

# $A^{\text{SAR}}$ : $\varepsilon$ -Optimal Graph Search for Minimum Expected-Detection-Time Paths with Path Budget Constraints for Search and Rescue (SAR)

Eric Mugford<sup>1</sup> and Jonathan D. Gammell<sup>1</sup>

**Abstract**—Searches are conducted to find missing persons and/or objects given uncertain information, imperfect observers and large search areas in Search and Rescue (SAR). In many scenarios, such as Maritime SAR, expected survival times are short and optimal search could increase the likelihood of success. This optimization problem is complex for nontrivial problems given its probabilistic nature.

Stochastic optimization methods search large problems by nondeterministically sampling the space to reduce the effective size of the problem. This has been used in SAR planning to search otherwise intractably large problems but the stochastic nature provides no formal guarantees on the quality of solutions found in finite time.

This paper instead presents  $A^{\text{SAR}}$ , an  $\varepsilon$ -optimal search algorithm for SAR planning. It calculates a heuristic to bound the search space and uses graph-search methods to find solutions that are formally guaranteed to be within a user-specified factor,  $\varepsilon$ , of the optimal solution. It finds better solutions faster than existing optimization approaches in operational simulations. It is also demonstrated with a real-world field trial on Lake Ontario, Canada, where it was used to locate a drifting manikin in only 150s.

## I. INTRODUCTION

Search and rescue (SAR) operations are inherently resource intensive and may include land-, air-, and/or marine-based searches. Aerial searches are commonly conducted from manned aircraft, which are expensive to procure and operate [1], limiting the number available. This limited availability means that aerial assets are often late to arrive on scene, especially in remote locations.

Unmanned aerial vehicles (UAVs) are an alternative to manned aircraft for SAR operations. UAVs can be launched quickly and can search an area faster than land- or marine-based assets without the traditional expenses of manned aircraft. Crewed aircraft also generally employ simple search patterns, such as parallel track, that can be reliably flown by pilots operating large, relatively unmanoeuvrable platforms [2]. UAVs are more manoeuvrable and can instead perform complex search patterns optimized for individual scenarios.

A meaningful optimization objective for selecting search patterns (i.e., searcher paths) is expected detection time (i.e., mean time to detection; MTTD). The MTTD of a searcher

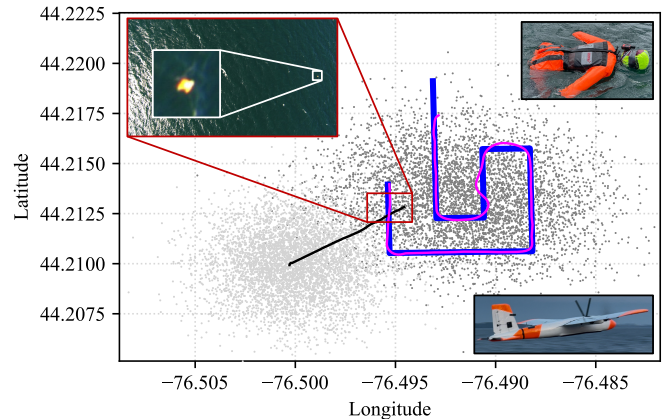


Fig. 1: Operational field experiment on Lake Ontario, Canada. A manikin (top right) drifted freely for one hour and was modelled with OpenDrift [7], with initial particle positions in light grey and the predicted position at search time in dark grey.  $A^{\text{SAR}}$  computed the minimum MTTD path, which was flown by a fixed-wing UAV (bottom right). Planned (blue) and flown (magenta) paths are shown up to their intersection with the manikin at 150s. The UAV image of the manikin is shown top left and its true trajectory (black) was calculated after the experiment from an attached GPS tracker.

path is the average time to first detect the target along the path for a given probability distribution of target locations. Minimizing MTTD may increase the probability of survival of a missing person since survival probability decreases with time, especially in cold or maritime environments.

Stochastic optimization algorithms can reduce MTTD given a starting position, searcher-path budget, prior target probability distribution and sensor model [3–5]. These algorithms use nondeterministic sampling to reduce the complexity of the path-constrained MTTD problem [6]. This allows them to search otherwise intractably large searcher-path problems but provides no formal guarantees on the quality of solutions found in finite time.

This paper instead presents  $A^{\text{SAR}}$ , an  $\varepsilon$ -optimal graph-search algorithm that can be applied to path-budget-constrained SAR planning.  $A^{\text{SAR}}$  bounds the search space by using an admissible heuristic which provides a lower bound on the MTTD objective function. It then uses a best-first search ordered by potential solution quality to find a solution that is within a user-defined suboptimality factor,  $\varepsilon$ , of the optimal path. This suboptimality factor allows the user to balance search speed and optimality guarantees.

$A^{\text{SAR}}$  is tested in realistic maritime search scenarios generated using OpenDrift [7], an open-source software package used by SAR professionals to model object drift.

<sup>1</sup>Estimation, Search, and Planning (ESP) Research Group, Queen’s University, Kingston ON, Canada. {19egm4, gammell}@queensu.ca

The authors would like to thank Ryan Berry, Grant Keefe, Jeremy Hayes, and Allan Taylor for their support and advice on this work as well as Prof. Sara Pérez-Carabaza for sharing her work. This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) [RGPIN-2024-06637] and Public Safety Canada’s Search and Rescue New Initiatives Fund (SAR NIF).




Fig. 2: The optimal searcher path (orange) calculated by  $A^{\text{SAR}}$  on a toy target distribution for a path length budget of 30, a perfect sensor and fixed starting position. The searcher maintains a target belief state that is updated and plotted with the searcher path.  $A^{\text{SAR}}$  plans a path to optimize the MTTD of a given object probability distribution. The belief state probability is reduced according to the sensor model when  $A^{\text{SAR}}$  searches a vertex.

We show that  $A^{\text{SAR}}$  finds better solutions faster than a leading MTTD path planning method [4] and traditional search patterns. The performance of  $A^{\text{SAR}}$  is also validated in a real-world experiment on Lake Ontario, Canada (Figure 1).

## II. RELATED WORK

Optimal searcher planning has been studied since the 1940s [8] and is provably NP-hard when maximizing the probability of detection and NP-complete when minimizing MTTD [6]. Stone [9] develops methods for the optimal distribution of resources between search areas to solve the MTTD problem. Eagle [10] extends these techniques to solve the searcher problem where the path is constrained to adjacent cells as a partially observable Markov decision process (POMDP). This method only works for small problems (e.g., three-by-three grids) and finds paths that maximize the probability of detection instead of minimizing MTTD. Subsequent research in probabilistic search planning has largely continued to focus on the development of methods to maximize detection probability [11–17]. A detailed review of probability of detection methods is available in [18].

