

Navigating Underactuated Agents by Hitchhiking Forecast Flows

Marius Wiggert¹, Manan Doshi², Pierre F.J. Lermusiaux² and Claire J. Tomlin¹

Abstract—In dynamic flow fields such as winds and ocean currents an agent can navigate by *going with the flow*, only using minimal propulsion to nudge itself into beneficial flows. This navigation paradigm of hitchhiking flows is highly energy-efficient. However, reliable navigation in this setting remains challenging as typically only forecasts are available which differ significantly from the true currents and the forecast error can be larger than can be handled by the actuation of the agent. In this paper, we propose a novel control method for reliable navigation of underactuated agents hitchhiking flows based on imperfect forecasts. In the spirit of Model Predictive Control our method allows for time-optimal replanning at every time step with only one computation per forecast. Using the recent Multi-Time Hamilton-Jacobi Reachability formulation we obtain a value function which is then used for closed-loop control. We evaluate the reliability of our method empirically over a large set of multi-day start-target missions in the ocean currents of the Gulf of Mexico with realistic forecast errors. Our method outperforms the baselines significantly, achieving high reliability, measured as the success rate of navigating from start to target, at low computational cost.

I. INTRODUCTION

We increasingly deploy autonomous systems in the air and oceans. Beyond airplanes and ships, there are emerging applications such as balloons in the stratosphere for delivering internet access [1], airships, ocean gliders and active drifters for collecting ocean data [2]–[5], floating solar farms storing energy in fuels [6], and floating seaweed farms for biomass and carbon sequestration [7], [8].

The overactuated navigation approaches used in ships and planes require significant power to overcome the drag forces inherent in navigating through fluids. The power required scales with $P = F_{Drag} \cdot \frac{\Delta s}{\Delta t} \propto A_C \cdot v^3$, where A_C is the cross sectional area and v the velocity relative to the surrounding fluid. This makes overactuated control prohibitively expensive for energy constraint applications such as long duration environmental monitoring systems and floating structures with large cross-sectional areas. Hence, we investigate an energy-efficient steering paradigm leveraging the winds and ocean currents around these systems: *navigating agents by hitchhiking complex flows*. By using the drift of these non-linear, time-varying flows for propulsion, only a minimal amount of energy is required to nudge the agents into beneficial flows, e.g., balloons going up and down to use

¹M.W., and C.J.T. are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, USA. For inquiries contact: mariuswiggert@berkeley.edu

²M.D. and P.F.J.L. are with the Department of Mechanical Engineering at the Massachusetts Institute of Technology, USA.

The authors gratefully acknowledge the support of the C3.ai Digital Transformation Institute, the NASA ULI Grant on Safe Aviation Autonomy, the DARPA Assured Autonomy Program, the SRC CONIX Center, and the ONR BRC program.

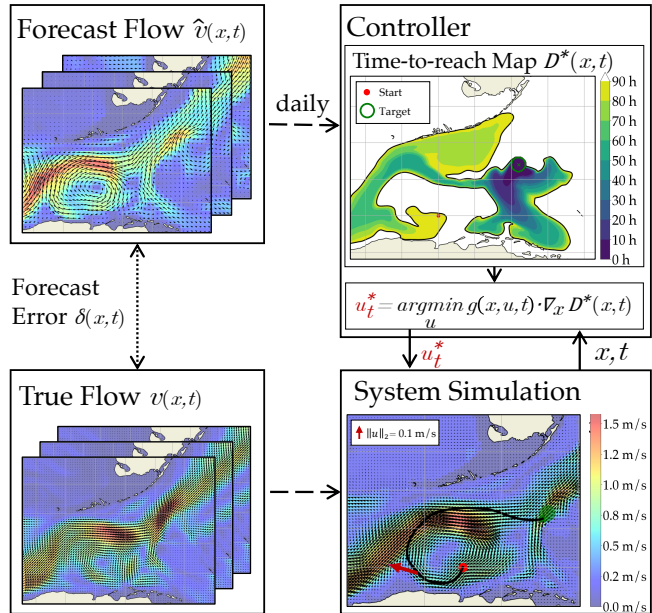


Fig. 1: Our method for reliable navigation leveraging flows is based on frequent replanning. For that the Time-to-reach map D^* is computed daily as new flow forecasts $\hat{v}(x, t)$ become available. Then for closed-loop operation, the time-optimal control u_t^* is calculated from D^* which is equivalent to replanning at every step. The system is simulated using the true flow $v(x, t)$ which differs from the forecast by the forecast error $\delta(x, t)$.

different winds or small horizontal propulsion in the oceans to navigate bifurcations in the flow. Because the required power scales cubically with the relative velocity, navigating with $\frac{1}{10}$ th of the speed means only $\frac{1}{1000}$ th of the power is required which enables a host of new applications.

From a control perspective, there are three core challenges when navigating by hitchhiking these flows. First, the dynamics of the wind and ocean flows are non-linear and time-varying. Second, in realistic scenarios only coarse forecasts of the flows are available and these differ from the true currents [9]–[12]. Third, when agents are underactuated they cannot easily compensate for these forecast errors with classical methods such as robust control.

There is a rich literature on planning time- and energy-optimal paths in flows both in the oceans [13]–[32] and the air [33]–[36]. For planning in *known flows* researchers have proposed Hamilton-Jacobi (HJ) reachability [13]–[15], non-linear programming [16], [17], evolutionary algorithms [18], and graph-based techniques such as A* [19], [20], [35], RRT [21], [34], and time-varying Dijkstra [22]. Non-linear programming, evolutionary algorithms, and graph-based meth-

ods suffer from discretization error and the non-convexity of the optimization problem which can cause solvers to get stuck in local minima and infeasible solutions. In contrast, HJ reachability is guaranteed to obtain time-optimal paths when the flows are known, as it solves the exact continuous-time control problem via dynamic programming.

As the wind and ocean forecasts are never perfect, different paradigms and their combinations have been explored for navigation when the true flow is not known. When we have access to a *realistic distribution of the flows* we can use probabilistic methods to optimize the expected energy or travel time of a policy using stochastic reachability [37] or Markov Decision Processes (MDP) [23], [33]. Previous work also explored risk-optimal path planning in this stochastic setting [24]–[26]. Unfortunately, for ocean surface currents, only daily deterministic forecasts are available from the leading ocean forecasting providers HYCOM [10] and EU Copernicus Marine [11]. There are ocean dynamics models with stochastic forcing to compute accurate probabilistic forecasts [12], [38], [39], however these are too computationally expensive for real-time. A heuristics approach is then to assume Gaussian noise [23] or to increase prediction uncertainty around high velocity currents [26] which does not capture the complexity of the forecast errors.

