Optimal Path Planning for Radiation Detection with Autonomous Underwater Gliders

Verónica Mendes Pedro

***Abstract*—Information on radioactivity in marine environments is sparse, lacking detailed data on source locations, intensities, and their impact on marine ecosystems. This work focuses on optimizing the trajectories of two underwater gliders to map these radiation sources, enabling future in-depth studies. The developed method uses an optimal control approach to find the most efficient path with minimal resource consumption. The approach involves a two-step path planning process: an initial general sweep without prior information aimed at covering the largest possible area with the least traveled distance, followed by a finer targeted sweep over potential locations to verify and refine the results.**

To test the initial step, both manually specified and optimized trajectories were assessed, with the latter outperforming the former. A turning angle limit was also tested and produced similar results. For very long trajectories, segment-based optimization was evaluated and the outcome was favorable. The simulations with two gliders achieved the expected result: they covered the area in a complementary manner. A map split between very high and very low probability of target location was considered. The vehicles behaved as expected, only sweeping the area with the highest probability.

For the targeted coverage tests, a normal distribution was used as a target model. When there was only one radiation source, all trajectories converged to it. When two targets were considered, in most cases the vehicles explored both sources, as desired.

***Index Terms*—coverage search, underwater glider, optimal control.**

I. Introduction

MARINE radioactivity has been present since the Earth's formation. This kind of naturally-occurring radioactivity usually poses a low, but long-term risk in ecosystems. In addition, since the last century, artificial radionuclides have been discharged into the oceans from various sources [1]. So, artificial radioactivity can cause more significant risks for marine environments and public health.

Currently, marine radioactivity monitoring is still very understudied and limited to very few sites. The H2020 EU project RAMONES - Radioactivity Monitoring in Ocean Ecosystems has the objective of long-term, continuous and *in situ* monitoring of radioactivity in the oceans. Hence, a cooperative system of radiological instruments constituted by a static benthic radiological laboratory (deployed on the seabed), two autonomous underwater gliders, and an autonomous surface craft are envisioned to autonomously and adaptively monitor radiation in underwater environments [2].

This work focuses on trajectory planning of the two underwater gliders for radiation sources detection that are in unknown locations. In order to do so efficiently, the path-planning problem is usually done in two steps: a general sweep, without prior information, that is faster and has less sensor precision. This constitutes a "coverage" problem, for which it is is important to sweep the largest possible area in the least amount of time. Then, once a potential location is found, it is necessary to do a finer sweep (by passing several times over the potential location at different directions and at lower speed) in order to check and possibly to rectify the results. This finer sweep can be seen as a targeted coverage problem, for which an optimized trajectory can also be found.

A. Underwater Gliders

Underwater gliders are a class of autonomous underwater vehicles, shaped like torpedos, with a length of approximately 2 meters and weigh about 50 kilograms. These are capable of very long missions (several months) and of traveling thousands of kilometers in a single deployment. In contrast to other autonomous underwater vehicles (AUVs), gliders are not propulsion-driven but buoyancy driven: the power needed to overcome the drag on the moving vehicle is provided by gravity in the form of positive or negative net buoyancy. As horizontal movement is only possible if the trajectory is inclined in the direction of vertical forces of gravity, the glider is equipped with wings that create a vertical lift force that is perpendicular to the trajectory and allows to convert buoyancy into horizontal movement.

The glider's heading can be controlled essentially in two ways: by deflecting a rudder or by rotating an eccentric mass on the longitudinal axis.

II. Coverage problem

According to Galceran and Carreras in [3], the coverage problem consist on passing all through points over an area or volume of interest while avoiding obstacles. Cao et al. [4] defined criteria for an ideal trajectory that performs a coverage operation:

- The vehicle must go through all points in an area;
- The vehicle must not have overlapping paths;
- Continuous and sequential operation without any repetition of paths is required;
- All obstacles must be avoided;
- Simple trajectories must be used;
- An "optimal" path is desired under the problem conditions.

In complex environments, it is not possible to meet all criteria. Therefore, it is necessary to prioritize the criteria according to the case at hand. If the AUV does not pass through all points in the area, the first criterion is not met and the coverage is incomplete. Alternatively, if it is possible to cover the whole area, the problem is said to be complete.

A. Glider Model

The position of the glider is the center of buoyancy, $\mathbf{x} = [x_1 \ x_2]^T$. The vehicle moves with translational velocity, V, relative to the inertial frame and the angular velocity is Θ, both expressed in the body frame. The AUV yaw angle is θ. The kinematic model can be described by

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} V cos(\theta) \\ V sin(\theta) \end{bmatrix}. \tag{1}$$

In theory, the control variable u would be the glider turning angles θ throughout the whole mission. However, as the glider has a relatively slow turning speed and the problem is discretized, the control variable corresponds to the integral between two time instants, t_1 and t_2, of the yaw angular velocity $\mathbf{\Theta}$,

$$u = \int_{t_1}^{t_2} \Theta(t) dt. \tag{2}$$

B. Sensor Model

The sensor model that is considered aims to detect nearly all radiation against the expected level of ambient noise within its detection radius r. Beyond the radius, the probability of detection gradually decreases to zero. A sensor radius of 1 m was used in all simulations. In order to ensure compatibility with MATLAB solvers, the probability of detection should not have derivatives with high norms due to the challenges coming from steep functions, not to mention from discontinuities. Therefore, the probability of target detection by the sensor (P) in a given time interval, which is the sensor instantaneous detection rate ($\gamma(t)$) , is described by the sigmoid function

$$P = \frac{e^{-Ax+Ar+5}}{e^{-Ax+Ar+5}+1}, \tag{3}$$

where A is a constant, set to satisfy the previous sensor model conditions, and d is the distance between the location of the glider and a possible target location, $\boldsymbol{\omega}$. Several values of A were tested and A = 2 was chosen.

