r(t))$ at time t is generated using the arithmetic spiral equation given below:

$$\begin{pmatrix} x_r(t) \\ y_r(t) \end{pmatrix} = b\theta \begin{pmatrix} \cos\theta \\ \sin\theta \end{pmatrix} \tag{1}$$

where b is a parameter that determines the distance between two consecutive spiral rounds. In order to assure that the targets are not missed in the spiral width region, while the searcher is following the spiral trajectory, we selected the spiral width to be proportional to the communication radius (R_{comm}), i.e., $b \leq \frac{R_{\text{comm}}}{2\pi}$. This implies that if the spiral width does not satisfy the aforementioned condition then there is no guarantee that the searcher will find the target.

There are two possible variants of the spiral search strategies: *inward*, and *outward* spirals. The inward spiral search encompasses the search region and moves inwards, towards the center of the region, thus minimizing the escape of the targets. Whereas, the outward spiral strategy, starts the search process from the center of the expected target region and expands outwards to minimize the search time for a greedy search.

IV. PERFORMANCE BOUNDS

We formalize our search problem by assuming the target to be a point object and the searcher to be a disk. Alternatively, we can consider the searcher to also be a point but to have a communication radius R_{comm} (defining a disc around the searcher), such that if the target is within the disc then the search terminates successfully. The target is considered to be initially confined within a region of radius r and moving with a constant velocity. Given these conditions, we analyze two performance bounds for guaranteed capture and minimum time capture, corresponding to inward and outward spiral, respectively. For simplicity of the following analysis we consider circular patterns with decreasing and increasing radii to represent inward and outward spiral patterns respectively.

Consider a two dimensional search region which is defined by a probability distribution function with value zero beyond radius r. A robot with maximum speed s_r is required to search for a drifting target with speed s_d initially within the search region. The robot covers the search area uniformly by clearing a band of width $b \leq \frac{R_{\text{comm}}}{2\pi}$ per circular round. The total number of circular rounds that the robot needs to complete for clearing the entire search region of radius r, is given by

$$n_s = \left\lceil \frac{r}{b} \right\rceil \tag{2}$$

The time taken to clear one circular round with radius r', is

$$\tau = \frac{2\pi r'}{s_r} \tag{3}$$

The total time taken by the robot to clear the complete search area is less than or equal to the product of the total number of circular rounds and the time taken to clear each round.

$$\tau_{tot} \leq n_s \tau \tag{4}$$

If the robot performs a circular search pattern with diminishing radius b for each circular round. Then, the total time is calculated as,

$$\tau_{tot} = \frac{2\pi n_s}{s_r} \sum_{i=0}^{(n_s-1)} (r - ib) \tag{5}$$

A. Guaranteed Capture

The worst case scenario for a guaranteed capture of the target is when the target is floating tangentially to the communication disc of the robot while the robot is moving away from the target. In this case, we can still guarantee the capture of the target, only if, the time taken by the robot to complete one circular round is equal to the time taken by the target to cross the width b, that is just before escaping the search region. We then define the capture speed of the robot, s_{cap}, as follows:

$$s_{cap} = \frac{2\pi r s_d}{b} \tag{6}$$

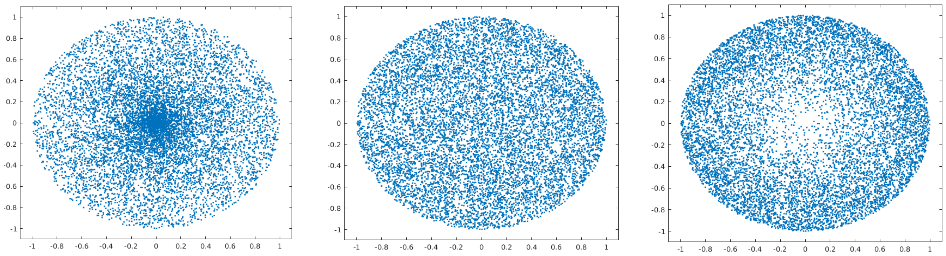

(a) Triangular Samples (b) Uniform Samples (c) V-shaped Samples

Fig. 3: Samples from three distributions on a unit circle centered at zero. We use these distributions to sample drifter (target) locations for simulating multiple trials.