Another paradigm is to plan on the deterministic forecast flows and follow the planned path with tracking algorithms which compensate for the drifts and can guarantee a tracking error bound [32]. While this is important in tight spaces around obstacles, minimizing tracking error does not necessarily lead to time and energy optimal paths. The Model Predictive Control (MPC) paradigm uses frequent replanning from the positions of the agent to handle the dynamics error [19], [40]. In the robust control methods we assume a bound on the forecast error to derive a path and controls that reach the target despite bounded adversarial disturbances using reach-avoid HJ reachability formulations [41] or approximations [19]. However, for underactuated agents often there exists no robust control for realistic forecast error bounds.

Lastly, recent work has explored the application of deep Reinforcement Learning (RL) to navigate in flows [28]–[30]. The controller trained by Gunnarson et. al. on vortical flows using only local current information navigates successfully in these flows but unsurprisingly it fails in other flow structures [29]. The Loon team trained an RL agent for station-keeping of balloons with forecasts as inputs and it performed well in long-duration real-world experiments after training on an immense distribution of flows [36].

In this work, we focus on the problem of *reliable* navigation of underactuated agents leveraging flows in the realistic setting when regular deterministic forecasts are available. We define reliability empirically as *the success rate of a controller in navigating from a start point to a target region over a set of start-target missions*, as developed in Section II-B. In this paper we make three core contributions: (1) We propose a control approach enabling full time-horizon replanning at every time step with a single computation. For that we build on the recent Multi-Time HJ Reachability formulation [42]

for computing a time-optimal value function (Time-to-reach) on the latest forecast and use it for closed-loop control. This can be thought of as full-time horizon MPC at every step; (2) We are the first to evaluate and compare the reliability of closed-loop control schemes for underactuated agents in ocean flows in the setting of daily forecasts with realistic forecast error. We evaluate performance across a large set of multi-day start-to-target missions distributed spatially across the Gulf of Mexico and temporally across four months using HYCOM and Copernicus Ocean Forecasts [10], [11] We compare several methods on this dataset across multiple metrics and find that our control architecture significantly outperforms other methods in terms of *reliability*; (3) We quantify how the reliability of various control methods is affected by the forecast error.

This paper is organized as follows: In Sec. II we define the problem; followed by Sec. III which details the proposed control architecture and our algorithm to compute the time-optimal value function. Sec. IV contains the closed-loop performance evaluation of our methods and baselines and we conclude with Sec. V and outline future work.

II. PROBLEM STATEMENT

In this section, we define the problem in terms of the agent and flow model and further explain and justify the notion of reliability as our performance measure.

A. Agent and Flow Dynamics

We consider an agent operating in a general time-varying, non-linear flow field $v(\mathbf{x}, s) \rightarrow \mathbb{R}^n$ where $\mathbf{x} \in \mathbb{R}^n$ represents the state, s the time and n the dimensionality of the domain e.g. $n = 2$ for a surface vehicle on the ocean and $n = 3$ for agents operating in the atmosphere or underwater. Let the agent’s actuation signal be denoted by \mathbf{u} from a bounded set $\mathbb{U} \in \mathbb{R}^{n_u}$ where n_u is the dimensionality of the control. Then the dynamics of the system with given initial conditions are governed by an Ordinary Differential Equation (ODE) of the following form

$$\dot{\xi}(s) = f(\xi(s), \mathbf{u}(s), s) = v(\xi(s), s) + g(\xi(s), \mathbf{u}(s), s), s \in [0, T] \quad (1)$$

where ξ represents the trajectory and $\xi(s) \in \mathbb{R}^n$ the state at time s . For more intuitive notation we use \mathbf{x} for the state whenever possible. The system dynamics I are further assumed to be continuous, bounded and Lipschitz continuous in ξ uniformly in \mathbf{u} [43]. The movement of the agent in the flow depends on its control $g(\mathbf{x}(s), \mathbf{u}(s), s)$ and the drift of the surrounding flow $v(\mathbf{x}(s), s)$. The control can be holonomic when the agent can directly actuate in each dimension e.g. $g(\mathbf{x}, \mathbf{u}(s), s) = \mathbf{u}$ or non-holonomic e.g. a balloon can only actuate up and down along the vertical axis. This makes the common assumption that the drift of the agent directly affect its state and neglects any inertial effects.

B. Problem Setting

The goal of the agent is to navigate *reliably* from a start state \mathbf{x}_0 to a target region $\mathcal{T} \in \mathbb{R}^n$ while being underactuated

$\max_u \|g(\mathbf{x}(s), \mathbf{u}(s), s)\|_2 \ll \|v(\mathbf{x}(s), s)\|_2$ most of the time. During operation the agent is given a forecast of the flow $\hat{v}(\mathbf{x}, s)$ which differs from the true flow $v(\mathbf{x}, s)$ by the stochastic forecast error $\delta(\mathbf{x}, s; \omega)$ where ω is a random variable. The error field δ can be characterized based on different metrics such as Root Mean Squared Error (RMSE) of the velocities or vector correlation [9]. The agent receives a new forecast at regular intervals (typically daily) that can be used to improve performance.

Our goal is reliable navigation of underactuated agents in realistic complex flows occurring in nature. The strongest notion of reliability is robustness to a bounded disturbance, which guarantees reaching the target despite a worst-case forecast error δ . However, proving robustness is not possible in our setting where the agent is significantly underactuated and the average forecast error is larger than the actuation. Nevertheless, we compare against a robust control baseline in Sec. IV. A weaker notion of reliability is a probabilistic bound, i.e. the agent reaches the target with high probability. Probabilistic bounds could be established by making strong assumptions on the distribution of the forecast error fields δ and using a simple flow field; this would render the results less meaningful for the realistic settings we consider.

For these reason we define *reliability* empirically as the success rate of a controller navigating from a start point to a target \mathcal{T} over a set of start-target missions \mathbb{M} in realistic flows. If the agent reaches the target \mathcal{T} within a maximum allowed time T_{max} the mission is successful, otherwise it failed. In our experiments we sample missions \mathbb{M} over a large spatial region and over a period of four months.