C. Target Models

The location of radiation sources is an uncertain parameter. Therefore, the target location, $\boldsymbol{\omega} \in \mathbf{\Omega}$, is assumed as random variable distributed over a designated search area, $\mathbf{\Omega} \subset \mathbb{R}^2$, according to a known probability density function $\phi(\boldsymbol{\omega}) : \mathbf{\Omega} \rightarrow \mathbb{R}$. Two target location distribution models are used to simulate two different possible situations:

- Uniform Prior: where there is no prior knowledge of the area, and, therefore, a uniform probability density function is used. For a generic map of length A and width B, $A, B \in \mathbb{R}$,
$$\phi(\boldsymbol{\omega}) = \frac{1}{AB}. \tag{4}$$
- 80/20 Prior: here, the probability density function divides the map in two: half the map has 20% probability of target location and the other half has 80%. This prior is useful if, for some reason, a part of the considered map is not significant enough to be explored. Some of these reasons might be because that low probability area has already been previously explored or because it contains hazards for the glider. An adjustment had to be made in order to make a smoother transition between the two distinct areas, as the optimization solver has trouble dealing with very steep or discontinuous functions. Thus, for $\boldsymbol{\omega} = (x, y)$ and for a generic map of length A and width B, $A, B \in \mathbb{R}$, the expression that describes these parameters is
$$\phi(\boldsymbol{\omega}) = \frac{1}{1 + 20e^{-(x+y-A)20+(x-y)^2}}. \tag{5}$$

D. Problem Definition and Discretization

1) Initial Guess: In order to initialize the numeric routine, an initial guess is usually necessary, which is typically a candidate solution. The solver then evaluates the objective function using this trajectory, and then it iteratively generates new candidate solutions that decrease the cost functional value, until it reaches a local minimum. Hence, a good initial guess is essential to decrease solution times [5].

2) Vehicle Model and Constraints: The AUV model used is the glider kinematics model, presented in II-A. This allows to know how the states evolve with time. However, to keep the notation simple, here the previously mentioned kinematics will be described with the general function h,

$$\dot{x} = h(x(t), u(t)), \tag{6}$$

for the mission timeframe $t \in [0, T_F]$ and with the initial condition $x(0) = \xi$.

Other constraints are related to physical actuator limitations, which affect the control inputs:

$$\boldsymbol{u}_{min} < \boldsymbol{u}(t) < \boldsymbol{u}_{max}, \forall t \in [0, T_F], \tag{7}$$

and to operational constraints that limit the state space explored:

$$\boldsymbol{x}_{min} < \boldsymbol{x}(t) < \boldsymbol{x}_{max}, \forall t \in [0, T_F]. \tag{8}$$

3) Problem Definition: In order to obtain the cost functional, the exponential detection model is used to quantify the search performance in continuous time. In this context, this quantity to be quantified and minimized is the residual risk: the probability that, by the end of the mission, the glider fails to detect all targets present in a map. S. Kragelund et al. [5] use this method for mine-hunting expeditions, and it is described in the following section. So as to get the cost functional, it is first necessary to get an instantaneous detection rate, $\gamma(t)$, as a function of the AUV trajectory, $\boldsymbol{x(t)}$, and of the target location, $\boldsymbol{\omega}$. $\gamma(t)$ used in this work is the sensor probability of target detection in (3). The location of the target was also defined in the previous section as a random variable distributed over a map, with a probability density function that is chosen depending on the situation at hand.

Having already defined the necessary models (glider movement, sensor and target models), the next step is to define

the objective function. For this purpose, two assumptions are introduced:

- The probability of detection in the short time interval $[t, t+\Delta t]$ is $\gamma(t)\Delta t$;
- Detection events in all non-overlapping time intervals are independent.

If $p(t)$ is the probability of detection at time t, and is $q(t) = 1 - p(t)$ the probability of detection failure, then under the previous assumptions it follows that

$$q(t+\Delta t) = q(t)(1-\gamma(t)\Delta t), \tag{9}$$

which leads to

$$\frac{q(t+\Delta t) - q(t)}{\Delta t} = -q(t)\gamma(t). \tag{10}$$

By taking the limit $\Delta t \to 0$, (10) becomes the differential equation

$$\dot{q}(t) = -q(t)\gamma(t) \tag{11}$$

that has the closed-form solution

$$q(t) = e^{-\int_0^t \gamma(\tau)d\tau}. \tag{12}$$

This allows to derive the exponential detection model

$$p(t) = 1 - q(t) = 1 - e^{-\int_0^t \gamma(\tau)d\tau}. \tag{13}$$

As, at the moment, the probability of target non-detection is defined in (12), it is possible to obtain the residual risk of failure to find the target at the end of a mission with duration T_F with

$$q(T_F) = e^{-\int_0^{T_F} \gamma(\tau)d\tau}, \tag{14}$$

the minimization of which is the aim of the optimal search.

As $\gamma(t)$ depends on the AUV trajectory and on the target location, which is a random variable, the result of (14) is also a random variable, which cannot be explicitly minimized. Therefore, the cost functional to minimize is its expected value

$$J = \mathbb{E}[q(T_F)] = \int_\Omega e^{-\int_0^{T_F} \gamma(\boldsymbol{x}(\tau),\boldsymbol{\omega})d\tau}\phi(\boldsymbol{\omega})d\boldsymbol{\omega}, \tag{15}$$

where Ω is the search area.

In order to adapt the functional in (15) for multiple search vehicles, n^v with $v = 2, 3, ...$, it is important to assume that their measurements are independent, and so

$$q^{n_v}(T_F) = \prod_{k=1}^{n_v} e^{-\int_0^{T_F} \gamma(\boldsymbol{x}^k(\tau),\boldsymbol{\omega})d\tau} = e^{-\int_0^{T_F} \Gamma(\tau,\boldsymbol{\omega})d\tau}, \tag{16}$$

where

$$\Gamma(t,\boldsymbol{\omega}) = \sum_{k=1}^{n_v} \gamma^k((\boldsymbol{x}^k(\tau),\boldsymbol{\omega}). \tag{17}$$

Therefore, as shown in [5],

$$J^{n_v} = \mathbb{E}[q^{n_v}(T_F)] = \int_\Omega e^{-\int_0^{T_F} \Gamma(\tau,\boldsymbol{\omega})d\tau}\phi(\boldsymbol{\omega})d\boldsymbol{\omega} \tag{18}$$

Now, it is possible to define the optimization problem at hand as

$$\min_{(\xi,u)} J(\xi, u) \tag{19}$$

subject to

$$\boldsymbol{u}_{min} < \boldsymbol{u}(t) < \boldsymbol{u}_{max}, \forall t \in [0, T_F] \tag{20}$$

and

$$\boldsymbol{x}_{min} < \boldsymbol{x}(t) < \boldsymbol{x}_{max}, \forall t \in [0, T_F], \tag{21}$$

where ξ is the initial position $x(0) = \xi$.

4) Problem Discretization: One significant concern is the feasibility of the solution, as most optimal control problems cannot be solved through analytical methods. Therefore, numerical methods are needed, often using a discretized version of the original problem. The aforementioned approach is feasible as long as it satisfies all the problem constraints in a finite number of nodes. This means that all the solutions with discrete trajectories are feasible if a vehicle can execute a smooth control trajectory $\boldsymbol{u}(t)$, generated by the discrete solution $\boldsymbol{u}(k)$, and produce a state trajectory $\boldsymbol{v}(t)$ sufficiently close to the discrete solution's $\boldsymbol{s}(k)$.