Lau et al. [19] use dynamic programming to compute an optimal searcher policy that minimizes MTTD for multiple targets in an indoor environment. This method scales poorly, assumes a perfect sensor and stationary targets and does not consider path budget constraints.

Sarmiento et al. [20] consider the MTTD problem for a robot searching for a target through a building with a uniform probability distribution over the target location. This finds approximate solutions by dividing the search problem in two levels. The high level planner decides the order of room visitation using depth-first search and the low level planner finds the robot-trajectory using the Newton-Raphson method. This method is not globally optimal and does not extend to the general discrete-graph MTTD problem.

Hollinger et al. [21] introduce the Multi-robot Efficient Search Path Planning problem, which approximates the MTTD problem with a time-discounted reward (TDR) objective function. It enumerates all possible paths up to a fixed planning depth then applies sequential allocation that scales linearly with the number of searchers. This finds near-optimal solutions to the finite-horizon TDR problem but provides no guarantees to the MTTD problem.

The cooperative graph-based model predictive search algorithm [22] uses receding-horizon planning and dynamic probability-weighted graphs to plan the searcher paths of multiple UAVs. It achieves faster detection than greedy or parallel-track baselines with proven finite-time discovery guarantees but uses a probability-of-detection and travel-cost reward function that does not guarantee global MTTD optimality.

Lanillos [3] uses constraint programming (CP) to find solutions to the TDR approximation of the MTTD problem. This finds the optimal TDR solution but provides no guarantees to the MTTD problem. It is also only applicable to small problems and assumes a stationary target and a perfect sensor.

Lanillos [3] also introduces a truncated version of the MTTD function which can be optimized over a finite planning horizon. They use cross entropy optimization (CEO) to find searcher paths that minimize this objective. CEO uses nondeterministic sampling to reduce the complexity and can solve larger problems than the CP method without unrealistic sensor and target assumptions but provides no guarantees on the quality of solutions found in finite time.

Perez-Carabaza et al. [4] use ant colony optimization (ACO) to find solutions to the MTTD problem, which also uses nondeterministic sampling. They calculate a heuristic to guide the search towards higher areas of target probability. ACO outperforms other approximate methods, including CEO, Bayesian optimization [5] and genetic algorithms [23] but provides no guarantees on solution quality in finite time.

A related problem in robotics is informative path planning (IPP), which seeks paths that maximize information collected from the environment given a budget constraint. Typical objectives include minimizing entropy or maximizing mutual information, both of which are order-invariant, unlike minimizing MTTD. Some techniques from IPP can be applied to MTTD searcher path planning despite these differences.

Existing solution approaches for IPP include sampling-based planning [24], [25], nondeterministic sampling [26], and graph-search methods [27–29]. Schlotfeldt et al. [29] calculate an admissible and consistent heuristic to inform an A\*-style search. The heuristic lower bounds the remaining uncertainty cost by substituting the true sensor information matrix with an upper bound over the robot’s reachable set. This heuristic is only applied to order invariant objectives and time-invariant environments.

A\* [30] has been applied to searcher path planning. Guo et. al. [31] use A\* to guide a UAV from its launch base to the target search area while avoiding obstacles. A\* was not used to find the searcher path inside the search area which was instead calculated using a coverage path planning (CPP) algorithm. Neither A\* nor the CPP algorithm optimized for probability of detection or MTTD.

Alternatively, A<sup>SAR</sup> does not require unrealistic assumptions about the target or sensor and optimizes MTTD directly. Unlike approximate methods, A<sup>SAR</sup> bounds the search with a heuristic calculated from the searcher's reachable set, as in [29]. Unlike [29], A<sup>SAR</sup> extends the heuristic for an order- and time-variant problem. A<sup>SAR</sup> orders its A\*-style search for a searcher path with this heuristic to find solutions that are guaranteed to be within a user-specified factor,  $\varepsilon$ , of the optimal solution, unlike [31].

### III. A\* FOR SEARCH AND RESCUE (A<sup>SAR</sup>)

A<sup>SAR</sup> is a graph search algorithm based on A\* [30] that finds  $\varepsilon$ -optimal searcher paths to minimize the finite-path MTTD objective function given an initial position, searcher path budget, prior probability distribution over target location, target movement and sensor models.

A\* is a best-first search algorithm ordered by potential solution quality that finds minimum-cost paths in a graph. It extends Dijkstra's algorithm [32] to evaluate vertices in order of potential solution quality as given by a heuristic estimate of the cost of a solution path passing through the vertex. If this heuristic never over estimates the true solution cost, then it is said to be admissible and A\* is guaranteed to find the optimal solution as efficiently as any other equivalent algorithm, in terms of the number of vertex expansions [30].

Weighted A\* [33] introduces an inflation parameter,  $\varepsilon > 1$ , to the heuristic. This biases the search towards the goal and often expands fewer vertices than A\* while sacrificing optimality guarantees. Weighted A\* is guaranteed to find a suboptimal solution no worse than the inflation factor times the optimum.

A<sup>SAR</sup> applies A\* to minimize the MTTD of a path through a belief state that evolves both with time and the route taken (i.e., time and order *variant*). It does this while considering a maximum path budget by representing the problem as a graph where each vertex couples the searcher position and the belief state of the target location.

A<sup>SAR</sup> maintains a priority queue of vertices ordered by summing the MTTD objective function to reach a vertex (i.e., objective-to-come) and a heuristic estimate of the remaining objective available (i.e., objective-to-go). The vertex with minimum value is removed from the queue at each iteration and its outgoing edges are expanded (i.e., its child vertices are added to the queue). A<sup>SAR</sup> does not have a specific goal and the goal is instead defined as any vertex that has an incoming path with length equal to the budget.

A<sup>SAR</sup> also weights the heuristic by an inflation factor,  $\varepsilon \geq 1$ , to prioritize expanding vertices near a goal. This often expands fewer vertices while guaranteeing that the returned path is within the inflation factor of the optimum, as in Weighted A\*. This parameter allows the user to balance optimality and computational efficiency.