III. METHOD

In this section, we first outline the motivation behind our method and then detail our closed-loop control strategy.

A. Motivation

In realistic settings, only deterministic flow forecasts are available to the agent. To ground our discussion we look at ocean currents which are on the order of magnitude of $0.5 - 2 \frac{m}{s}$ and underactuated agents with limited actuation $\|\max(g(\mathbf{u}, \mathbf{x}, s))\|_2 = 0.1 \frac{m}{s}$. The forecast error of the ocean current forecasts by the HYCOM model is estimated to be $RMSE(\delta) = 0.2 \frac{m}{s}$ after extensive validation analysis [9], [44]. In this challenging setting we can neither use probabilistic methods without making unrealistic assumptions about the error distribution δ nor can we apply robust control as the average error δ is larger than the actuation of the agent. How than can we achieve reliable navigation in this setting?

Our approach builds on the MPC paradigm of regular replanning with deterministic dynamics to compensate for imperfect knowledge of the dynamics and achieve reliable navigation. Intuitively, the higher the frequency of replanning, the more we can adapt the control to the dynamics experienced. While we could use any planning algorithm for non-linear time-varying dynamics, as mentioned in the related work, HJ reachability is the state-of-the-art, as it guarantees finding the optimal solution and can even handle

time-varying obstacles and targets [13], [41], [43]. While we can obtain time-optimal trajectories from both classic forward and backward HJ reachability, the value function of backwards reachability is more useful. The backwards reachability objective is to minimize the distance of the agent to the target set at a terminal time. A key insight is that the value function of this objective can be used for closed-loop control as for every state and time in the domain we can extract the optimal control that minimizes this objective. This provides a notion of replanning at every step even when the deterministic dynamics are not accurate. However, there is a problem with directly applying classic backwards HJ reachability: the value function that is minimized is the distance to the target at a *fixed terminal time*. This poses the problem of which terminal time to choose to calculate the value function? If we choose it too distant in the future the system will "loiter", not making progress towards the goal. If we choose it too close to the current time, it might be impossible to reach the target and the agent will minimize the distance to the target in the short term, potentially at the cost of long term progress. We can compute the earliest possible arrival time, using forward reachability, and then compute backwards reachability from that time to get our value function for closed loop control (a baseline in IV). A more elegant approach is the recent *multi-time reachability* formulation which requires only one backwards computation and produces a value function that yields the time-optimal control everywhere, not just at the zero level-set, as the classic reachability value function. We found that using time-optimal control is a good proxy for reliability in our setting.

B. Multi-Time Reachability for Closed-Loop Control

In the following we summarize the multi-time reachability technique for completeness. Details for more general systems and applications are available in [42].

Multi-time reachability uses dynamic programming to derive a controller that (a) if possible, will get the system to the target in the minimum time, and (b) if not, will get as close to the target as possible. In order to achieve this behavior, we define the following cost function J :

$$J(\mathbf{x}, \mathbf{u}(\cdot), t) = \underbrace{d(\xi_{t,\mathbf{x}}^{\mathbf{u}(\cdot)}(T), \mathcal{T})}_{\text{Terminal distance from target set}} - \underbrace{\int_t^T \mathbb{I}_{\mathcal{T}}(\xi_{t,\mathbf{x}}^{\mathbf{u}(\cdot)}(s)) ds}_{\text{Time spent in target set}} \quad (2)$$

where $d(x, \mathcal{T})$ is a distance metric of a point x to the target set \mathcal{T} , $\xi_{t,\mathbf{x}}^{\mathbf{u}(\cdot)}(s)$ is the position of the agent at time s , when starting at \mathbf{x} at time t with dynamics given in (1). $\mathbb{I}_{\mathcal{T}}(x)$ is the identity function that is 1 when the state is in the target $x \in \mathcal{T}$ and 0 otherwise. The consequence of this cost function is that if the agent can reach the destination, minimizing the cost implies reaching as quickly as possible. If the agent cannot reach the destination, the optimal control will attempt to reduce the terminal distance to the target.

Note that for underactuated systems we want to consider the dynamics of the agent only until it reaches the target.

Within the target we switch off the dynamics to reward the agent for staying in the target.

Given this cost function, we use the principle of dynamic programming to derive an HJ Partial Differential Equation (PDE) whose viscosity solution is [42]:

$$\frac{\partial J^*(\mathbf{x}, t)}{\partial t} = \begin{cases} 1 & \mathbf{x}(t) \in \mathcal{T} \\ -\min_{\mathbf{u}} [\nabla_{\mathbf{x}} J^* \cdot f(\mathbf{x}, \mathbf{u}, t)] & \text{otherwise} \end{cases} \quad (3)$$

$$J^*(\mathbf{x}, T) = d(\mathbf{x}, \mathcal{T})$$

Since the value of J^* contains information about the minimum time it takes to reach the destination, we can extract an informative time-to-reach map \mathcal{D}^* from it:

$$\mathcal{D}^*(\mathbf{x}, t) = T + J^*(\mathbf{x}, t) - t, \quad \forall(\mathbf{x}, t) \text{ s.t., } J^*(\mathbf{x}, t) \leq 0 \quad (4)$$

If the target can be reached from \mathbf{x} starting at t (implied by $J^*(\mathbf{x}, t) \leq 0$), then $\mathcal{D}^*(\mathbf{x}, t)$ is the minimum duration required to reach \mathcal{T} . Which means that inside the target $\mathcal{D}^*(\mathbf{x}, t) = 0 \forall t, \mathbf{x} \in \mathcal{T}$. The optimal control $\mathbf{u}^*(\mathbf{x}, t)$ is then the value that minimizes the Hamiltonian.

$$\mathbf{u}^*(\mathbf{x}, t) = \arg \min_{\mathbf{u} \in \mathbb{U}} f(\mathbf{x}, \mathbf{u}, t) \cdot \nabla_{\mathbf{x}} J^*(\mathbf{x}, t) \quad (5)$$

$$= \arg \min_{\mathbf{u} \in \mathbb{U}} g(\mathbf{x}, \mathbf{u}, t) \cdot \nabla_{\mathbf{x}} J^*(\mathbf{x}, t) \quad (6)$$

Where (6) follows from (5) because $v(\mathbf{x}, t)$ does not depend on \mathbf{u} . Note that $\nabla_{\mathbf{x}} \mathcal{D}^*(\mathbf{x}, t) = \nabla_{\mathbf{x}} J^*(\mathbf{x}, t)$ which means we can use $\mathcal{D}^*(\mathbf{x}, t)$ and $J^*(\mathbf{x}, t)$ interchangeably to obtain the optimal control. This formulation solves the two key issues with classic backwards reachability as described above: a) we do not need to fix an arrival time *a priori*, instead we can run backwards Multi-Time Reachability from a large maximum time backwards until the current time t ; and b) the control obtained from this value function provides the time-optimal control at all states making it more useful for reliable navigation.