In order to discretize the problem, it is first necessary to define a parametric "information state" [6] [7] that represents the probability of target detection up to a time t,

$$\dot{z}(s;a) = \gamma(x(s), y(s;\omega)), \tag{22}$$

where the normalized time-frame is $t \in [0,1]$, $\omega \in \Omega$ and $\{x(s) : 0 \le s \le t\}$.

Then, using Euler's method, it is possible to describe the state model in (6) as

$$x_N^\eta((j+1)/N) - x_N^\eta(j/N) = \frac{1}{N}h(x_N^\eta(j/N), u^\eta(j/N)), \tag{23}$$

where $\eta = (\xi, u)$, $N \in \mathbb{N}$, $x_N^\eta(0) = \xi$ and for $j = 0, 1, ..., N-1$.

It is now possible to discretize (22) by Euler's method

$$z_N^\eta((j+1)/N;\omega) - z_N^\eta(j/N;\omega) = \frac{1}{N}\gamma(x_N^\eta(j/N),\omega), \tag{24}$$

for $z_N^\eta(0,\omega) = 0$. This parametric information state can be adapted to L AUVs

$$z_N^\eta((j+1)/N;\omega) - z_N^\eta(j/N;\omega) = \frac{1}{N}\sum_{l=1}^{L}\gamma^l(x_N^{\eta,l}(j/N),\omega). \tag{25}$$

Now, it is essential to introduce a numerical integration scheme, I_M. Since the search space $\Omega \in \mathbb{R}^2$ is two-dimensional, the corresponding discretization parameter is also two-dimensional: $M = (M_1, M_2)^T \in \mathbb{N}x\mathbb{N}$, for a function

$\Psi : \Omega \to \mathbb{R}^2$, with $\Psi \in C^p(\Omega, \mathbb{R})$, meaning that Ψ is p-times continuously differentiable from Ψ to $\mathbb{R}$. Therefore, I_M is a mapping from $C^p(\Omega, \mathbb{R})$ to $\mathbb{R}$

$$I_M(\Psi) := \sum_{i=1}^{M_1} \sum_{j=1}^{M_2} W_{ij} \Psi(\omega ij), \tag{26}$$

where $W_{ij} \in (R)$ are weights.

Now, it is possible to write the cost functional (15) as its discrete approximation

$$J = I_M(e^{-z_N(1;.)\phi(.)}). \tag{27}$$

E. Simulation Parameters

The following sections elaborate on aspects made different or chosen for the simulations.

1) Initial Guess: The goal of the optimization is to get the "longest" trajectory possible, or the trajectory that sweeps the maximum map area. With that in mind, the initial condition chosen was the exact opposite: a small geometrical shape (in this case a hexagon) that has the whole chosen length of the trajectory "wrapped" in its edges. In this way, the algorithm is forced to work with this difficult and very inefficient initial condition, which is a useful first evaluation of the optimized path: if the optimized trajectory is fully "unraveled", then it is obviously better than the initial condition.

2) Simulation Constraints: In the following simulations two main constraints were considered, that were described previously in a general manner.

- Map size: the constraint described in (21) is key to the vehicle not leaving the stipulated map area. This is used across all simulations.
- Turning angle limitations: this constraint mimics the real glider actuator limitations, as in (20), that limits the AUV turning angle in order not to exceed a module of 30º.

3) Integration Scheme: The integration scheme in (26) is the two-dimensional Simpson's rule.

4) Turning Angle Penalization: As gliders move at slow speeds, their angular velocities are also low. Therefore, a new parameter was added to the cost functional

$$J_{new} = J + C(\sum_{i=1}^{N} u_i)^2, \tag{28}$$

where C is a constant, whose value is studied later on. This new factor penalizes significant turning angle variations. Even though this added penalization is important, it always increases very significantly the cost value. This phenomenon entails that the best trajectory might not be the one with the lower cost functional, when comparing different values of C.

5) Energy Units (EU): One of the main concerns in a coverage search mission is energy efficiency, due to its direct relation to mission length. To this goal, an energy parameter is added. The power consumption generated by the horizontal rudder is proportional to the absolute value of the rudder turning angle in one time-step [8], therefore, the energy spent during one trajectory is proportional to the integral of the power during the mission length,

$$Energy \propto \int_0^{T_f} \dot{\Theta}(t) dt. \tag{29}$$

As this is only a proportional value and not a value of energy itself, in the future this metric is going to be designated generally as "energy units", EU.

6) Manually Specified Trajectories: These trajectories are common deterministic search patterns which are usually used as a means of comparison when developing an optimized path. The trajectories used were the lawnmower and the boxspiral, which are represented in Fig. 1.

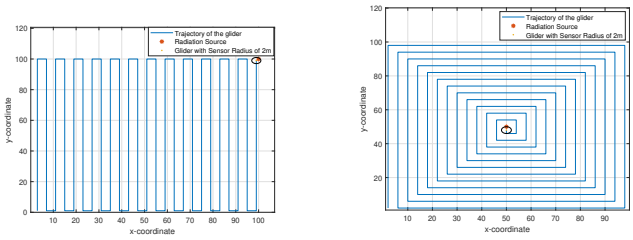

((a)) Lawnmower **((b))** Boxspiral

Fig. 1: Manually Specified Trajectories

F. Results: Uniform Prior

During the following simulations, the mission length was 120 seconds with a linear velocity of 20 cm/s. The time-step considered was 2.5 seconds and, therefore, the complete trajectory was 220 steps long. The map used was $41m \times 41m$.

1) One Glider: The following simulations are the result of the optimization algorithm applied. As described in (28), an angle variation penalization parameter (C) was added to the cost functional. In Fig. 2, the trajectories for several values of C can be seen. These have a colorful background that shows the probability of the sensor to detect a radiation source, if there is one.

Another way to approximate further the simulations to the reality of AUV movement is to add the optimization constraint that limits the control variable, u, to $\pm 30°$. The resulting trajectories can be seen for several values of C in Fig. 3.