### A. Notation

Probabilistic searcher planning often represents the location of the target with a discrete probability map where each element stores the likelihood that it contains the target. The map is modelled as a finite, graph,  $G = (V, E)$ , where  $V$  is the set of vertices and  $E \subseteq V \times V$  is the set of edges, where a bidirectional edge exists between all spatially adjacent vertices (e.g., 4- or 8-connected).

The initial probability distribution over the target location is denoted as  $p$ . The search for the target occurs in discrete time and over a finite time horizon,  $t \in (1, \dots, T)$ , where  $T$  denotes the searcher path budget.

The searcher can only search for the target in one vertex at each time step. The searcher path,  $\sigma$ , is a sequence of edges,  $(v_i, v_j) \in E$ , such that  $\sigma = ((v_0, v_1), (v_1, v_2), \dots, (v_{n-1}, v_n))$ . A solution is a path such that the number of edges is equal to the searcher path budget, i.e.,  $n = T$ .

A<sup>SAR</sup> maintains a priority queue,  $Q$ , of states. Let  $A$  be any set and let  $B$  and  $C$  be subsets of  $A$ , i.e.,  $B, C \subseteq A$ . The notation  $B \stackrel{\pm}{\leftarrow} C$  is used for  $B \leftarrow B \cup C$  and  $B \leftarrow C$  for  $B \leftarrow B \setminus C$ .

### B. Problem Preliminaries

The belief state that represents the probability that the target is present and has not been detected at a vertex,  $v$ , at a time,  $t$ , given the searcher path,  $\sigma$ , is denoted as  $b(v, t | \sigma)$ . Note that the belief state is only affected by subset of the path that occurred before the query time. The belief state is initialized as the initial probability distribution over the target location,  $p$ , at the beginning of the search, i.e.,  $\forall v \in V, b(v, 1) = p(v)$ . The belief state is not a probability distribution as the total mass in the belief at any time is the probability that the target has yet to be detected [4].

Search targets often move during the search (e.g., ocean drift). The motion of the target from one vertex,  $v_i$ , to another,  $v_j$ , at a time,  $t$ , is modelled by  $P(v_j | v_i, t)$ . The belief state evolves through time via a motion model,

$$\bar{b}(v_j, t | \sigma) = \sum_{v_i \in V} P(v_j | v_i, t) b(v_i, t - 1 | \sigma), \quad (1)$$

where  $\bar{b}$  is the *a priori* belief state that accounts for the target motion to time  $t$  but only the search up to time  $t - 1$ .

The probability that the searcher detects the target at a vertex,  $v$ , if it is present (i.e., glimpse probability) is given by  $q(v)$ . Glimpse probability is assumed to be independent and identically distributed (i.i.d.) between timesteps and no false-positive detections are considered. When a vertex,  $v$ , is searched at time,  $t$ , its *a priori* belief is decreased by the glimpse probability,  $b(v, t | \sigma) = (1 - q(v)) \bar{b}(v, t | \sigma)$ .

We define a state,  $x \in X$ , as the tuple consisting of a vertex,  $v$ , the partial path taken to the vertex,  $\sigma_x$ , and the stored belief state,  $b_x[\cdot]$ , i.e.,  $x := \{v, \sigma_x, b_x[\cdot]\}$ . The cost of a state, denoted as  $t(x)$ , is the travel time required to reach that state along the path. The stored belief state,  $b_x[\cdot]$ , is represented as an array with each entry corresponding to the belief value,  $\forall u \in V, b_x[u] = b(u, t(x) | \sigma_x)$ .

### C. Problem Formulation

Two common objective functions in probabilistic searcher planning are maximizing the probability of detection,  $P_d$ , and minimizing the MTTD. The probability of detection of a path of length  $T$  is defined as

$$P_d(\sigma) = \sum_{v \in V} (1 - b(v, T | \sigma)), \quad (2)$$

where  $b(v, T | \sigma)$  is the probability that the target is at vertex  $v$  and has not been detected by the searcher path.

The expected detection time of a path is defined as,

$$\begin{aligned} E[D(\sigma)] &= \sum_{t=1}^{\infty} t P(D(\sigma) = t) \\ &= \sum_{t=1}^{\infty} t q(\sigma(t)) \bar{b}(\sigma(t), t | \sigma), \end{aligned} \quad (3)$$

where  $D(\sigma)$  is the detection time of the searcher taking the path,  $\sigma$ , and  $\sigma(t)$  is the vertex occupied by the path at time  $t$ . This objective function requires a sum over all time and when approximated to a finite time results in a suboptimal expected detection time [3]. When (3) is rewritten as

$$E[D(\sigma)] = \sum_{t=1}^{\infty} P(D(\sigma) > t), \quad (4)$$

it can be truncated to

$$J(\sigma) = \sum_{t=1}^T P(D(\sigma) > t) = \sum_{t=1}^T \sum_{v \in V} b(v, t | \sigma), \quad (5)$$

where  $J(\sigma)$  is the finite-path-budget MTTD objective function of a path,  $\sigma$ , and the sum,  $\sum_{v \in V} b(v, t | \sigma)$ , is the probability that the target has not been detected at time  $t$  by the searcher following the path [3]. A path that minimizes this truncated version is optimal for the path budget,  $T$ , even though it returns a lower estimate of the expected target detection time [3]. Formally the MTTD problem is then to find the optimal path,  $\sigma^*$ , such that,

$$\sigma^* = \arg \min_{\sigma} J(\sigma) = \arg \min_{\sigma} \sum_{t=1}^T \sum_{v \in V} b(v, t | \sigma). \quad (6)$$

The truncated MTTD objective function is non-Markovian since the portion of the objective from a vertex to the end of a path is dependent on the partial path up to that vertex.