The key insight here is that, the value function allows us to extract the time-optimal control for each state \mathbf{x} at each time t thereby enabling frequent replanning at low computational cost. This use of the reachability value function differs significantly from the ways it has been used in the literature. Much of the previous work in reachability-based control has focused on least restrictive control for safety specifications [43], where the terminal cost encodes constraints that must be satisfied for safety, and the solution of the corresponding HJ PDE provides both a reachable set of states that satisfy the constraints, as well as the control \mathbf{u}^* to apply at the boundary of the reachable set in order to stay within this safe set. Such a least restrictive framework has been applied to both safe trajectory planning [45]–[47] and learning-based control [48], [49].

In this work, we propose to solve this Multi-time reachability problem once per forecast, which are received at regular (daily) intervals (Fig. 1 and Alg. 1). The most recent time-to-reach map is then used for closed-loop control which by construction provides reliability through what can be thought of as time-optimal replanning at every time-step.




Fig. 2: In time-invariant forecast flows \hat{v} the time-to-reach map $\mathcal{D}^*(\mathbf{x}, t)$ is static. The level-sets of $\mathcal{D}^*(\mathbf{x}, t)$ indicate how fast an agent can reach the target from a specific state \mathbf{x} . With holonomic actuation the optimal control is the spatial gradient $\mathbf{u}_t^* = \nabla_{\mathbf{x}} \mathcal{D}^*(\mathbf{x}, t)$. Applying \mathbf{u}_t^* closed loop leads to reaching the target region \mathcal{T} even if the true currents are different $v_t \neq \hat{v}_t$. This schema is equivalent to time-optimal replanning at every step which leads to its reliability.

For an intuitive example, take the setting in Fig. 2: at \mathbf{x}_0 the agent applies \mathbf{u}_0^* based on the planned time-to-reach map. As the true currents v are different than the forecast \hat{v} , the agent finds itself at \mathbf{x}_1 , a different state than expected. Based on our forecast the time-optimal control from this state onwards can again directly be computed from the time-to-reach map.

There are two advantages of this method over classic MPC replanning based on fast non-linear programming or graph-based methods: (a) much higher replanning frequencies are possible because deriving the optimal control from the value function is computationally cheap compared to solving the optimal control problem from a new state at every time step, and (b) less discretization error because the HJ PDE solves the continuous-time problem whereas the classic MPC methods rely on discretization in time and space to enable fast planning in the loop.

Algorithm 1: Multi-Time HJ Closed-loop Schema

Input: Forecast Flow(s) $\hat{v}(\mathbf{x}(s), s)$, $t = 0$, $\mathbf{x}_t = \mathbf{x}_0$

while $t \leq T_{max}$ and $\mathbf{x}_t \notin \mathcal{T}$ **do**

if new forecast available **then**

 compute time-to-reach map \mathcal{D}^*

 Use latest \mathcal{D}^* for control

$\mathbf{u}_t^* = \arg \min_{\mathbf{u} \in \mathbb{U}} g(\mathbf{x}_t, \mathbf{u}, t) \cdot \nabla_{\mathbf{x}} \mathcal{D}^*$

$\mathbf{x}_{t+1} = \mathbf{x}_t + \int_t^{t+1} f(\mathbf{u}, \mathbf{x}(s), s) ds$

IV. EXPERIMENTS

In this section we evaluate our control schema of using multi-time HJ reachability for closed-loop control and compare it to baseline methods on realistic ocean currents.

A. Experimental Set-Up

We investigate the reliability of various controllers in navigating a two dimensional Autonomous Surface Vehicle (ASV) with holonomic actuation of fixed magnitude $\|g(\mathbf{u}, \mathbf{x}, t)\|_2 = \|\mathbf{u}\|_2 = 0.1 \frac{m}{s}$. The control is the thrust angle θ and the ASV is navigating in strong ocean currents $v(\mathbf{x}, t) \in [0.3 \frac{m}{s}, 2 \frac{m}{s}]$ which it wants to hitchhike to get to the targets. In the following we describe how we ensure realistic ocean forecast simulation and obtain a large set of missions. Then we explain the baselines and evaluation metrics.

a) *Realistic Ocean Forecast Simulation:* The Ocean forecast data we employ are the HYCOM forecast and hindcast [10] and hindcasts from Copernicus [11] for the Gulf of Mexico region. To simulate realistic conditions we provide the control methods daily with the HYCOM forecast as it becomes available while simulating the system dynamics with the hindcast as the true flow $v(\mathbf{x}, x)$ (Fig. 1). We investigate two settings (a) planning on HYCOM forecasts and simulating on HYCOM hindcasts (HYCOM-HYCOM) and (b) planning on HYCOM forecasts and simulating on Copernicus hindcasts (HYCOM-Copernicus) (IV-B).