If the speed of the robot is s_r, then we need $s_r > s_{cap}$, for a guaranteed capture. Specifically,

$$s_r > \frac{2\pi r s_d}{b} \tag{7}$$

such that, if the above condition for robot speed s_r is not satisfied, then the target can escape the search region. In addition, given the maximum speed of the robot, the estimated average target speed and the communication radius, we can calculate the search radius for a guaranteed capture using Eq. 6. The guaranteed search radius can then be used as the initial radius for the inward spiral search strategy.

B. Minimum Time Capture

The time to capture the target, from Eq. 5, can be minimized, when, $i = (n_s - 1)$, such that,

$$\tau_{min} = \frac{2\pi n_s}{s_r}(r - (n_s - 1)b) \tag{8}$$

From Eq. 2 we know that, $r = n_s b$, which can be substituted in the above equation to obtain:

$$\tau_{min} = \frac{2\pi b n_s}{s_r} \tag{9}$$

Hence, to minimize the time to capture, the robot should start with an initial radius, $r_{min} = b$ and incrementally expand outwards by a factor b. This gives us a circular pattern with an increasing radius similar to outward spiral pattern. In the following sections we validate the above analytical results for spiral search strategies and compare its performance with global and local maxima search strategies.

V. CONTROLLED SIMULATION

A. Setup

We evaluate the search strategies with our analytical results on a real-time simulator that we developed for our field trials. An instance of the simulator is shown in Fig. 1. We pre-selected a search region of 100 meters radius for our simulations. The maximum speed of the robot is set to 1.2 m/s which is also the maximum speed of the ASV used in our field trials. The target's speed is set to 0.2 m/s based on the drifter data collected in open water during the field trials. The maximum communication range of the robot is simulated to be 5 meters in radius to mimic the degraded WiFi range caused due to poor radio signal transmissions from the drifting target.

The initial locations of the target (drifter) is simulated within the search region based on three spatially representative probability distributions namely, *uniform*, *triangular*, and $|x|$ (*v-shaped*) distributions as shown in Fig. 3. The uniform distribution provides target locations that are unbiased in the search region, whereas the triangular distribution samples the locations that are biased towards the center of the search region, and the v-shaped distribution selects biased locations near the circumference of the search region. We simulate the initial locations of the drifter (x_d, y_d) by randomly sampling the radius $r \in [0, 1]$ and $\theta \in [0, 2\pi)$ from two independent uniform distributions. By sampling r and θ in this way, we obtain a larger concentration of points in the center of the search region, resulting in a triangular distribution (Eq. 11 shows the transformation to Cartesian coordinates). We then transform the sampled points using the following *integration by substitution* rule to obtain the uniform and v-shaped distributions.

$$g(q) = f(p(q)) \left| \frac{dp}{dq} \right| \tag{10}$$

Where, the inverse transformation function $p(q)$ transforms the distribution $f(p)$ into the distribution $g(q)$. In our case, $g(q)$ is uniform and v-shaped distributions. The respective transformations are given in Eq. 12 and 13. The difter's motion is simulated using beta distribution based on our previous work [15].

- Triangular sample

$$(x_d, y_d) = (rcos\theta, rsin\theta) \tag{11}$$

- Uniform sample

$$(x_d, y_d) = (\sqrt{r}cos\theta, \sqrt{r}sin\theta) \tag{12}$$

- v-shaped ($|x|$) sample

$$(x_d, y_d) = (r^{1/3}cos\theta, r^{1/3}sin\theta) \tag{13}$$

The initial location of the searcher is set to a fixed point within the search region for a fair comparison of all the search strategies. For the global-maxima and heuristic local-maxima search strategies, an underlying probability distribution is required to represent the searcher's belief about the