### D. Path Budgeted Search (Algorithm 1)

$A^{\text{SAR}}$  maintains a priority queue of states,  $Q$ , ordered by the  $\varepsilon$ -weighted evaluation function,  $f(x) = g(x) + \varepsilon \hat{h}(x)$ , where  $\varepsilon$  is the suboptimality factor,  $g(x)$  is the portion of the objective function accrued to reach the state (i.e., the objective-to-come) and  $\hat{h}(x)$  is the admissible heuristic estimate of the objective function to be accrued from the state (i.e., the objective-to-go). The objective-to-come is calculated by summing the objective over the path to the vertex,

$$g(x) = \sum_{t=1}^{t(x)} \sum_{v \in V} b(v, t | \sigma_x). \quad (7)$$

### Algorithm 1: $A^{\text{SAR}}(x_{\text{start}}, T, \varepsilon)$

---

```

1  $Q \leftarrow \{x_{\text{start}}\}; g_{\text{best}} \leftarrow \infty; x_{\text{top}} \leftarrow x_{\text{start}};$ 
2 while  $t(x_{\text{top}}) < T$  do
3    $Q \leftarrow \{x_{\text{top}}\};$ 
4   for  $(v_{\text{top}}, v_{\text{new}}) \in E$  do
5      $t_{\text{new}} \leftarrow t(x_{\text{top}}) + t(v_{\text{top}}, v_{\text{new}});$ 
6     if  $t_{\text{new}} \leq T$  then
7        $\sigma_{\text{new}} \leftarrow \sigma_{x_{\text{top}}} \oplus (v_{\text{top}}, v_{\text{new}});$ 
8       for  $v \in V$  do
9          $\bar{b}_{\text{new}}[v] \leftarrow \sum_{u \in V} P(v | u, t_{\text{new}}) b_{x_{\text{top}}}[u];$ 
10         $b_{\text{new}}[v_{\text{new}}] \leftarrow (1 - q(v_{\text{new}})) \bar{b}_{\text{new}}[v_{\text{new}}];$ 
11        for  $v \in (V \setminus v_{\text{new}})$  do
12           $b_{\text{new}}[v] \leftarrow \bar{b}_{\text{new}}[v];$ 
13         $x_{\text{new}} \leftarrow \{v_{\text{new}}, \sigma_{\text{new}}, b_{\text{new}}[\cdot]\};$ 
14        if  $g(x_{\text{new}}) + \varepsilon \hat{h}(x_{\text{new}}) < g_{\text{best}}$  then
15           $Q \leftarrow \{x_{\text{new}}\};$ 
16          if  $g(x_{\text{new}}) < g_{\text{best}}$  and  $t_{\text{new}} \equiv T$  then
17             $x_{\text{best}} \leftarrow x_{\text{new}};$ 
18             $g_{\text{best}} \leftarrow g(x_{\text{new}});$ 
19         $x_{\text{top}} \leftarrow \underset{x \in Q}{\operatorname{argmin}} \{g(x) + \varepsilon \hat{h}(x)\};$ 
20 return  $x_{\text{best}}$ 

```

---

The admissible heuristic estimate of the objective-to-go is always less than the unknown optimal objective-to-go,

$$\hat{h}(x) \leq \sum_{t=t(x)+1}^T \sum_{v \in V} b(v, t | \sigma_x \oplus \sigma^*), \quad (8)$$

where  $\sigma^*$  is the best possible path achievable from the current state with the remaining budget and  $\oplus$  denotes the concatenation of two paths. The heuristic estimate of the objective-to-go is calculated using Algorithm 2.

The best solution objective function value,  $g_{\text{best}}$ , is initially set to infinity (Line 1). The state with minimum  $\varepsilon$ -weighted evaluation function value (i.e.,  $f$ -value),  $x_{\text{top}} := \arg \min_{x \in Q} f(x)$ , is removed from the queue at each iteration (Lines 3, 19). The successor states corresponding to the current state are then generated and evaluated (Lines 4–18).

The path cost of each successor (i.e., new) state,  $t_{\text{new}}$ , is calculated by adding the path cost of the current state,  $t(x_{\text{top}})$ , and the path cost of the edge between the current state and the successor state,  $t(v_{\text{top}}, v_{\text{new}})$  (Line 5). If the new state's path cost is less than or equal to the search budget,  $t_{\text{new}} \leq T$ , the stored *a priori* belief state,  $\bar{b}_{\text{new}}[\cdot]$ , is constructed using the motion model,  $P(v_i | v_j, t)$  (Lines 8–9). The stored belief state,  $b_{\text{new}}[\cdot]$ , is then constructed from the *a priori* belief state and by accounting for the effect of searching the new vertex,  $v_{\text{new}}$  (Lines 10–12).

If the  $f$ -value of the successor state is less than the best solution,  $f(x_{\text{new}}) = g(x_{\text{new}}) + \varepsilon \hat{h}(x_{\text{new}}) < g_{\text{best}}$ , then it could possibly improve the solution and the state is added to the queue,  $Q$  (Lines 14–15). The addition of this check reduces memory usage by only adding states that could possibly improve the solution to the queue and does not

affect the solution since any state that fails the check would never be removed from the queue before the terminal state. The best solution,  $g_{\text{best}}$ , is only updated when a goal state is added to the queue with a lower objective-to-come than the current best solution (Lines 16–18). A goal vertex is any vertex that uses the entire path budget, i.e.,  $t_{\text{new}} = T$ .

$A^{\text{SAR}}$  ends when a goal vertex is removed from the queue, as in  $A^*$ . The first goal vertex removed from the queue will have an  $f$ -value equal to the best solution objective function value, i.e.,  $f(x_{\text{goal}}) = g_{\text{best}} + 0$ , and every vertex left in the queue must have a higher  $f$ -value,  $\forall x \in Q, g_{\text{best}} \leq g(x) + \varepsilon \hat{h}(x)$ . The admissible heuristic never overestimates the objective-to-go, i.e.,  $\forall x, g(x) + \hat{h}(x) \leq g(x) + h(x)$ , which ensures that the best solution's objective function value is guaranteed to be within the suboptimality factor,  $\varepsilon$ , of the optimal solution, as in Weighted  $A^*$ .

### E. Heuristic Calculation (Algorithm 2)

$A^{\text{SAR}}$  calculates a heuristic estimate of the objective-to-go in a manner similar to [29] to order the search by solving a relaxed problem where at each future step,  $t > t(x)$ , the searcher may search any vertex reachable within  $t - t(x)$  steps from the current location without being constrained by path continuity. The resulting sequence of vertices is a lower bound on the path-constrained MTTD objective function (Theorem 1). It is calculated by greedily selecting the vertex that maximizes the probability of detecting the target,  $q(v)b[v]$ , at each time step.

The heuristic value,  $\hat{h}$ , is initially set to zero (Line 1). The stored heuristic belief state,  $b[\cdot]$ , is initialized as the stored belief state of the inputted state,  $b_x[\cdot]$  (Lines 2–3). The heuristic is calculated by iterating through each time step in the remaining budget i.e.,  $t \in (t(x) + 1, t(x) + 2, \dots, T)$  (Line 4). At each time step, the *a priori* belief state,  $\bar{b}[\cdot]$ , is constructed using the motion model,  $P(v_i | v_j, t)$  (Lines 5–6). The vertex with the highest probability of detection that is reachable from the current vertex is then found (Line 7). The stored heuristic belief state,  $b[\cdot]$ , is finally updated to include the effect of searching the selected vertex (Lines 8–10). The heuristic value is updated by calculating the probability that the target has not been detected by time  $t$  and appending it to the running value (Line 11). The final heuristic value is returned after iterating through all time steps up to the path budget,  $T$  (Line 13).