To estimate how realistic our simulations are we compare the simulated forecast error δ across our start-target mission set \mathbb{M} with the HYCOM forecast error as estimated by Metzger et. al. using extensive drifter buoy data [9]. In Fig 3 we visualize two metrics, the velocity RMSE and the vector correlation, where 2 represents perfect correlation and 0 no correlation [44]. We find that the HYCOM-HYCOM setting underestimates the forecast error, especially in the first 24h where the forecast is perfect. The HYCOM-Copernicus setting is realistic as the simulated forecast error is of similar magnitude as the actual HYCOM forecast error.

b) *Large Representative Set of Missions:* To obtain a set of start-target missions \mathbb{M} we first fix 18 regularly spaced starting times t_i between November 2021 and February 2022. For each starting time t_i we uniformly sample 16 start points $x_{Start}^{t_i, j}$ spatially over the Gulf of Mexico. In our underactuated setting many start-target missions are impossible even if the true currents are known. Hence, for the test set \mathbb{M} we need to ensure each mission is fundamentally feasible given the true currents. To generate only feasible missions from each starting points $x_{Start}^{t_i, j}$ we calculate the forward reachable set (FRS) starting at t_i for a maximal time-horizon of $T_{max} = 120h$ using HJ reachability. The FRS is the set of all states x_s at time s for which there exists a control signal $u(\cdot)$ such that $\xi_{t_i, x_{Start}^{t_i, j}}^{u(\cdot)} = x_s$. To get a variety of mission durations we sample 4 relative times $\Delta t_k \in [20, 120]h$ and sample a target point $x_{T}^{i, j, k}$ from within the forward reachable set at $s = t_i + \Delta t_k$. Then $x_{T}^{i, j, k}$ is guaranteed to be reachable from $x_{Start}^{t_i, j}$ and the target \mathcal{T} is the circular region of radius

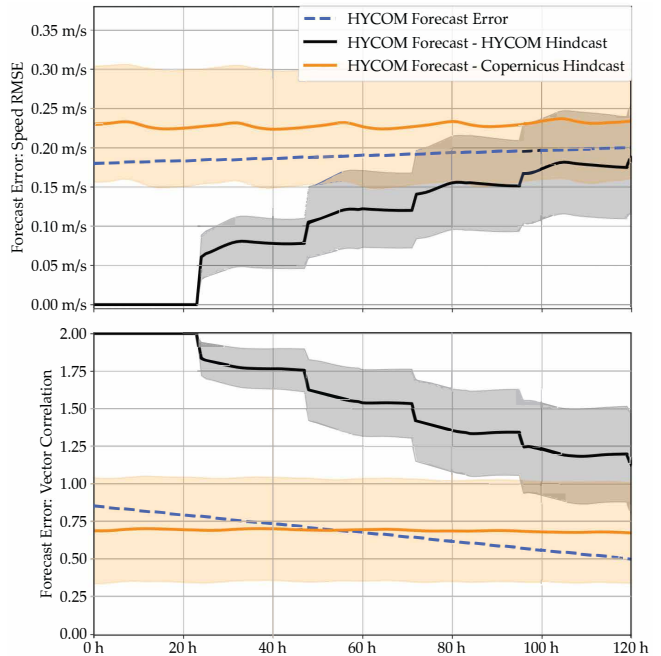


Fig. 3: The simulated forecast errors in our experiments are of similar magnitude as the actual HYCOM forecast errors [9]. The graphs show the mean and standard deviation of two forecast error metrics, velocity RMSE and vector correlation, over our mission set \mathbb{M} for the two simulation settings.

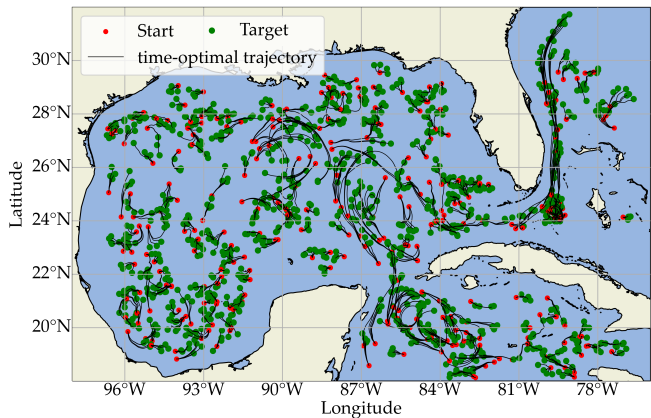


Fig. 4: We evaluate our method over a large set of 1152 start-target missions spatially distributed over the Gulf of Mexico and temporally from November 2021 until February 2022.

$r_{\mathcal{T}} = 0.1^\circ$ around it. This gives us a large and diverse start-target set \mathbb{M} of 1152 missions ranging in duration from 20h to 120h as visualized in Fig. 4, together with their respective time-optimal trajectories based on the true currents.

c) *Baseline Controllers:* We compare the performance of our multi-time HJ reachability closed-loop strategy with four baselines. The first baseline is the HJ Closed-Loop method which is the same as our method except that it employs the classic backwards reachability value function. The backwards reachability value function is calculated based on the earliest possible arrival time and used for closed-loop control, as explained in Sec. III. Our second

baseline (Robust Hamilton-Jacobi-Isaacs (HJI) Closed-Loop) is a robust controller which is equivalent to HJ Closed-Loop but instead of solving for the classic backward reachability value function it solves for the reach-avoid value function with a small bounded disturbance of $d = 0.05 \frac{m}{s}$ [41]. This means that the control extracted from it in closed-loop is a conservative control that can arrive at \mathcal{T} or get as close as possible to \mathcal{T} at the earliest arrival time despite worst-case disturbance d . The third method (HJ + Waypoint Tracking) we compare with is based on the idea of tracking waypoints to compensate for the forecast error. It plans the time-optimal path on each new forecast using HJ and at each time actuates towards the next waypoint in this plan. Lastly, we compare against a Naive-to-Target approach which ignores the forecasts altogether and always actuates towards the center of the target region. Note that we do not compare against any graph-based methods such as RRT and A* because the only approximate the time-optimal control problem which HJ reachability solves in continuous time and space. Hence, we expect them to be strictly worse.

d) Evaluation Metrics: The key performance metric for our controllers is *reliability*, defined as the success rate of a controller over the set of missions \mathbb{M} (Sec. III). To further elucidate the controller performances we look at two additional metrics. Lastly, we investigate how the reliability of a controller changes with the forecast error.

To evaluate how fast a controller reaches a target we compute the time-optimality ratio ρ . We define it as the time it took the controller to reach \mathcal{T} divided by the fastest it could have reached \mathcal{T} under perfect knowledge of the currents. It is calculated over the set of all missions in which the controller succeeds ($\tilde{\mathbb{M}}_{ctrl}$) as

$$\rho = \frac{1}{|\tilde{\mathbb{M}}_{ctrl}|} \sum_{\mathbb{M}_{ctrl}} \frac{\Delta t_{arrival\ ctrl}}{\Delta t_{arrival\ best-in-hindsight}} \quad (7)$$

A value of 1 means the controller was as fast as possible and 1.1 means it took 10% longer than the fastest possible.

