FGDWI: An Online Feasibility-Gated Dynamic Waypoint Insertion Policy for Value-Aware ASV Navigation in Forecast Flows

Guy Ludford*‡, Pablo Borja*, Nieves Garcia†, Jake Lewis‡, Dena Bazazian*
*School of Engineering, Computing and Mathematics (SECaM),
Faculty of Science and Engineering, University of Plymouth,
Plymouth, UK
†School of Biological and Marine Sciences (SoBMS),
Faculty of Science and Engineering, University of Plymouth,
Plymouth, UK
‡Oshen Ltd, Plymouth, UK
{guy.ludford,jake.lewis}@oshendata.com
{pablo.borja,nieves.garcia,dena.bazazian}@plymouth.ac.uk

***Abstract*—Autonomous surface vehicles (ASVs) must reliably reach mission objectives under time-varying currents and winds while, in many applications, prioritising where to sample to maximise scientific return. Observation-value heuristics that ignore feasibility may select unreachable targets, while feasibility-only planners may achieve navigation goals but deliver limited observational impact. We propose *Feasibility-Gated Dynamic Waypoint Insertion* (FGDWI), a lightweight decision loop that combines feasibility-gated target selection with a receding-horizon waypoint-insertion policy that samples candidate waypoints, filters them using fast forecast-rollout feasibility checks, and ranks feasible options with a value-per-cost score. We evaluate robustness under parametric forecast error using Monte Carlo trials across multiple flow scenarios and blob-map value fields. Results in synthetic jet–eddy environments show that FGDWI frequently improves value-rate when high-value structure is reachable via feasible detours, but typically with a reliability penalty relative to direct guidance. The method therefore operates as a tunable reliability–value trade-off policy rather than a universal replacement for direct-to-goal, and degrades gradually when terminal reachability becomes marginal.**

***Index Terms*—Autonomous surface vehicles, flow-aware planning, waypoint insertion, targeted sampling, feasibility gating, receding-horizon control**

I. Introduction

Autonomous surface vehicles (ASVs) offer persistent, low-logistics ocean observation capabilities and are increasingly considered a complementary component of global observing systems [1], [2]. In energetic ocean environments (regions with strong, spatially structured currents and variable winds where advective flow speeds can be comparable to (or exceed) vehicle control authority), navigation performance depends critically on the ability to exploit forecast forcing rather than continuously oppose it [3], [4].

Figure 1 shows the C-Star ASV platform that motivates our interest in lightweight, flow-aware mission planning.

This paper introduces *Feasibility-Gated Dynamic Waypoint Insertion* (FGDWI). At each step the policy certifies feasibility using short forecast rollouts, avoiding full reachable-set or Hamilton–Jacobi computation, and ranks candidates by value-per-cost evaluated at the *forecast arrival time*. When terminal reachability becomes marginal, it falls back to direct-to-goal.

A. Related Work and Problem Positioning

The problem we consider sits at the intersection of two literatures: flow-aware routing, which targets reliable or time-efficient arrival in time-varying currents [3], [5]–[7]; and targeted observation and informative path planning, which selects observations to improve forecast skill or information gain [8]–[13]. We aim to select high-value sampling opportunities while explicitly preserving terminal feasibility under forecast flows, without resorting to expensive global optimisation or full reachable-set computation.

Flow-aware routing and reachability. Time-optimal navigation in flows traces back to Zermelo's navigation problem [14]. Modern ocean-forecast-aware routing includes time-optimal planning in dynamic flows [5], [6], risk-aware planning with predictive ocean models [7], and "hitchhiking" formulations for underactuated agents that replan to manage imperfect forecasts [3]. Related approaches compute reachability or arrival-time fields using Hamilton–Jacobi formulations [15] or fast marching methods [16], [17].

Targeted sampling and informative path planning. Targeted observations have a substantial history [8], including ensemble-transform and sensitivity-based approaches [9]. In robotics, informative path planning formalises observation-driven trajectory selection using tractable proxies for information gain [10], [11]. Ocean observing network evaluation and the synthesis of observations with data assimilation remain active areas [12], [13], [18].

Fig. 1. C-Star autonomous surface vehicle (ASV) platform context for the flow-aware waypointing problem considered in this paper.