Some observations about the previous simulations can be made. Firstly, some manually-specified patterns were tested: a boxspiral, a narrow lawnmower, and a wider lawnmower. As in optimized trajectories the increase of C significantly increases cost value due to the turning angle penalization, the optimized trajectories (with and without control variable) are only comparable to the manually specified paths when C = 0. This evaluation showed that paths computed with the optimization algorithm have a lower final cost functional value than the lowest-valued trajectory, the narrow lawnmower. In terms of energy spent, the boxspiral is slightly more energy efficient than the optimized trajectory that requires the least energy (when C = 0.01). However, by analyzing the cost-benefit of both the energy efficiency and the cost functional

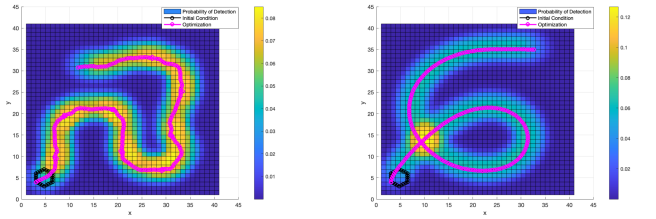

((a)) C = 0 ((b)) C = 0.0005

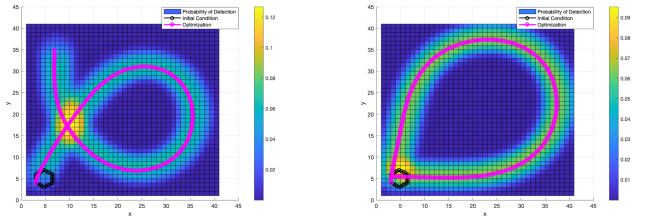

((c)) C = 0.001 ((d)) C = 0.01

Fig. 2: Optimized trajectories for 1 AUV with a uniform prior probability density function.

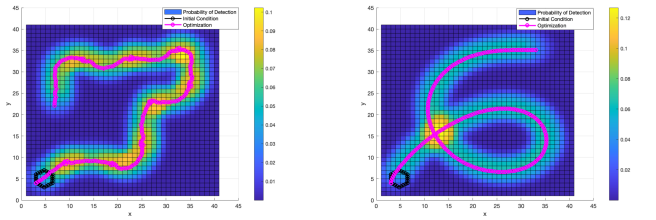

((a)) C = 0 ((b)) C = 0.0005

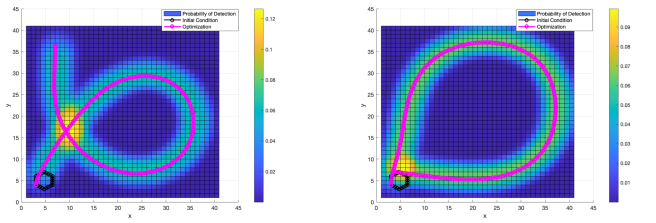

((c)) C = 0.001 ((d)) C = 0.005

Fig. 3: Optimized trajectories with control variable constraint for 1 AUV with a uniform prior probability density function.

value, the optimized trajectories have better performance: even with the added value of the highest cost functional penalization (C = 0.01), the final value of J is less than any of the cost values of the manually specified trajectories, only with a slight increase in energy spent.

Then, it is important to analyze the paths resulting from the optimization algorithm without the control variable constraint. Firstly, as C increases the cost functional value increases as well, as expected. By observing Fig. 2, it is evident that with the increase of the turning angle penalization, the trajectories become smoother lines and the turning angles decrease. This phenomenon is consistent with the decrease in glider energy as C increases. The trajectory when C = 0.01 consumes less

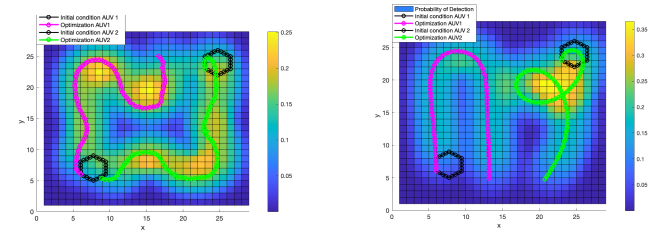

((a)) C = 0 ((b)) C = 0.001

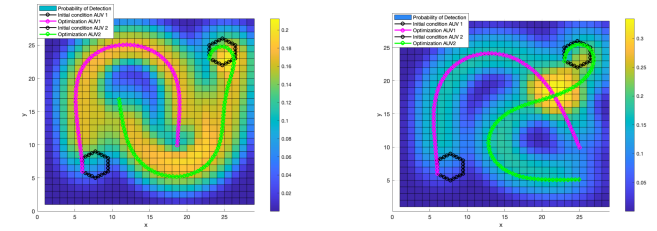

((c)) C = 0.0025 ((d)) C = 0.0105

Fig. 4: Optimized trajectories for 2 AUVs with a uniform prior probability density function.

94.69% energy than when C = 0.

Lastly, the trajectories with control variable constraint are analyzed. These paths also have the previously described expected trajectory characteristics and cost functional values. Comparing this situation with the one without control variable constraint, the first thing to note is that the cost functional values are similar for both the paths with and without the control variable constraint. By comparing Figs. 2(a) and 3(a), it can be concluded that the biggest difference in the paths lies when C = 0. In this case, it can be observed that when the constraint is used, the turning angles are smaller and trajectory generally follows (when possible due to the map limits) the same direction, which does not happen without the constraint. In addition, the energy consumption with and without constraint is very similar but slightly better with the constraint; and the energy savings are also marginally better, mainly the case when C = 0.005 saving (when compared to C = 0) 94.41% without constraint, and with constraint saving 94.99% of energy.

2) Two Gliders: As the RAMONES project operates with two underwater gliders, simulations were performed to reflect this situation. In this specific case, it is more relevant to study the trajectories of the gliders in regard to each other. This does not mean that the cost functional value or the gliders ability to sweep an area decreases in importance. However if two gliders explore a very wide area in a very limited time, the problem in hand becomes equivalent to having only one glider and a wide space, as studied the previous section. Therefore, the map size in the following simulations was decreased to $29m \times 29m$ and the initial trajectory was proportionally reduced to 84 steps. In this way, it is evident if the optimized path makes the gliders explore the map on a complimentary manner. In Fig. 4 the optimized trajectories can be seen.