**Theorem 1** (Admissibility). *The estimate of the objective-to-go,  $\hat{h}(\cdot)$ , never overestimates the optimal objective-to-go,  $h^*(\cdot)$ , and is admissible, i.e.,*

$$\forall x \in X, \hat{h}(x) \leq h^*(x).$$

*Proof sketch.* Define a relaxation in which at each future step,  $t > t(x)$ , the searcher may search any vertex reachable within  $t - t(x)$  steps from the current location without considering path continuity. The optimal value of this relaxed problem underestimates the optimal objective to go,

$$h_{\text{relax}}^*(x) \leq h^*(x), \quad (9)$$

---

### Algorithm 2: Heuristic( $x$ )

---

```

1  $\hat{h} \leftarrow 0$ ;
2 for  $v \in V$  do
3    $b[v] \leftarrow b_x[v]$ ;
4 for  $t \in (t(x) + 1, t(x) + 2, \dots, T)$  do
5   for  $v \in V$  do
6      $\bar{b}[v] \leftarrow \sum_{u \in V} P(v | u, t)b[u]$ 
7      $v_{\text{best}} \leftarrow \operatorname{argmax}_{v \in V} \{q(v)\bar{b}[v] \mid t(v_x, v) \leq t\}$ ;
8      $b[v_{\text{best}}] \leftarrow (1 - q(v_{\text{best}}))\bar{b}[v_{\text{best}}]$ ;
9     for  $v \in (V \setminus v_{\text{best}})$  do
10       $b[v] \leftarrow b[v]$ 
11       $\hat{h} \leftarrow \hat{h} + \sum_{v \in V} b[v]$ 
12 return  $\hat{h}$ 

```

---

because all possible solutions to  $h(x)$  lie within the set of possible solutions to  $h_{\text{relax}}(x)$ .

The probability mass of the target location is preserved through time and only lowered when the searcher searches a vertex,

$$\sum_{v \in V} \bar{b}(v, t + 1) = \sum_{v \in V} b(v, t), \quad (10)$$

$$\sum_{v \in V} b(v, t) = \sum_{v \in V} \bar{b}(v, t) - q(u)\bar{b}(u, t). \quad (11)$$

If there exists some larger reduction in probability,  $q(v)b(v, \tau) > q(v)b(v, t)$ , at some later time,  $\tau > t$ , then exchanging the searches so that larger reduction in probability occurs first reduces (or leaves unchanged) all future probability and strictly reduces the objective. The optimal solution to the relaxed problem is therefore greedy; at each step,  $t$ , choose the vertex maximizing the reduction in probability mass,  $q(v)\bar{b}(v, t)$ . Algorithm 2 implements this greedy solution, and the heuristic value is therefore the optimal value of the relaxed problem,

$$\hat{h}(x) = h_{\text{relax}}^*(x). \quad (12)$$

Since the optimal solution to the relaxed problem lower bounds the optimal objective-to-go,  $h_{\text{relax}}^*(x) \leq h^*(x)$ , the heuristic is admissible,  $\hat{h}(x) \leq h^*(x)$ .  $\square$

## IV. EXPERIMENTS

The  $A^{\text{SAR}}$  algorithm was developed as part of a project focusing on maritime SAR and experiments were conducted in the maritime domain.

### A. Drift Model

SAR planning professionals frequently employ Monte Carlo particle-drift models to predict object movement [7], [34], which are then used to plan searches for missing vessels or objects at sea. We generated realistic test scenarios with OpenDrift [7], an open-source framework for modelling oceanic drift. We used its Leeway model, derived from the

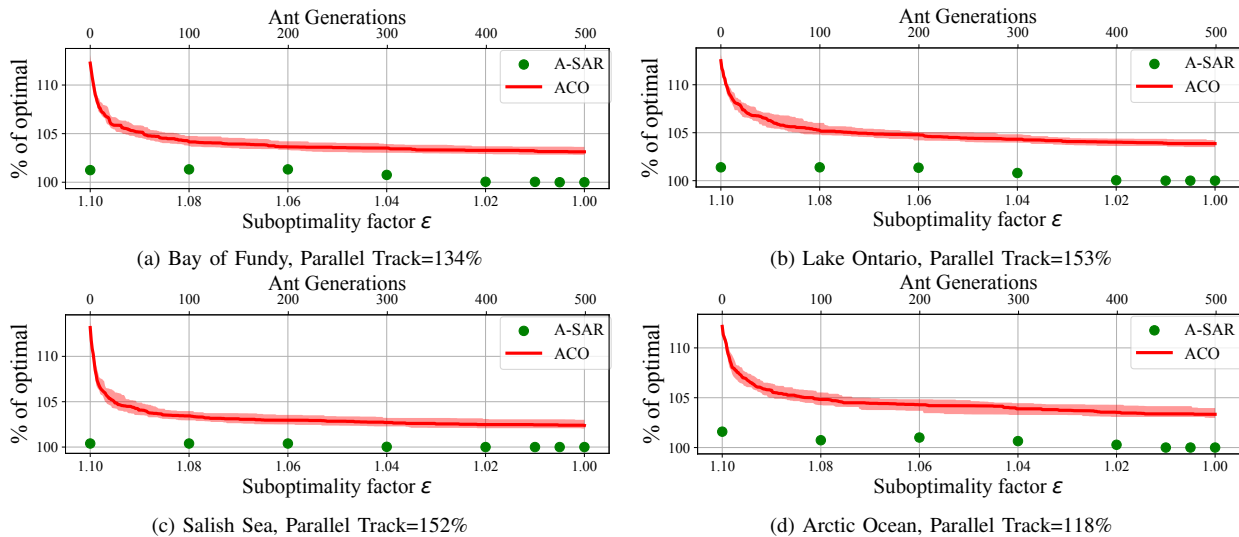


Fig. 3: Experimental results for Bay of Fundy (3a), Lake Ontario (3b), Salish Sea (3c) and Arctic Ocean (3d) scenarios comparing an approach from [4] and  $A^{SAR}$ . The results are expressed as a percentage of the optimal solution. The median values of the ACO planner after 100 trials are plotted along the ant generations axis with nonparametric 99% confidence intervals. The results of  $A^{SAR}$  are plotted along the suboptimality factor,  $\epsilon$ , axis. The parallel track (PT) baseline results for each scenario are included in the figure labels.  $A^{SAR}$  outperforms ACO and parallel track in all tested scenarios for all suboptimality factors while providing optimality guarantees.