The minimum distance to target measures the closest the controller got to \mathcal{T} during the full simulation horizon T_{max} in degree latitude, longitude. The minimum distance to target is positive, except for when the controller succeeds in a mission, and then it is 0.

Lastly, our diverse set of missions \mathbb{M} allows us to investigate how the reliability of various controllers changes with increasing forecast error. For that we calculate the average forecast error $RMSE(\delta)$ for each mission spatially over a regional box containing the start and target region and temporally over the 5 day horizon of each forecast. We then group the missions by their $RMSE(\delta)$ into 20 bins and calculate for each controller the success rate across each bin. To make the trends more visible we fit a weighted linear regression over these 20 success rate - $RMSE(\delta)$ points and weight each bin by its number of missions.

B. Experimental Results

a) Simulation Setting HYCOM-HYCOM:

In this setting we evaluate the controller performance over

$|\mathbb{M}| = 1152$ start-target missions and run the simulation for $T_{max} = 150h$. If the controller reaches the target region within that time, it is successful otherwise it failed. Our control approach achieves a success rate of 99% and outperforms the baselines (Table I). The time to reach the target is on average only 3% higher than the fastest possible, here HJ Closed-Loop performs slightly better with 2%. The Naive-to-Target controller succeeds only in 84.9% of missions. These results highlight that highly reliable navigation is possible with low (short-term) forecast error. However, as of now the actual forecast error of current ocean models is significantly higher (Fig 3) which makes this an optimistic performance estimate. To evaluate the statistical significance of our results we perform a one-sided two sample z proportion test for each of the controllers against the Naive-to-Target baseline. With Γ denoting the success rate of a controller, our null hypothesis is $H_0 : \Gamma_{Naive-to-Target} = \Gamma_{controller}$ and the alternative hypothesis is $H_A : \Gamma_{Naive-to-Target} < \Gamma_{controller}$. We obtain that the success rate of all controllers is higher than Naive-to-Target in a statistically significant way (p-values Multi-Time HJ CL $p = 3.9e^{-36}$, HJ CL $p = 2.8e^{-27}$, Robust HJI CL $p = 1.36e^{-22}$, HJ + Waypoint Tracking $p = 3.65e^{-15}$).

b) Simulation Setting HYCOM-Copernicus:

To ensure comparability across settings we take the same set of missions \mathbb{M} , however with the Copernicus hindcasts only 837 of the missions are fundamentally feasible, which makes the set \mathbb{M} smaller. Again our control approach achieves a success rate of 82.3% and outperforms the baselines but with less of a margin than in HYCOM-HYCOM (Table I). We perform the same statistical significance test and find only our approach has statistically significant higher success rate than Naive-to-Target ($p = 0.012$). We want to emphasize that a 4.4% increase over the baseline in this challenging setting with large forecast errors is a sizeable improvement. Moreover, the results from HYCOM-HYCOM indicate that performance improves with better forecasts. Our closed-loop control schema is easily extendable to include learning about the currents while operating in them and using this information to improve its estimate of the future currents and thereby improve performance.

Figures 5 show how the success rate of the controllers changes for missions with varying forecast errors. As expected we see that the success rate decreases with increasing forecast error with different slopes for different controllers. The performance of our Multi-Time HJ Closed-Loop approach decreases slower than the baselines. In the HYCOM-HYCOM setting it is almost unaffected by higher forecast error. Note that we would expect Naive-to-Target to be indifferent to the forecast error as it does not consider the forecast. However, for Naive-to-Target we observe a significant drop in performance for missions with high forecast error. We hypothesize that the forecast error is higher in regions with complex and strong currents, conditions which are inherently more challenging to navigate. This could explain why Naive-to-Target fails more frequently with higher forecast error.

	Success Rate	Time-Opt.	Min. Dist.
Plan on HYCOM Forecasts – true flows HYCOM Hindcasts			
Multi-Time HJ CL	99.0%*	1.0332	0.060°
HJ Closed-Loop	97.6%*	1.0207	0.104°
Robust HJI CL	96.6%*	1.0346	0.154°
HJ + Waypt. Tracking	94.7%*	1.1412	0.060°
Naive-to-Target	84.9%	1.0752	0.068°
Plan on HYCOM Forecasts – true flows Copernicus Hindcasts			
Multi-Time HJ CL	82.3%*	1.129	0.107°
HJ Closed-Loop	79.1%	1.106	0.110°
Robust HJI CL	78.6%	1.111	0.106°
HJ + Waypt. Tracking	67.1%	1.373	0.091°
Naive-to-Target	77.9%	1.052	0.078°

TABLE I: We compare the performance of multiple controllers in two forecast settings. Our Multi-Time HJ Closed-Loop (CL) controller outperforms the baselines in reliability, the key focus of this work. The * marks statistically significant higher success rate compared to Naive-to-Target. The average time-optimality ratio indicates how fast the controller reached the target relative to the best possible in-hindsight. Minimum distance measures the closest the controller got to the target region in degree lat, lon.

V. CONCLUSION AND FUTURE WORK

In this work we have demonstrated that planning with Multi-Time HJ Reachability on daily forecasts and using the value function for closed-loop control enables reliable navigation of underactuated agents leveraging complex flows. The reliability of our method stems from the fact that the optimal control extracted from the value function at every time step is equivalent to full horizon time-optimal replanning. There are two key advantages over classic MPC with non-linear programming or graph-based methods. First, our method has lower computational cost as it only requires computing the value function once per forecast instead of having to solving an optimal control problem at every time-step. Second, our method solves the continuous-time optimal control problem and does not require spatial or temporal discretization typically employed by MPC for fast computation. This leads to high reliability at low computational cost which enables reliable autonomy for resource-constrained systems navigating in flows in the atmosphere and the oceans.