The same simulation conditions as previously were tested, but now with the constraint that limits the turning angle to

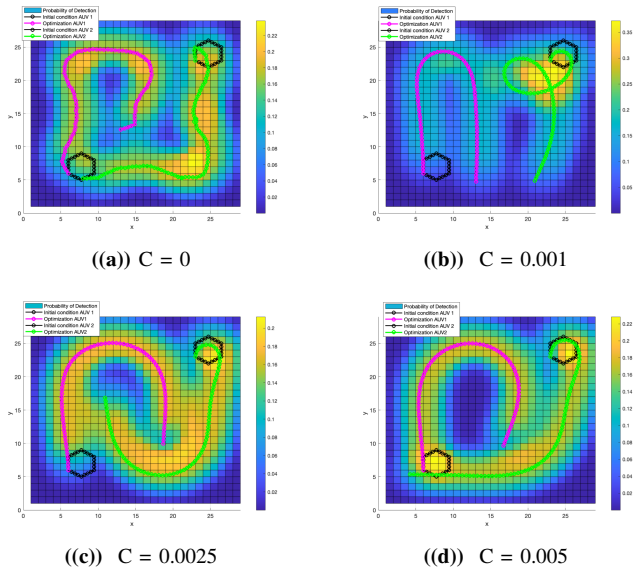

((a)) C = 0 ((b)) C = 0.001

((c)) C = 0.0025 ((d)) C = 0.005

Fig. 5: Optimized trajectories with control variable constraint for 2 AUVs with a uniform prior probability density function.

$\pm 30°$. The respective trajectories are in Fig. 5.

As previously, manually specified paths were tested. Of these trajectories the best performing one is the wide boxspiral. Both trajectories generated with the optimization algorithm (without the turning angle constraint and with the constraint) have lower cost functional value than the cost of the wide boxspiral. As a result, the optimally-generated trajectories have, in fact, better performance than the manually specified ones. However, in terms of energy consumption, the manually specified trajectories perform significantly better than the optimized ones: the most energy efficient manual trajectories are the lawnmowers, which have a significantly lower energy consumption than both optimized trajectories, with and without control input constraint.

By analyzing Fig. 4, the expected results can be observed. Firstly, the cost functional values increase with the increase of C. Then, the resulting trajectories observed clearly become smoother and less sinuous with the increase of C, which is an effect of the decrease of energy consumption with the increase of C. In addition, both glider trajectories do not make a redundant path with each other and explore the map in a complimentary way.

In the case with control variable constraint, the simulation characteristics are similar: with the increase of C the cost functional value increases and the trajectories have less curves (as can be observed in Fig. 5). Comparing with the previous situation, the cost functional values with the constraint are slightly lower. Additionally, the energy consumption is mostly lower and that the percentage of energy saved with high values of C, comparing to C = 0, is also higher.

3) Segment-based Optimization: Optimizing long trajectories in big maps is very computationally heavy. One solution to this issue is to divide the desired trajectory into small segments, and optimize them one by one, keeping in each optimization step the previous probability of detection by the sensor of the swept area. The trajectory obtained is only semi-optimized, due to the fact that the algorithm optimizes one segment at a time only with the knowledge of the path so far, and so cannot take into account the full length. The final trajectory can be seen in Fig. 6. The first trajectory is 72 steps long, and each iteration trajectory has a 72-step increment until the last trajectory, which is 648 steps long. After 9 algorithm iterations, the energy spent is 539.7112 energy units.

As the path length increases with each iteration, the initial cost functional value does not reflect what would be the cost of a path as long as the final one. Therefore, an initial condition path was computed with the full trajectory length, yielding a cost functional value of 1.0290. In the last iteration, the value of J was reduced to 0.9427.

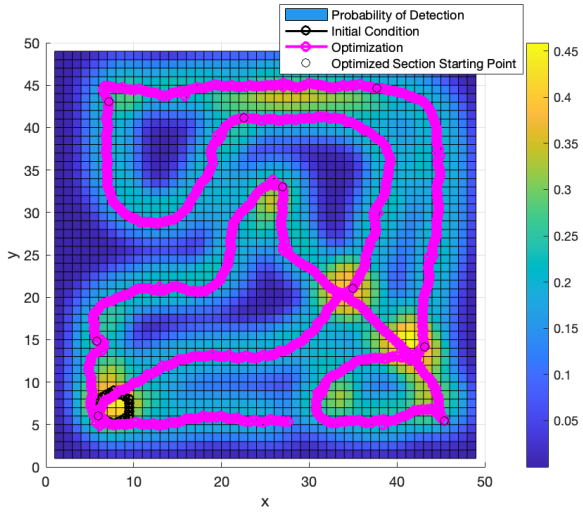

Fig. 6: Segment-based optimized path.

This method allows to get an semi-optimized path when there is lack of computational power or time, and obtains an almost full coverage approach in big maps. In Table I are the side-by-side comparisons between the segment-based optimization and the regular optimization with one glider and uniform prior (when C = 0). Comparing both cases, as expected the final cost functional of the segment optimization is lower than that of the regular optimization, due to the higher coverage obtained. Surprisingly, the energy spent per step is lower in this case than in the regular optimization. As the segments are proportionally much smaller than the side of the map, it is easier to generate straighter trajectory segments than in the classic optimization, which decreases amount the energy spent.

TABLE I: Comparison between the regular optimization done with one AUV and uniform prior (C = 0) and the optimization in segments.

	Normal Optimization	Segment-Based Optimization
Number of Steps	220	648
Initial J	0.9351	1.0290
Final J	0.9281	0.8869
EU	224.715	613.8990
EU/Step	1.1236	0.9474

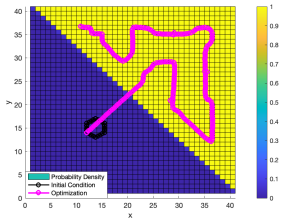

((a)) Prior probability density of target location and AUV trajectories for C = 0

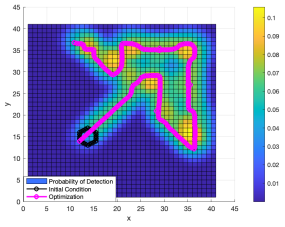

((b)) Sensor probability of target detection and AUV trajectories for C = 0.

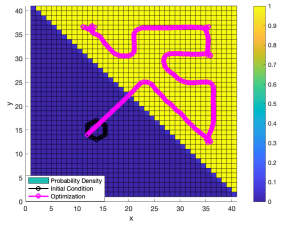

((c)) Prior probability density of target location and AUV trajectories for C = 0.01

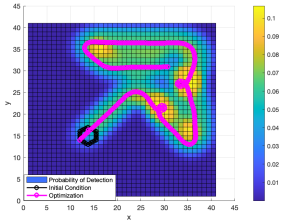

((d)) Sensor probability of target detection and AUV trajectories for C = 0. C = 0.01

Fig. 7: Optimized AUV trajectories for a 80/20 prior with 1 AUV.