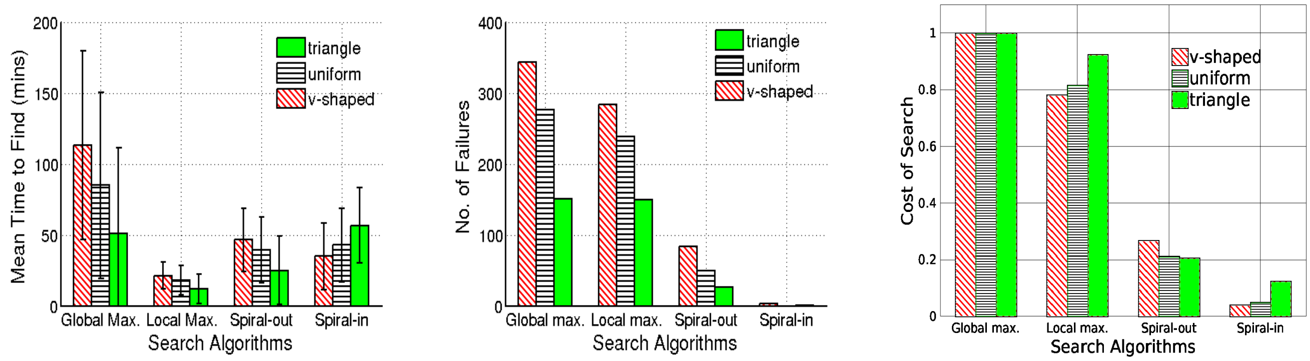

(a) Mean Time to Find

(b) Failure Rate

(c) Cost

Fig. 4: Single target search performance comparison for different target distributions and search strategies.

target. In our experiments, a Gaussian distribution is chosen to represent this belief. For the spiral search strategy, the spiral pattern is pre-generated using Eq. 1.

The following sub-sections present experimental results based on simulations for single-target and multi-target captures. Given the initial target distribution and the communication range of the robot, we recorded the mean time to capture the target and the number of capture failures for 1000 trials for each search strategy. We also propose a cost function which is a combination of the mean time to find and the failure rate.

B. Single Target Search

An instance of the single target capture is illustrated in Fig. 2. The results from the simulation experiments are discussed below.

1) Mean Time to Find (MTTF): The mean-time-to-find the target for the three search strategies is presented in Fig. 4(a). The plot includes a two-sided standard deviation computed over multiple trials of the experiment. The global-maxima search has the highest MTTF due to its overlapping search trajectories. Whereas, the heuristic local-maxima strategy has the least MTTF because of its behavior to clear the local search region and gradually increase the horizon. The inward and outward spirals illustrate an intermediate performance.

It can be observed that the triangular target distribution has the least MTTF for all the search strategies except inward spiral strategy. This is due to the fact that the triangular distribution samples the target locations closer to the center of the search region which also coincides with the mean of the target belief distribution and is the starting point for outward spiral. In contrast, the inward spiral strategy out-performs the outward spiral search strategy when the targets are distributed according to the v-shaped distribution. This is because the target location samples from a v-shaped distribution are closer to the circumference of the search region. Since, the inward spiral strategy starts the search from the circumference of the search region, it follows the reverse trend when compared to other search strategies.

2) Failure Rate: The failure rate follows the same trend as MTTF for the different target probability distributions. However, for the search strategies, it is interesting to note in Fig. 4(b), that the spiral search strategies have the smallest number of failures when compared to global-maxima and heuristic local-maxima strategies. Since the inward spiral strategy guarantees the target capture under certain conditions, (from Section IV), the number of failures for this strategy is close to zero.

3) Cost Analysis: The mean-time-to-find and failure rate, define the performance of search strategies independently. However, for a comprehensive performance score, we combined these two factors to formulate a cost (C), that is a function of time to find the target (t_i) and time spent on missed targets.

$$C = \sum_{i=1}^{n} (\mathbf{1}_{\mathbf{Found_i}} * t_i + \mathbf{1}_{\mathbf{Found_i^C}} * \beta) \tag{14}$$

Where,

$$\mathbf{1}_{\mathbf{Found_i}} = \begin{cases} 1 & \text{if target } i \text{ is found} \\ 0 & \text{if target } i \text{ is not found,} \end{cases}$$

$$\mathbf{1}_{\mathbf{Found_i^C}} = \mathbf{1} - \mathbf{1}_{\mathbf{Found_i}} \tag{15}$$

In Eq. 14, n is the total number of trials, β is the penalizing factor for missing the target. $\beta = t_{max} + \epsilon$, where t_{max} is the maximum time spent by search strategies on missed targets, and $\epsilon > 0$ assures that the failures are more expensive than the successes. The value of ϵ is application specific and for our application we have chosen $\epsilon = t_{max}$ to penalize capture failures. In field experiments, t_{max} could be the battery-time of the search vehicle.