US Coast Guard’s SAR model [35–37], which uses empirical coefficients to describe how objects react to wind and current. By combining these coefficients with environmental data, OpenDrift produces trajectory ensembles that define the initial belief state of object location,  $p$ . In our implementation, particles define the belief state. Each particle carries an associated probability that it has not been detected (PND). A particle’s PND is reduced by the glimpse probability when the searcher path passes over its cell. The belief of a cell is obtained by summing the PNDs of all particles it contains.

### B. Simulated Experiments

We used OpenDrift to create four distress scenarios which are summarized in Table I. These scenarios occur in the Bay of Fundy (Atlantic Ocean), Lake Ontario (North American Great Lakes), Salish Sea (Pacific Ocean) and Hudson’s Bay (Arctic Ocean). Four different OpenDrift search objects were considered: a person in water (PIW) in an unknown state (PIW-1), a conscious PIW positioned vertically and wearing a PFD (PIW-2), a fishing vessel and a deep-ballast life raft with unknown capacity and loading. These parameters were selected to ensure that  $A^{SAR}$  was assessed across a broad spectrum of realistic maritime SAR conditions.

The position and time of the distress (i.e., the last known position; LKP) is specified in all scenarios. Table I contains the LKP, distress time, commence search point (CSP), commence search time (CST), wind and current models used in the drift model, sweep width and the path budget. The difference between distress time and CST represent the drift time of the object and vary for the four different scenarios. The sweep width is the spacing between search legs in a parallel track search and is determined based on the sensor model, search object and weather conditions [2].

The probability map for each scenario was created by

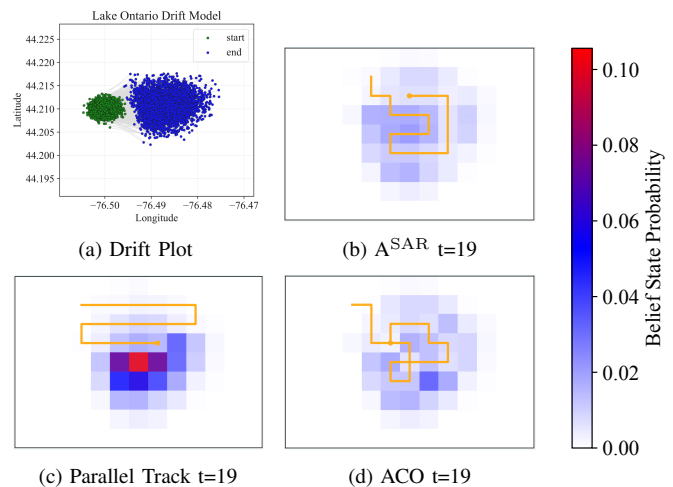


Fig. 4: The OpenDrift model and the paths found by  $A^{SAR}$ , an ACO approach [4] and parallel track on the Lake Ontario scenario. Fig 4a shows the OpenDrift Monte Carlo particle drift used to model object location. Figs. 4b, 4c, and 4d show the  $A^{SAR}$ , parallel track, and ACO searcher paths after 19 steps, respectively.  $A^{SAR}$  outperforms parallel track and ACO by prioritizing areas with high probability first and conducting the search efficiently, avoiding unnecessary revisits to recently explored vertices when moving toward other high-probability regions.

overlaying a grid on the drift plot and counting the number of particles in each cell for each time step. The cell size was chosen as the appropriate sweep width for the search object and visibility, given that the search was conducted from fixed wing aircraft at 300 ft, travelling at 20 m/s [2]. These searcher parameters were chosen as they best match the UAV used in the field trial in Section IV-C. We use a sensor model with the glimpse probability being 78% for all vertices, i.e.,  $\forall v \in V, q(v) = 0.78$ , in our experiments [2], although any sensor model could be used with  $A^{SAR}$ .

We compared the performance of  $A^{SAR}$  to the max-min

TABLE I: Testing Scenarios

Scenario	Search Object	LKP	Distress Time	CSP	CST	Wind Model	Current Model	Sweep Width	$T$
Bay of Fundy	PIW-1	N 45.2 W 65.3	2025-02-14 00:00 UTC	N 45.283 W 65.072	2025-02-14 09:00 UTC	HRRR [38]	GoMOFS [39]	0.4Nm	55
Lake Ontario	PIW-2	N 44.21 W 76.5	2025-08-26 14:45 UTC	N 44.215 W 76.496	2025-08-26 15:45 UTC	HRRR [38]	LOOFS [40]	0.1Nm	49
Salish Sea	Fishing Vessel	N 48.3 W 123.07	2025-08-24 00:00 UTC	N 48.195 W 123.187	2025-08-24 11:00 UTC	HRRR [38]	CIOPS Salish Sea [41]	1.6Nm	41
Arctic Ocean	Deep Ballast Life Raft	N 62.85 W 90.6	2025-08-29 00:00 UTC	N 62.786 W 90.447	2025-08-30 00:00 UTC	GFS [42]	ESPC-D-V02 [43]	0.5Nm	60

TABLE II:  $A^{\text{SAR}}$  Execution Time (seconds)

Scenario	Suboptimality factor $\varepsilon$							
	1.1	1.08	1.06	1.04	1.02	1.01	1.005	1.0
Bay of Fundy	0.10	0.11	0.12	0.17	0.29	1.97	5.45	19.73
Lake Ontario	0.09	0.09	0.10	0.13	0.19	0.30	0.52	0.86
Salish Sea	0.05	0.05	0.07	0.22	0.75	1.32	1.70	2.24
Arctic Ocean	0.10	0.10	0.10	0.50	0.52	3.09	10.78	35.26

ACO variant in [4]. We implemented the node encoding with the best parameters from [4].

We restrict motion to the four cardinal neighbours, excluding diagonals (i.e., 4 connected). This reflects the fixed-wing UAV application where the cell width is the sweep width of a downward-facing camera and diagonal motion would leave parts of the cell unobserved.