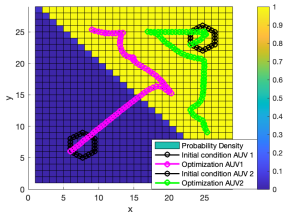

((a)) Prior probability density of target location and AUV trajectories for C = 0

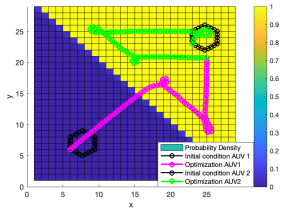

((b)) Sensor probability of target detection and AUV trajectories for C = 0.1

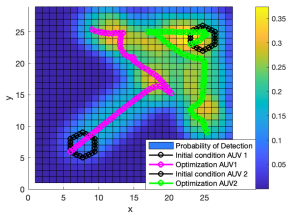

((c)) Prior probability density of target location and AUV trajectories for C = 0

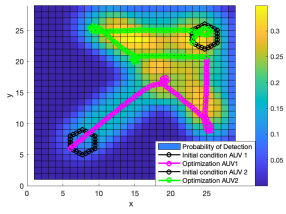

((d)) Sensor probability of target detection and AUV trajectories for C = 0.01

Fig. 8: Optimized AUV trajectories for a 80/20 prior with 2 AUVs.

G. Results: 80/20 Prior

In this section, the map is divided in half: the first part has a target location probability of 0.2 and the second part a probability of 0.8. This allows to confirm that the algorithm will indeed sweep only the highest probability side, which might be applied, for example, in a situation where half the map was already swept, and so it is of very little significance. The initial position was purposely chosen in the lower target location probability section to force the AUV to move to the highest probability area. As in this case the prior probability density values are significantly higher than with the uniform prior simulations, the cost functional values are considerably higher.

The turning angle constraint of $\pm 30^{\circ}$ was not possible to be satisfied in the following simulations. Due to the very limited space of the high probability area contemplated in these simulations, the glider is always forced to take sharper turns at one point or another, therefore this possibility in not reflected in this work.

1) One Glider: For the situation with a single glider, it is important that the vehicle sweeps as much of the high probability area as possible. In Fig. 7, the simulation resulting paths are shown. When there was a uniform prior, the probability of target location was equal in all area of the map, thus the prior probability of target location was not relevant, and it was not presented in the figures. However, this is not the case anymore. Thus, these simulations depict two figures for each turning angle penalization factor (C): the left hand side is the prior probability of target location, and the right hand side is the probability of sensor detection of the target.

These simulations depict expected results. Firstly, it is evident in Fig. 7 that the glider only goes through the lower probability of target location area to get to the higher probability area and only the latter is completely swept. Then, the cost functional values increase with C. Lastly, with the increase of the cost functional values, the trajectory becomes smoother and less sinuous. However, in contrast with the uniform prior simulations, it is not possible to eliminate the path loops completely with a long trajectory and the curves almost completely due to the very limited area of high probability. Despite this phenomenon, there is a decrease in energy consumption as C increases. The energy saved when comparing to C = 0 always increases, reaching values as high as 74.26%.

2) Two Gliders: As in the previous case, here the map also displays a section with low values of probability of target location, and another area where this value is high. In order to test the complimenting glider trajectories, the initial position of each glider was chosen in different sections: one glider starts in the 80% target location probability and the other AUV in the alternate section. This way, it is evident that the first glider stays and search the 80% section, while the second glider takes the shortest path to the same section and searches it. Similarly to the previous simulations, in Fig. 8, the AUVs trajectories can be seen along with the probability density function of target location and the probability density function of sensor detection.

As formerly, these simulations depict expected results. As noticeable in Fig. 8, the first glider only goes through the lower probability of target location area to get to the higher probability area and the second glider stays in the latter area. Both gliders sweep thoroughly the higher probability area. The cost functional values increase with C. As in the situation with one AUV, although with the increase of C the trajectories become smoother and less sinuous, here the limited map area available along with the long trajectories of both AUVs

make these paths have more trajectory loops and curves. This phenomenon explains the reason why EU does not always decrease with the increase of C.

III. Targeted Coverage Problem

This chapter delves into search with prior knowledge of the map. The radiation sources are modeled using Gaussian functions.

A. New Target Models

At this point, the first sweep of the map was already done, and there is knowledge of an approximate target location. The goal of the targeted coverage missions is to re-acquire data in the interest locations and to validate the obtained results.

For this purpose, it is necessary to model possible targets. The prior probability density function used for a single target on the map is the 2D Gaussian

$$\phi(\omega, \mu, \mathbf{\Sigma}) = \frac{1}{\sqrt{|\mathbf{\Sigma}|(2\pi)^2}} e^{-\frac{1}{2}(\omega-\mu)\mathbf{\Sigma}^{-1}(\omega-\mu)^T}, \tag{30}$$

where the μ is the mean vector and $\mathbf{\Sigma}$ is the covariance matrix. The mean vector was chosen so that the target would be at the center of the map (29 m $\times$ 29 m) and the covariance was chosen according to the size of the map. Therefore, the values used were

$$\mu = \begin{bmatrix} 14.5 \\ 14.5 \end{bmatrix}, \mathbf{\Sigma} = \begin{bmatrix} 12 & 10 \\ 10 & 12 \end{bmatrix}. \tag{31}$$

For the case with two targets on the same map, two symmetrical Gaussians were created using (30): $\phi_1(\omega, \mu_\mathbf{1}, \mathbf{\Sigma})$ and $\phi_1(\omega, \mu_\mathbf{2}, \mathbf{\Sigma})$, with

$$\mu_\mathbf{1} = \begin{bmatrix} 9.67 \\ 19.33 \end{bmatrix}, \mu_\mathbf{2} = \begin{bmatrix} 19.33 \\ 9.67 \end{bmatrix}. \tag{32}$$

The resulting probability density for two targets is

$$\phi(\omega) = \frac{1}{2}\phi_1(\omega) + \frac{1}{2}\phi_2(\omega). \tag{33}$$

B. Results: One Target

As in this chapter the prior information is the most relevant, the background of all simulations represents the prior probability density function of target location and not, like in previous simulations, the probability of the vehicle sensor to detect a target.