A comparison between the normalized cost for three search strategies for the three target distributions is presented in Fig. 4(c). We observe that the global maxima search strategy has the highest cost, implying that it has the worst combination of search time and success rate. The next strategy is the heuristic local maxima which has the shortest time to find but with a very high failure rate. The outward and inward spiral search strategies achieve a better trade-off between the mean-time-to-find the target and the success rate.

(a) Global Maxima (b) Heuristic Local Maxima (c) Spiral

Fig. 5: Multi target search path using the proposed search strategies.

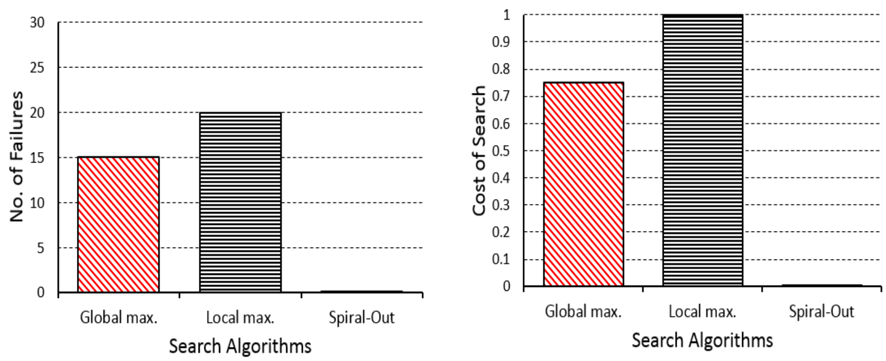

(a) Number of Failures for search strategies (b) Cost of search for search strategies

Fig. 6: Multi-target number of failures and cost comparison for different search strategies.

C. Multi-Target Search

The most critical aspect of multi-target search is to decide when to switch between the search regions for different targets. This decision process naturally evolves for the global-maxima and heuristic local-maxima search strategies. However, for the spiral strategy, the switching time between the target search regions need to be decided explicitly. In our experiments for spiral strategy, we use regular interval switching times. This switching method guarantees that equal effort is spent in searching all the targets and covering higher probability regions in the initial search period.

We simulated 4 targets for 100 trials within the search region of radius 200 meters. The target locations are sampled from triangular distribution using Eq. 11. The searcher's belief distribution is represented by 4 Gaussian distributions with means at the centers of the target distributions. The searcher's initial location was fixed for all trials for consistency. Sample search paths executed by the robot for the three search strategies are illustrated in Fig. 5. It can be observed that the global maxima search executes repetitive overlapping paths while transitioning from one target to another. The heuristic local maxima strategy covers the target regions exhaustively, by slowly expanding its horizon, thus taking more time to find the targets. The outward spiral strategy performs structured transitions between the search regions of the targets while initially covering the higher probability regions.

The number of failures for each of the search strategies is presented in Fig. 6(a). These results are reported for 100 trials. The heuristic local search strategy has the highest failure rate (20%). This can be explained by the sequential nature of the local search strategy which allows the targets to escape from the initial search region. The global search strategy is next in order for failure rate (15%) as it generates longer and repetitive trajectories, thus losing the targets. The spiral strategy has the best performance with 100% success. The cost analysis of the search strategies follow the same trend as the failure rate and is shown in Fig. 6(b).