We evaluated the performance of our method in realistic ocean currents over a large set of multi-day start-to-target missions distributed spatially across the Gulf of Mexico and temporally across four months. In the setting of using forecasts from HYCOM and simulating the true currents with HYCOM hindcasts, our method achieves a 99% success rate and outperforms the baselines. However, this setting underestimates the actual HYCOM forecast error. Hence, we also evaluated our method in a setting with forecast errors that reflect more realistic operations: planning on HYCOM forecasts and simulating the true currents with Copernicus hindcasts. In this more challenging setting our method again outperforms the baselines achieving a 82.3% success rate. While we showcased our method on 2D ocean currents, we want to emphasize that it is directly applicable to other

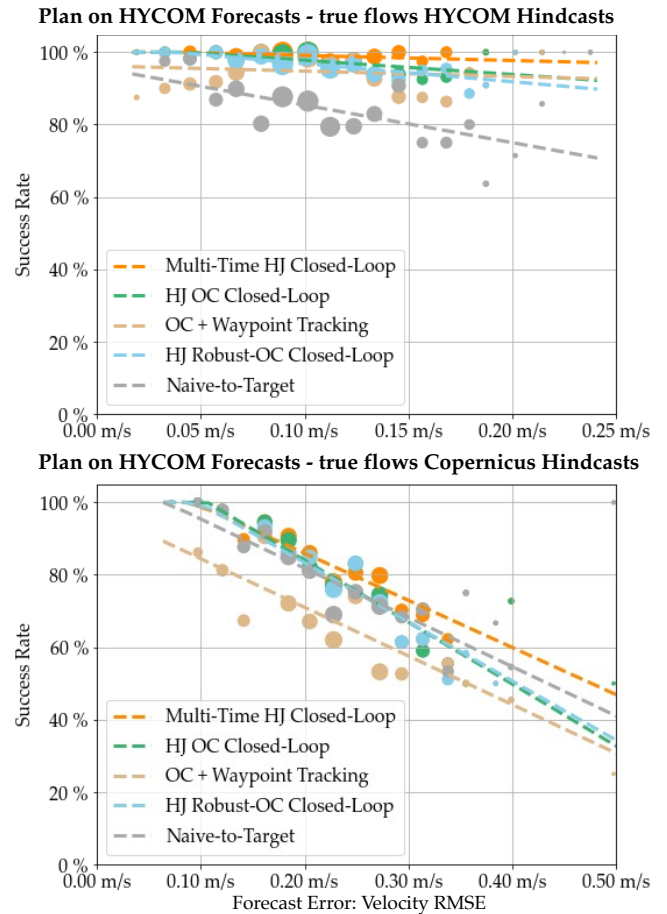


Fig. 5: The reliability of different controllers decreases for missions with higher forecast error. The dots represent the average success rate over a subset of missions with mean forecast error plotted on the x axis. The size indicates the number of missions in the respective subset. The lines indicate the trends and represent the linear regression of the success-rate to forecast error points weighted by their number of missions.

flows e.g. 3D flows in the air and underwater. Lastly, we demonstrated that our method is less affected by increasing forecast errors than the baselines.

The stark difference between 99% and 82.3% success rate in the two simulation settings highlights that more accurate forecasts increase reliability. Hence, we anticipate that if the agent learns about its surrounding flow field using online flow measurements to improve its estimate of the future flows, it can increase reliability significantly above the current 82.3% in realistic settings. Therefore, our future work we will focus on this learning of the surrounding currents which can directly be integrated in our method. To further validate our method, we plan to perform field tests with multiple autonomous surface vehicles in the ocean.

REFERENCES

- [1] "Project loon," X, the moonshot factory. [Online]. Available: <https://x.company/projects/loon/>