The results of the experiments are plotted in Figure 3. A baseline was set using a parallel track search created by a certified maritime search planner for each scenario. We then ran  $A^{\text{SAR}}$  and the ACO method from the same starting position and with the same budget as the length of the parallel track search. A plot of the Lake Ontario scenario’s modelled drift and resulting searcher paths is shown in Figure 4.

The performance of the ACO method was measured according to the number of ant generations, a number that increases with time, instead of computational time because our implementation of the ACO was not performance optimized. We plot the median solution quality with nonparametric 99% confidence interval of 100 trials in each experiment.  $A^{\text{SAR}}$  is deterministic and finds one result for any given problem and suboptimality factor,  $\varepsilon$ , with solution quality increasing as the suboptimality factor decreases. The solution quality is plotted as a function of the suboptimality factor alongside the ACO results.  $A^{\text{SAR}}$ ’s wall clock time is presented in Table II for each scenario and suboptimality factor.  $A^{\text{SAR}}$  was implemented in C++ and run on a 3.2 GHz M1 Pro processor.

### C. Lake Ontario Field Experiment

A field experiment was conducted on Lake Ontario to validate the performance of  $A^{\text{SAR}}$  on a real-world search object. The experiment involved dropping a manikin in position N44.210, W76.500 on 2025-09-08 at 19:15 UTC. The manikin was then allowed to drift freely for one hour. A drift model was created for the manikin using OpenDrift and the same wind and current models as the Lake Ontario scenario

in Table I. This drift model was used to create the probability map and motion model in the same manner as the previous experiments.  $A^{\text{SAR}}$  was then run to create a searcher path which was executed by a fixed wing UAV (Figure 1). This UAV is equipped with downward facing RGB and infrared cameras with a sweep width of 0.1 Nm. Figure 1 shows the drift model, drift track, searcher path and point of intersection of the manikin and searcher path of the experiment.

## V. DISCUSSION

Figure 3 shows that  $A^{\text{SAR}}$  finds a better solution in terms of MTTD than the ACO planner in [4] and the parallel-track patterns. The  $\varepsilon=1.0$  configuration achieved truncated MTTD objective values that were 3.18% lower on average than the median solutions produced by ACO and 39% lower on average than the parallel tracks searches. Even the weighted setting,  $\varepsilon=1.1$ , produced a solution less than 101.5% of the optimal truncated MTTD objective on average and outperformed the median ACO solution in all experiments.

$A^{\text{SAR}}$  finds these better solutions faster than ACO and provides formal solution quality guarantees. Our implementation of ACO was not performance optimized, so exact ACO run times are not included, but the experimental section of [4] runs ACO for 50s before convergence on problems with significantly shorter path budgets. The number of ants per generation is proportional to path budget and the time required for ant to construct a path increases with path budget. This implies that significantly more computational time would be required for convergence on real scenarios.  $A^{\text{SAR}}$  can find the optimal path in less time and found paths within 1.15% of optimum on average within 0.1 seconds. This shows that  $A^{\text{SAR}}$  can be used in real-world scenarios.

The Lake Ontario field trial confirms that this performance transfers to a real-world search target. The fixed-wing UAV located the drifting manikin in 150 seconds, within 20% of the detection time forecast by the planner despite localization uncertainty.

## VI. CONCLUSION

This work introduced  $A^{\text{SAR}}$ , an A\*-style graph-search algorithm that computes  $\varepsilon$ -optimal MTTD searcher trajectories under path budget constraints in dynamic environments.  $A^{\text{SAR}}$  bounds the search space with an admissible heuristic that provides a lower bound on the MTTD objective function. It then uses a best-first search ordered by potential solution quality to find a solution that is within a user-defined suboptimality factor,  $\varepsilon$ , of the optimal path.

Simulated experiments using OpenDrift models showed that A<sup>SAR</sup> finds better solutions faster than a leading MTTD searcher planning method and significantly outperforms traditional search patterns. A field trial validated A<sup>SAR</sup>'s performance on a real-world search object. Future work will include anytime planning/replanning, multiagent coordination and the inclusion of searcher kinematic constraints.