VI. FIELD TRIALS

A. Setup

We evaluated our proposed search strategies in Lake Okanagan, Canada over a search radius of 50 meters. The searcher robot used for field experiments is an Autonomous Surface Vehicle (ASV), Kingfisher from Clearpath Robotics. The search targets used are in-house designed floating data collectors referred to as the *drifters*. The maximum operating speed of the ASV is 1.2 m/s. The average drifter speed under natural wind and current conditions is measured to be 0.2 m/s. The battery life of the ASV is 3 hours. On ASV, we used a simple PID controller to achieve smooth trajectories following a set of waypoints. This setup was done 50 meters

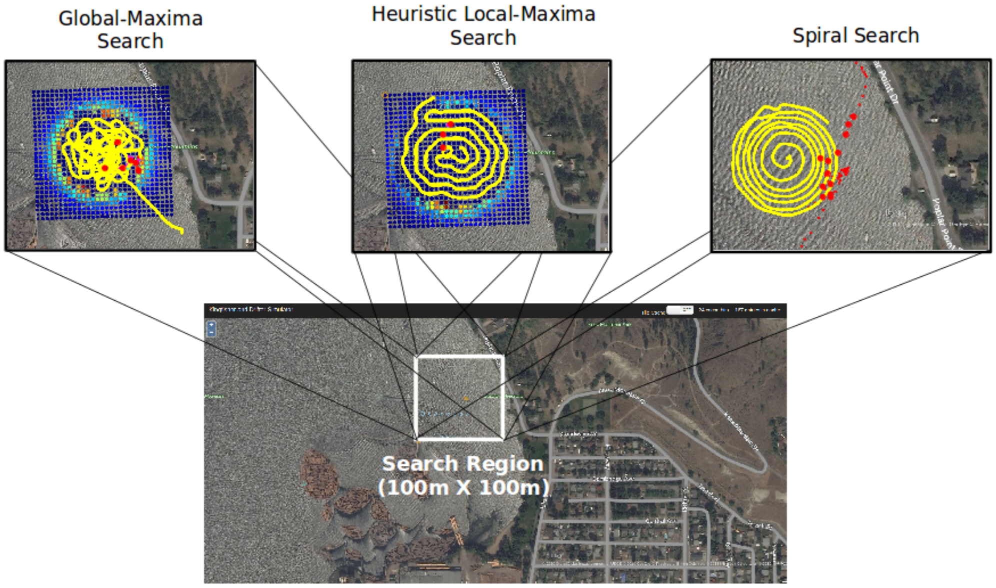

Fig. 7: The search region where field experiments were conducted and the paths generated by three search strategies : Global Maxima, Heuristic Local Maxima, and Spiral Search. Boat's search path is shown in yellow and drifter's discrete locations is shown with red dots.

away from the shore where the wind was calm but the currents were moderate due to the disturbances from nearby docking ships. An illustration of our deployment for the search experiment is presented in Fig. 8, and the deployment area is shown in Fig. 7. During the field experiments, the maximum velocity of the boat, measured drifter velocity, and the size of the search area did not satisfy the conditions for guaranteed capture of the target (Eq. 6). Hence, we could not validate the inward spiral search strategy on the real robot.

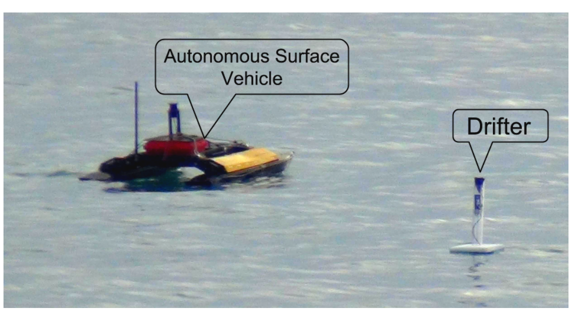

Fig. 8: Autonomous Surface Vehicle, Kingfisher and search target, drifter.

1) Autonomous Surface Vehicle: The searcher robot, Kingfisher, is a catamaran style, differential drive boat. It has an on-board computer with 2.4 GHz Wireless interface to communicate with the base station and the drifters.