- [2] M. Manikandan and R. S. Pant, "Research and advancements in hybrid airships - a review," *Progress in Aerospace Sciences*, vol. 127, p. 100741, 2021.
- [3] D. Meyer, "Glider technology for ocean observations: A review," *Ocean Science Discussions*, pp. 1–26, 2016.
- [4] A. Molchanov *et al.*, "Active drifters: Towards a practical multi-robot system for ocean monitoring," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 545–552.
- [5] P. F. J. Lermusiaux *et al.*, "A future for intelligent autonomous ocean observing systems," *Journal of Marine Research*, vol. 75, no. 6, pp. 765–813, Nov. 2017, the Sea. Volume 17, The Science of Ocean Prediction, Part 2.
- [6] B. D. Patterson *et al.*, "Renewable CO₂ recycling and synthetic fuel production in a marine environment," *Proceedings of the National Academy of Sciences*, vol. 116, no. 25, pp. 12 212–12 219, 2019.
- [7] "Floating seaweed farms," Phykos PBC. [Online]. Available: <https://www.phykos.co>
- [8] M. S. Bhabra *et al.*, "Optimal harvesting with autonomous tow vessels for offshore macroalgae farming," in *OCEANS 2020 IEEE/MTS*. IEEE, Oct. 2020, pp. 1–10.
- [9] E. Metzger *et al.*, "Validation test report for the global ocean forecast system 3.5-1/25 degree hycom/cice with tides," NAVAL RESEARCH LAB WASHINGTON DC WASHINGTON United States, Tech. Rep., 2020.
- [10] E. P. Chassignet *et al.*, "Us godae: global ocean prediction with the hybrid coordinate ocean model (hycom)," *Oceanography*, vol. 22, no. 2, pp. 64–75, 2009.
- [11] EU Copernicus Marine Service Information, "Global ocean 1/12 physics analysis and forecast updated daily product [data set]." [Online]. Available: <https://doi.org/10.48670/moi-00016>
- [12] P. F. J. Lermusiaux *et al.*, "Quantifying uncertainties in ocean predictions," *Oceanography*, vol. 19, no. 1, pp. 92–105, 2006.
- [13] T. Lolla *et al.*, "Path planning in time dependent flow fields using level set methods," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 166–173.
- [14] D. N. Subramani *et al.*, "Time-optimal path planning: Real-time sea exercises," in *Oceans '17 MTS/IEEE Conference*, Aberdeen, Jun. 2017.
- [15] P. F. J. Lermusiaux *et al.*, "Optimal planning and sampling predictions for autonomous and Lagrangian platforms and sensors in the northern Arabian Sea," *Oceanography*, vol. 30, no. 2, pp. 172–185, Jun. 2017, special issue on Autonomous and Lagrangian Platforms and Sensors (ALPS).
- [16] W. Zhang *et al.*, "Optimal trajectory generation for a glider in time-varying 2d ocean flows b-spline model," in *2008 IEEE International Conference on Robotics and Automation*. IEEE, 2008, pp. 1083–1088.
- [17] D. Jones and G. A. Hollinger, "Planning energy-efficient trajectories in strong disturbances," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 2080–2087, 2017.
- [18] M.-C. Tsou and H.-C. Cheng, "An ant colony algorithm for efficient ship routing," *Polish Maritime Research*, 2013.
- [19] V. T. Huynh *et al.*, "Redictive motion planning for auvs subject to strong time-varying currents and forecasting uncertainties," in *2015 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2015, pp. 1144–1151.
- [20] D. Kularatne *et al.*, "Going with the flow: a graph based approach to optimal path planning in general flows," *Autonomous Robots*, vol. 42, no. 7, pp. 1369–1387, 2018.
- [21] D. Rao and S. B. Williams, "Large-scale path planning for underwater gliders in ocean currents," in *Australasian conference on robotics and automation (ACRA)*. Citeseer, 2009, pp. 2–4.
- [22] G. Mannarini and L. Carelli, "Visir-1. b: Ocean surface gravity waves and currents for energy-efficient navigation," *Geoscientific Model Development*, vol. 12, no. 8, pp. 3449–3480, 2019.
- [23] D. Kularatne *et al.*, "Optimal path planning in time-varying flows with forecasting uncertainties," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 4857–4864.
- [24] D. N. Subramani and P. F. Lermusiaux, "Risk-optimal path planning in stochastic dynamic environments," *Computer Methods in Applied Mechanics and Engineering*, vol. 353, pp. 391–415, 2019.
- [25] A. A. Pereira *et al.*, "Risk-aware path planning for autonomous underwater vehicles using predictive ocean models," *Journal of Field Robotics*, vol. 30, no. 5, pp. 741–762, 2013.
- [26] G. A. Hollinger *et al.*, "Learning uncertainty in ocean current predictions for safe and reliable navigation of underwater vehicles," *Journal of Field Robotics*, vol. 33, no. 1, pp. 47–66, 2016.
- [27] G. Knizhnik *et al.*, "Flow-based control of marine robots in gyre-like environments," *arXiv preprint arXiv:2203.00796*, 2022.
- [28] M. Buzzicotti *et al.*, "Optimal control of point-to-point navigation in turbulent time dependent flows using reinforcement learning," in *International Conference of the Italian Association for Artificial Intelligence*. Springer, 2020, pp. 223–234.
- [29] P. Gunnarson *et al.*, "Learning efficient navigation in vortical flow fields," *Nature communications*, vol. 12, no. 1, pp. 1–7, 2021.
- [30] R. Chowdhury and D. N. Subramani, "Physics-driven machine learning for time-optimal path planning in stochastic dynamic flows," in *International Conference on Dynamic Data Driven Application Systems*. Springer, 2020, pp. 293–301.
- [31] W. Sun *et al.*, "Pursuit-evasion games in dynamic flow fields via reachability set analysis," in *2017 American Control Conference (ACC)*. IEEE, 2017, pp. 4595–4600.
- [32] R. Rout *et al.*, "Modified line-of-sight guidance law with adaptive neural network control of underactuated marine vehicles with state and input constraints," *IEEE transactions on control systems technology*, vol. 28, no. 5, pp. 1902–1914, 2020.
- [33] W. H. Al-Sabban *et al.*, "Wind-energy based path planning for unmanned aerial vehicles using markov decision processes," in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 784–789.
- [34] A. Chakrabarty and J. Langelaan, "Uav flight path planning in time varying complex wind-fields," in *2013 American control conference*. IEEE, 2013, pp. 2568–2574.
- [35] M. Otte *et al.*, "Any-time path-planning: Time-varying wind field+ moving obstacles," in *2016 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2016, pp. 2575–2582.
- [36] M. G. Bellemare *et al.*, "Autonomous navigation of stratospheric balloons using reinforcement learning," *Nature*, vol. 588, no. 7836, pp. 77–82, 2020.
- [37] D. N. Subramani *et al.*, "Stochastic time-optimal path-planning in uncertain, strong, and dynamic flows," *Computer Methods in Applied Mechanics and Engineering*, vol. 333, pp. 218–237, 2018.
- [38] P. F. J. Lermusiaux *et al.*, "Progress and prospects of U.S. data assimilation in ocean research," *Oceanography*, vol. 19, no. 1, pp. 172–183, 2006.
- [39] P. Lermusiaux *et al.*, "Real-time probabilistic coupled ocean physics-acoustics forecasting and data assimilation for underwater GPS," in *OCEANS 2020 IEEE/MTS*. IEEE, Oct. 2020, pp. 1–9.
- [40] S. V. Raković and W. S. Levine, *Handbook of model predictive control*. Springer, 2018.
- [41] J. F. Fisac *et al.*, "Reach-avoid problems with time-varying dynamics, targets and constraints," in *Proceedings of the 18th international conference on hybrid systems: computation and control*, 2015, pp. 11–20.
- [42] M. Doshi *et al.*, "Hamilton-jacobi multi-time reachability," in *2022 IEEE 61th Annual Conference on Decision and Control (CDC)*. IEEE, 2022.
- [43] S. Bansal *et al.*, "Hamilton-jacobi reachability: A brief overview and recent advances," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE, 2017, pp. 2242–2253.
- [44] D. Crosby *et al.*, "A proposed definition for vector correlation in geophysics: Theory and application," *Journal of Atmospheric and Oceanic Technology*, vol. 10, no. 3, pp. 355–367, 1993.
- [45] S. L. Herbert *et al.*, "Fastrack: A modular framework for fast and guaranteed safe motion planning," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE, 2017, pp. 1517–1522.
- [46] K. Leung *et al.*, "On infusing reachability-based safety assurance within probabilistic planning frameworks for human-robot vehicle interactions," in *Int. Symp. on Experimental Robotics*, 2018.
- [47] P. Rivera-Ortiz and Y. Diaz-Mercado, "On guaranteed capture in multi-player reach-avoid games via coverage control," *IEEE Control Systems Letters*, 2018.
- [48] J. F. Fisac *et al.*, "A general safety framework for learning-based control in uncertain robotic systems," *IEEE Transactions on Automatic Control*, vol. 64, no. 7, pp. 2737–2752, 2018.
- [49] Y. S. Shao *et al.*, "Reachability-based trajectory safeguard (rts): A safe and fast reinforcement learning safety layer for continuous control," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3663–3670, 2021.