## REFERENCES

- [1] "Report 7: Operating and Maintenance Support for Military Equipment—National Defence, of the Fall 2016 Reports of the Auditor General of Canada," House of Commons, Parliament of Canada, Report of the Standing Committee on Public Accounts 7, Jun. 2017.
- [2] "CAMSAR: Canadian Aeronautical and Maritime Search and Rescue Manual, Combined Edition – Volumes I, II and III," Department of National Defence and Canadian Coast Guard, Technical Manual B-GA-209-001/FP-001 DFO 5449, Sep. 2014.
- [3] P. Lanillos, "Minimum time search of moving targets in uncertain environments," Ph.D. dissertation, 2013.
- [4] S. Pérez-Carabaza, E. Besada-Portas, J. A. Lopez-Orozco, and J. M. de la Cruz, "Ant colony optimization for multi-UAV minimum time search in uncertain domains," *Applied Soft Computing*, vol. 62, p. 789–806, Jan 2018.
- [5] P. Lanillos, J. Yañez-Zuluaga, J. J. Ruz, and E. Besada-Portas, "A Bayesian approach for constrained multi-agent minimum time search in uncertain dynamic domains," *Proceedings of the 15th annual conference on genetic and evolutionary computation*, p. 391–398, Jul 2013.
- [6] K. E. Trummel and J. R. Weisinger, "Technical note—the complexity of the optimal searcher path problem," *Operations Research*, vol. 34, no. 2, p. 324–327, Apr 1986.
- [7] K.-F. Dagestad, J. Röhrs, Ø. Breivik, and B. Ådlandsvik, "OpenDrift v1.0: A generic framework for trajectory modelling," *Geoscientific Model Development*, vol. 11, no. 4, p. 1405–1420, Apr 2018.
- [8] B. O. Koopman, "A theoretical basis for method of search and screening," Jan 1946.
- [9] L. D. Stone, *Theory of optimal search*. Academic Press, Inc, 1975.
- [10] J. N. Eagle, "The optimal search for a moving target when the search path is constrained," *Operations Research*, vol. 32, no. 5, p. 1107–1115, Oct 1984.
- [11] T. Stewart, "Search for a moving target when searcher motion is restricted," *Computers & Operations Research*, vol. 6, no. 3, p. 129–140, Jan 1979.
- [12] J. N. Eagle and J. R. Yee, "An optimal branch-and-bound procedure for the constrained path, moving target search problem," *Operations Research*, vol. 38, no. 1, p. 110–114, Feb 1990.
- [13] R. F. Dell, J. N. Eagle, G. H. Alves Martins, and A. G. Santos, "Using multiple searchers in constrained-path, moving-target search problems," *Naval Research Logistics*, vol. 43, no. 4, p. 463–480, Jun 1996.
- [14] F. Bourgault, T. Furukawa, and H. F. Durrant-Whyte, "Optimal search for a lost target in a Bayesian world," *Springer Tracts in Advanced Robotics*, p. 209–222, 2006.
- [15] H. Lau, "Optimal search in structured environments," Ph.D. dissertation, 2007.
- [16] J. Berger and N. Lo, "An innovative multi-agent search-and-rescue path planning approach," *Computers & Operations Research*, vol. 53, p. 24–31, Jan 2015.
- [17] P. N. Kasthurirangan, L. Nguyen, M. Perk, N. Chakraborty, and J. S. Mitchell, "Provable methods for searching with an imperfect sensor," *ICRA*, May 2025.
- [18] M. Raap, M. Preuß, and S. Meyer-Nieberg, "Moving target search optimization – a literature review," *Computers & Operations Research*, vol. 105, p. 132–140, May 2019.
- [19] H. Lau, S. Huang, and G. Dissanayake, "Optimal search for multiple targets in a built environment," *IROS*, p. 3740–3745, 2005.
- [20] A. Sarmiento, R. Murrieta-Cid, and S. Hutchinson, "An efficient motion strategy to compute expected-time locally optimal continuous search paths in known environments," *Advanced Robotics*, vol. 23, no. 12–13, p. 1533–1560, Jan 2009.
- [21] G. Hollinger, S. Singh, J. Djughash, and A. Kehagias, "Efficient multi-robot search for a moving target," *IJR*, vol. 28, no. 2, p. 201–219, Feb 2009.
- [22] J. R. Riehl, G. E. Collins, and J. P. Hespanha, "Cooperative search by UAV teams: A model predictive approach using dynamic graphs," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 47, no. 4, p. 2637–2656, 2011.
- [23] L. Lin and M. A. Goodrich, "UAV intelligent path planning for wilderness search and rescue," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, p. 709–714, Oct 2009.
- [24] G. A. Hollinger and G. S. Sukhatme, "Sampling-based robotic information gathering algorithms," *IJR*, vol. 33, pp. 1271–1287, August 2014.
- [25] B. Moon, S. Chatterjee, and S. Scherer, "TIGRIS: An informed sampling-based algorithm for informative path planning," in *IROS*, 2022, pp. 5760–5766.
- [26] G. Hitz, E. Galceran, M.-E. Garneau, F. Pomerleau, and R. Siegwart, "Adaptive continuous-space informative path planning for online environmental monitoring," *Journal of Field Robotics*, vol. 34, no. 8, pp. 1427–1449, 2017.
- [27] C. Chekuri and M. Pal, "A recursive greedy algorithm for walks in directed graphs," in *FOCS*. Pittsburgh, USA: IEEE, 2005, pp. 245–253.
- [28] J. Binney and G. S. Sukhatme, "Branch and bound for informative path planning," in *ICRA*. St. Paul, USA: IEEE, 2012, pp. 2147–2154.
- [29] B. Schlotfeldt, N. Atanasov, and G. J. Pappas, "Maximum information bounds for planning active sensing trajectories," *IROS*, p. 4913–4920, Nov 2019.
- [30] P. Hart, N. 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, p. 100–107, 1968.
- [31] Y. Guo, J. Yi, and J. Hao, "UAV maritime search and rescue based on improved A\* and RCPP algorithms for path planning," *IEEE 5th International Conference on Power, Intelligent Computing and Systems (ICPICS)*, p. 873–879, Jul 2023.
- [32] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, no. 1, p. 269–271, Dec 1959.
- [33] I. Pohl, "Heuristic search viewed as path finding in a graph," *Artificial Intelligence*, vol. 1, no. 3–4, p. 193–204, Jan 1970.
- [34] T. M. Kratzke, L. D. Stone, and J. R. Frost, "Search and rescue optimal planning system," *13th International Conference on Information Fusion*, p. 1–8, Jul 2010.
- [35] A. A. Allen, "Leeway Divergence," U.S. Coast Guard Research and Development Center, Technical Report CG-D-05-05, 2005.
- [36] A. A. Allen and J. V. Plourde, "Review of Leeway: Field Experiments and Implementation," U.S. Coast Guard Research and Development Center, Technical Report CG-D-08-99, 1999.
- [37] O. Breivik, A. A. Allen, C. Maisondieu, J.-C. Roth, and B. Forest, "The leeway of shipping containers at different immersion levels," *Ocean Dynamics*, vol. 62, pp. 741–752, 2012.
- [38] D. C. Dowell *et al.*, "The high-resolution rapid refresh (HRRR): An hourly updating convection-allowing forecast model. part I: Motivation and system description," *Weather and Forecasting*, vol. 37, no. 8, pp. 1371 – 1395, 2022.
- [39] Z. Yang, P. Richardson, Y. Chen, J. Kelley, E. Myers, F. Aikman, M. Peng, and A. Zhang, "Model development and hindcast simulations of NOAA's Gulf of Maine Operational Forecast System," *Journal of Marine Science and Engineering*, vol. 4, no. 4, p. 77, Nov 2016.
- [40] P. Y. Chu, J. G. Kelley, G. V. Mott, A. Zhang, and G. A. Lang, "Development, implementation, and skill assessment of the NOAA/NOS Great Lakes Operational Forecast System," *Ocean Dynamics*, vol. 61, no. 9, p. 1305–1316, May 2011.
- [41] S. MacDermid, F. Dupont, J.-P. Paquin, and Y. L. Shin, "Coastal ice-ocean prediction system-west (CIOPS-West) version 2.3.0," Environment and Climate Change Canada, Tech. Rep., jun 2024.
- [42] National Centers for Environmental Prediction, "The GFS atmospheric model," National Centers for Environmental Prediction, Office Note 442, 2004, accessed: Feb. 18, 2026.
- [43] N. Barton *et al.*, "The Navy's earth system prediction capability: A new global coupled atmosphere-ocean-sea ice prediction system designed for daily to subseasonal forecasting," *Earth and Space Science*, vol. 8, no. 4, p. e2020EA001199, 2021.