2) Drifter: The drifters are equipped with a miniPC (Android MK-802), a GPS receiver (Adafruit) and a 5V/2A battery pack as shown in the Fig. 9. The miniPC is capable of storing onboard data and communicating with the searcher via its in-built WiFi antenna. The drifters are designed to be neutrally buoyant such that they floated upright while receiving WiFi and GPS signals without interference.

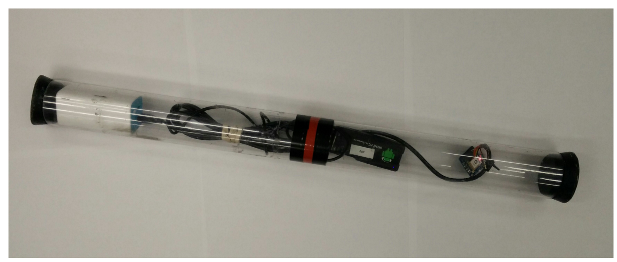

Fig. 9: Search target: the Drifter.

B. Results

The result of the field trials are presented in Fig. 7 with the illustration of ASV's search path and the drifters' waypoints. For the global and heuristic local maxima search strategies, the local and global maxima of the prior belief, Gaussian distribution, is followed by the ASV and the distribution is updated in real-time. For the spiral search strategy, the spiral pattern is pre-generated according to Eq. 1.

In order to get a complete and extensive comparison of the three search strategies, we used the actual search path of the ASV as the input to our real-time simulator. The drifter's initial location was simulated using the three probability distributions: uniform, triangular and v-shaped. The drifter's trajectory was again generated using the beta distribution. We executed 500 trials per search strategy and recorded the mean-time-to-find, success rate and the cost. These results are presented in Fig. 10. Once again, we observe that the outward spiral strategy has the least number of failures, followed by global-maxima, and heuristic local-maxima. The cost analysis in Fig. 10(b) shows that the spiral strategy performs very well with the drifters sampled from a triangular distribution than those sampled from v-

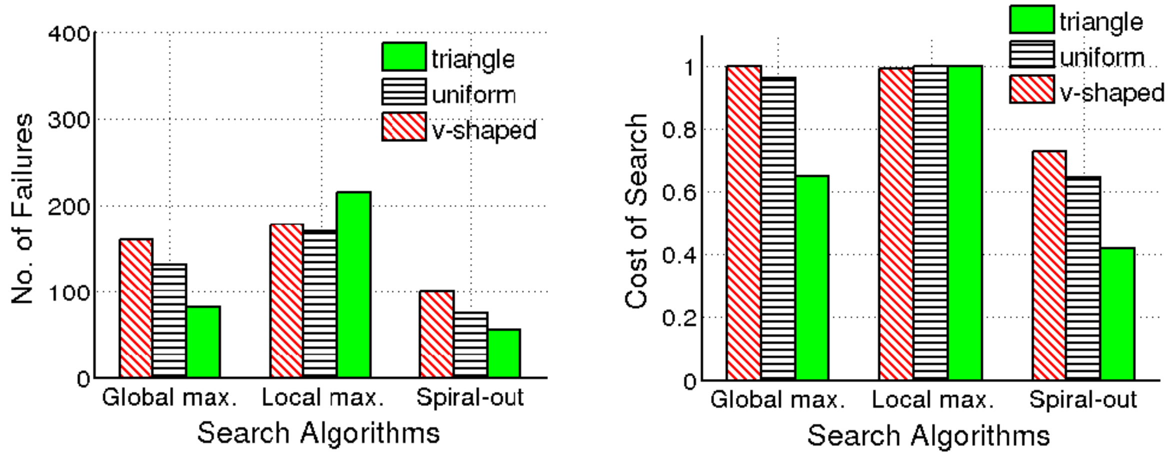

Fig. 10: Comparing the number of failures and cost for lake trials.

shaped distribution. This can be attributed to the fact that the outward spiral strategy covers the center of the search region first allowing the drifters sampled from v-distribution to escape. Nevertheless, the spiral strategy performs better than any other search strategy with all three target sampling distributions.