1) One Glider: The paths that resulted from the optimization algorithm for several angle variation penalization values can be seen in Fig. 9

Firstly, a boxspiral was tested for this case. The boxspiral was picked over the lawnmower due to its trajectory "following" the shape of the Gaussian better: the path created with the boxspiral passes through more area occupied by the Gaussian than the lawnmower, which has a trajectory that goes through a lot of "empty" space. Comparing the functional cost values of the boxspiral with the cost value of the optimized path when C = 0, it is once again obvious that the optimization

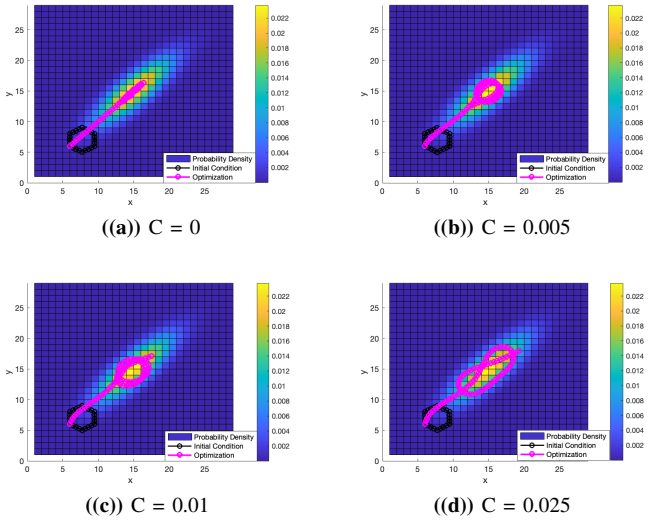

((a)) C = 0 ((b)) C = 0.005

((c)) C = 0.01 ((d)) C = 0.025

Fig. 9: Optimized trajectories for one target and one AUV.

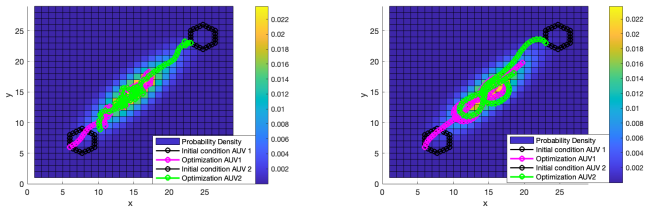

((a)) C = 0 ((b)) C = 0.01

Fig. 10: Optimized trajectories for one target and two AUVs.

algorithm generates the path with best performance. In terms of energy efficiency, the boxspiral shows a better result than all the optimized trajectories except when C = 0.025. However, this path not only has lower energy consumption, but it also has a lower cost functional value, even with the highest turning angle penalization.

As expected, the cost functional values increase with the increase of the turning angle penalization. In Fig. 9, it is evident that when C has a very low value, the trajectory is very redundant with very sharp turns. With higher angle variation penalization, the trajectory becomes wider. This latter path is preferable because not only the glider uses less energy but more area around the radiation source is explored (and not only the cells with higher intensity). As predicted, the trajectory energy consumption decreases with the increase of C. This decrease can be as large as 87.30% of the energy spent when C = 0.

2) Two Gliders: This case is particular due to the fact that there are more searchers than targets. Therefore, both searchers paths will converge to the same high target location probability area, as shown in Fig. 10.

As expected, the cost functional values increase with the increase of the turning angle penalization. In Fig. 10, it is clear that the gliders explore the same area. Although this is the expected behavior, considering the prior chosen, the vehicle paths are redundant and prone to crashes. As the

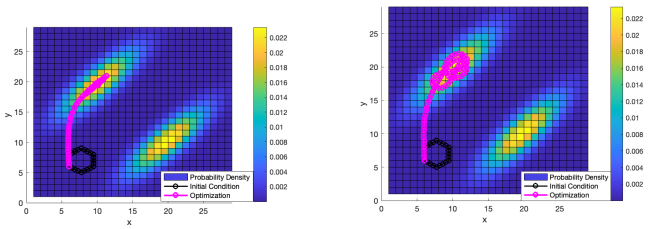

((a)) C = 0 ((b)) C = 0.005

((c)) C = 0.007 ((d)) C = 0.07

Fig. 11: Optimized trajectories for two targets and one AUV.

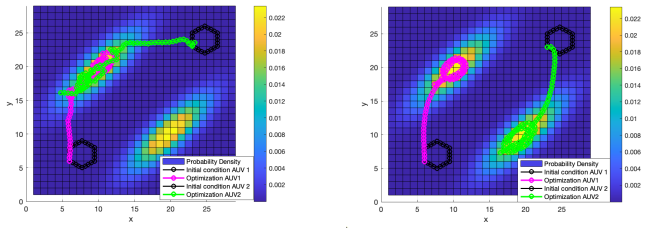

((a)) C = 0 ((b)) C = 0.0025

((c)) C = 0.007 ((d)) C = 0.02

Fig. 12: Optimized trajectories for two targets and two AUVs.

value of C increases, this phenomenon is less exacerbated, as the trajectories become wider. This also causes the energy consumption values to decrease as C increases, and from C = 0 to C = 0.01 the energy saved is as high as 75.72%.

C. Results: Two Targets

In the following sections two targets will be considered. Therefore, the prior probability density used is one with two peaks, as in (33).

1) One Glider: In this case, it is interesting to study not only the situation when, as expected, the glider trajectory converges to one of the targets, but also the situation in which, when the turning angle variation penalization increases, the trajectory goes through both targets. This phenomenon can be seen in Fig. 11.

These simulations have an interesting development with the increase of C, as can be seen in Fig. 11 : when C = 0 and C = 0.005 the glider path converges to a single radiation source; then, for C = 0.007 the trajectory goes through both targets, but makes sharper turns when going through the radiation center and has a very straight trajectory between the two targets; and lastly, for C = 0.07 the path goes through both targets in one single very wide curve. The energy consumption follows this phenomenon: for C = 0 it is very high due to the very narrow curves, then for C = 0.005 the energy value is already 77.20% lower due to the wider curves. As C increases and the trajectories are smother, the energy consumption decreases as much as 93.84% than when C = 0. As the goal is to go through both targets while using the least amount of energy, the paths with higher angle variation penalization values are preferable. Finally, as expected, the cost functional values increases with the increase of the turning angle penalization.

2) Two Gliders: As previously, several turning angle variation penalization values were studied. In this case, not only it is relevant the way the gliders interact with each other, but with the targets as well, which varies greatly with the values C. This phenomenon is evident in Fig. 12.

Two boxspirals (one for each target) were tested. The cost functional value of the manually specified trajectory used is significantly higher than its comparable optimized trajectory, when C = 0. Additionally, even taking into account the turning angle penalization, the boxspiral cost functional value is also higher than any other optimized trajectory tested. In terms of energy consumption, the boxspirals present a low consumption. This value is lower than any other optimized path value, with the exception of when C = 0.02.