VII. CONCLUSIONS

In this paper, we present an analysis of three search strategies and compare their performance in terms of their success rate and the cost of search. We found the outward spiral strategy to consistently outperform the global-maxima and heuristic local-maxima strategies for both single-target and multi-target experiments in simulations and open water experiments. The analytical results for spiral search strategy is shown to guarantee the target capture and minimize the capture time. The corresponding simulation results validate our hypothesis that inward spiral patterns provide guaranteed capture, whereas outward spirals minimize the capture time.

For future work, we would like to enhance the performance of the spiral search strategy for multiple targets. We plan to do this by designing a decision making process that can choose the time to transit between the targets instead of our current fixed time switching strategy.

VIII. ACKNOWLEDGEMENT

This work was supported by the Natural Sciences and Engineering Research Council (NSERC) through the NSERC Canadian Field Robotics Network (NCFRN). The authors would specially like to acknowledge all the members of the Mobile Robotics Lab at McGill University for their technical discussions and motivation.

REFERENCES

[1] E. Gelenbe, N. Schmajuk, J. Staddon, and J. Reif. Autonomous search by robots and animals: A survey. *Robotics and Autonomous Systems*, 22(1):23–34, 1997.

[2] S. Alpern and S. Gal. The theory of search games and rendezvous. pages 165–178, 2003.

[3] I. Rekleitis, A.P. New, E.S. Rankin, and H. Choset. Efficient boustrophedon multi-robot coverage: an algorithmic approach. *Annals of Mathematics and Artificial Intelligence*, 52(2):109–142, 2008.

[4] S. Burlington and G. Dudek. Spiral search as an efficient mobile robotic search technique. In *Proceedings of the 16th National Conf. on AI, Orlando Fl.* Citeseer, 1999.

[5] Gregory Dudek and Michael Jenkin. *Computational principles of mobile robotics*. Cambridge university press, 2010.

[6] Ricardo A Baezayates, Joseph C Culberson, and Gregory JE Rawlins. Searching in the plane. *Information and computation*, 106(2):234–252, 1993.

[7] E González, P Aristizábal, and M Alarcón. Backtracking spiral algorithm: a mobile robot region filling strategy. In *Proceeding of the 2002 international symposium on robotics and automation, Toluca, Mexico*, pages 261–266, 2002.

[8] Frédéric Bourgault, Tomonari Furukawa, and Hugh F Durrant-Whyte. Optimal search for a lost target in a bayesian world. In *Field and service robotics*, pages 209–222. Springer, 2006.

[9] L.D. Stone. *Theory of Optimal Search*. Computational Methods for Modeling of Nonlinear Systems. Academic Press, 1975.

[10] Tomonari Furukawa, Frederic Bourgault, Benjamin Lavis, and Hugh F Durrant-Whyte. Recursive bayesian search-and-tracking using coordinated uavs for lost targets. In *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*, pages 2521–2526. IEEE, 2006.

[11] Tomonari Furukawa, Lin Chi Mak, Hugh Durrant-Whyte, and Rajmohan Madhavan. Autonomous bayesian search and tracking, and its experimental validation. *Advanced Robotics*, 26(5-6):461–485, 2012.

[12] Pablo Lanillos, Eva Besada-Portas, Jose Antonio Lopez-Orozco, and Jesus Manuel de la Cruz. Minimum time search in uncertain dynamic domains with complex sensorial platforms. *Sensors*, 14(8):14131–14179, 2014.

[13] Zeyn A Saigol. *Automated planning for hydrothermal vent prospecting using AUVs*. PhD thesis, University of Birmingham, 2011.

[14] Jnaneshwar Das, Frédéric Py, Thom Maughan, Tom OReilly, Monique Messié, John Ryan, Gaurav S Sukhatme, and Kanna Rajan. Coordinated sampling of dynamic oceanographic features with underwater vehicles and drifters. *The International Journal of Robotics Research*, 31(5):626–646, 2012.

[15] Malika Meghjani, Florian Shkurti, Juan Camilo Gamboa Higuera, Arnold Kalmbach, David Whitney, and Gregory Dudek. Asymmetric rendezvous search at sea. In *Computer and Robot Vision (CRV), 2014 Canadian Conference on*, pages 175–180. IEEE, 2014.