The trajectories generated by the optimization algorithm differ widely with different values of C, as can be seen in Fig. 12. When C = 0, the least favorable result occurs: both gliders explore one single target, which not only leaves the other one unswept, but also the trajectories are redundant with high collision risk and have a very high energy consumption. When C = 0.0025, the most intuitive situation happens: each vehicle explores one source, here the energy spend decreases as much as 68.66%. When C = 0.007, an interesting development happens, as the gliders start by each one exploring a different target, but eventually one of them goes to the other's target. Lastly, when C = 0.02, both gliders explore both targets, without any collision risk and using minimal energy, up to 85% less than when C = 0. This latter path is the preferable one. As expected, the cost functional values increase with the increase of C.

IV. Conclusion

This study introduces trajectory planning solutions for a coverage search problem using a model-based trajectory optimization method. The successful development, implementation, and testing of an optimal control algorithm for generating optimized trajectories were presented. The results were as intuitively anticipated and aligned with the existing literature.

The primary objective was to explore completely unknown environments, testing both optimally-generated trajectories and manually-specified ones. As anticipated, the optimized trajectories outperformed the manually specified ones. Various angle variation penalization values, C, were tested. It was observed that a higher value of C results in a more efficient mission path: broader turns require less energy, while trajectories with very small C values are more redundant. The control variable limit of 30° is effective in cases with very small C and are slightly more energy efficient, but otherwise yields similar results. In scenarios involving two vehicles, the expected outcome was achieved: they covered the area in a complementary manner. Similar to previous cases, very low C values produced redundant and inefficient trajectories. Manually-specified trajectories with two AUVs were also tested but demonstrated poorer performance.

An alternative approach was explored: segment-based optimization. This method proves effective for optimizing extensive paths on large maps, ensuring substantial coverage with minimal computational resources. The resulting cost functional value was notably low, attributed to the extensive coverage achieved, and the energy consumption remained within regular levels. No strategies were explored to reduce the computation time of other simulated cases, as the simulations with a more complete map-coverage or who's path converges to an area with higher probability of radiation source location only take a few minutes to run, whereas the case with a long trajectory and a very wide map takes many hours (even reaching several days).

In the case of the 80/20 prior probability, the results were as anticipated: the area explored corresponded to the high probability region of target location, achieving maximum coverage in this section of the map. In this context, the influence of C was less significant due to the limited high probability area and the extended trajectory. The effect of increasing C was noticeable not by eliminating turns in the path, but by reducing their number while increasing their sharpness and overall reducing the path energy consumption. Similarly, in scenarios involving two AUVs, they explored the high probability area in a complementary manner, as expected. The variation of C did not always ensure the expected energy consumption values, due to the long trajectories in a very limited space.

Regarding simulations with a prior probability of target location, the outcomes were as predicted: trajectories converged to the point with the highest probability of target location, where the sensor detected the highest radiation intensity. For a single vehicle and one target, a manually specified trajectory was tested for comparison with optimized trajectories and, as before, it performed suboptimally. Comparing optimized trajectories, those with higher C values exhibited wider and more energy-efficient curves. As sensors had very high detection probability within their radius (similar in size to the entire source), sweeping a broader area around the center proved beneficial for better coverage. In the case of one target for two AUVs, the anticipated result was achieved: both trajectories converged to the target's center.

Lastly, in scenarios involving two targets and one vehicle, lower C values caused the path to converge to only one source, while larger values enabled the glider to cover both sources, as desired. With two targets and two gliders, manually specified trajectories were tested and were proven to be inferior to optimized trajectories. Moreover, as C increased, trajectories exhibited different behaviors: from both AUVs exploring one target, to each glider exploring a target, and finally both vehicles exploring both targets. The energy consumption decreased significantly with the increase of C.

In conclusion, this study yielded the expected results for the trajectory planning of two underwater gliders. Generally, higher angle variation penalization values proved beneficial for optimizing trajectories.

ACKNOWLEDGMENT

I would like to extend my sincere gratitude to the H2020 FET Proactive EU RAMONES [g.a. 101017808] for their support during the research and development of this thesis.

REFERENCES

[1] T. J. Mertzimekis, P. Nomikou, E. Petra, P. Batista, D. Cabecinhas, A. Pascoal, L. Sebastiño, J. Escartín, K. Kebkal, K. Karantzalos, A. Mallios, K. Nikolopoulos, and L. Maigne, "Radioactivity monitoring in ocean ecosystems (ramones)," in *Proceedings of the Conference on Information Technology for Social Good*, ser. GoodIT '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 216–220. [Online]. Available: https://doi.org/10.1145/3462203.3475906

[2] T. Mertzimekis, V. Lagaki, I. Madesis, G. Siltzovalis, E. Petra, P. Nomikou, P. Batista, D. Cabecinhas, A. Pascoal, L. Sebastiao, J. Escartin, K. Kebkal, K. Karantzalos, V. Douskos, A. Mallios, and K. Nikolopoulos, "Ramones and environmental intelligence: Progress update," in *Proceedings of the 2022 ACM Conference on Information Technology for Social Good*, ser. GoodIT '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 244–249. [Online]. Available: https://doi.org/10.1145/3524458.3547255

[3] E. Galceran and M. Carreras, "A survey on coverage path planning for robotics," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1258–1276, 2013. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S092188901300167X

[4] Z. L. Cao, Y. Huang, and E. L. Hall, "Region filling operations with random obstacle avoidance for mobile robots," *J. Field Robotics*, vol. 5, pp. 87–102, 1988. [Online]. Available: https://api.semanticscholar.org/CorpusID:5131874

[5] S. Kragelund, C. Walton, I. Kaminer, and V. Dobrokhodov, "Generalized optimal control for autonomous mine countermeasures missions," *IEEE Journal of Oceanic Engineering*, vol. 46, no. 2, pp. 466–496, 2021.

[6] J. Foraker, J. Royset, and I. Kaminer, "Search-trajectory optimization: Part i, formulation and theory," *Journal of Optimization Theory and Applications*, vol. 169, pp. 1–20, 06 2015.

[7] ——, "Search-trajectory optimization: Part ii, algorithms and computations," *Journal of Optimization Theory and Applications*, vol. 169, pp. 1–18, 06 2015.

[8] H. Xu, G. cheng Zhang, Y. shan Sun, and S. Pang, "Energy-saving control of long-range autonomous underwater vehicle vertical plane based on human simulating intelligent control method," *International Journal of Advanced Robotic Systems*, vol. 17, no. 5, p. 1729881420944744, 2020. [Online]. Available: https://doi.org/10.1177/1729881420944744