B. Contributions

- We introduce FGDWI, an online receding-horizon policy that gates candidate waypoints by forecast-rollout feasibility and ranks the feasible set by value-per-cost evaluated at the forecast arrival time, avoiding global optimisation and reachable-set computation.
- We evaluate FGDWI over a 47,520-mission sweep (11 value layouts $\times$ 6 flow scenarios $\times$ 6 forecast-noise levels $\times$ 30 seeds $\times$ 4 methods), reporting mission-weighted aggregates and case-level deltas.
- We characterise the regime in which FGDWI is most beneficial (off-corridor high-value structure with feasible detours) and quantify the reliability–value trade-off against direct-to-goal.

For example, an ASV may need to opportunistically sample a transient front or high-gradient region while still meeting a fixed recovery window; FGDWI explicitly trades observation value against deadline-feasible reachability under forecast flows.

The remainder of the paper is organised as follows. Section II defines the problem setup and modelling assumptions. Section III presents FGDWI and the baseline policies. Section IV describes the simulation design and evaluation metrics. Section V reports the empirical results, including aggregate trade-offs and regime dependence. Section VI discusses practical implications and limitations. Section VII concludes and outlines future directions.

II. Preliminaries and Problem Setup

We consider a single ASV operating in a 2D domain over a finite mission horizon $T_{\max}$. The vehicle starts at x_0 and must reach a goal region centred at x_g within radius r_g before the deadline. In addition to reaching the goal, the ASV collects high-impact observations along the way, represented by a lightweight value proxy.

A. Notation and Conventions

We represent positions as $x \in \mathbb{R}^2$ in a fixed Cartesian frame with x increasing eastward and y northward; axes are shown in km in plots while the simulator uses metres internally. Headings $\psi(t)$ are measured counter-clockwise from the positive x-axis (0° east, 90° north), and the low-level controller aims directly toward the current waypoint. Wind-driven propulsion is modelled using empirically derived velocity polars in true wind speed (TWS) and true wind angle (TWA). Planning uses forecast forcing $(\hat{u}, \hat{w})$; execution uses perturbed "true" forcing to model forecast error.

At each rollout/execution step, true wind speed is $\mathrm{TWS} = \|w(x,t)\|$. The wind direction is $\angle w = \operatorname{atan2}(w_y, w_x)$ and the true wind angle is the smallest angle between the commanded heading and wind direction mapped to $[0°, 180°]$, where 0° denotes headwind and 180° tailwind; boat speed is then obtained by polar-table interpolation.

B. Vehicle and Flow Model

The ASV experiences ocean currents and winds that vary in space and time. We use a kinematic model in which the vehicle's ground velocity is the sum of the ambient current and a wind-driven through-water velocity determined by a tabulated polar:

$$\dot{x}(t) = u(x,t) + v_b(\psi(t), w(x,t)). \tag{1}$$

Here $u(x,t)$ is the ocean current velocity, $w(x,t)$ is the wind velocity, and $v_b(\cdot)$ is the through-water velocity obtained from the polar tables (optionally with a constant thruster term). Planning uses forecast forcing $(\hat{u}, \hat{w})$, while execution uses perturbed "true" forcing to represent forecast error.

C. Observation Value Proxy

We model observation impact with "uncertainty blobs": spatially localised regions of high scientific value (e.g., displaced fronts, eddies, or thermal anomalies) where in-situ observations are most informative for downstream assimilation or mission objectives. The parameters (A_i, c_i, σ_i) would be supplied by an upstream information-gain layer in deployment; here they are set per-scenario, with amplitudes A_i reduced when sampled to reflect diminishing returns. We use a Gaussian-mixture form for analytic tractability:

$$V(x,t) = \sum_{i=1}^{M} A_i(t) \exp\left(-\frac{\|x - c_i\|^2}{2\sigma_i^2}\right). \tag{2}$$

In our simulations, value is accumulated either (i) continuously by integrating $V(x(t), t)$ along the executed trajectory (integral mode), or (ii) discretely at fixed sampling cadence Δt_s (sample mode). After each sample, blob amplitudes whose centres lie within radius r_s of the vehicle are reduced multiplicatively: $A_i \leftarrow (1-\rho)A_i$, modelling diminishing returns from repeat sampling.

D. Success and Metrics (Overview)

A run is successful if the vehicle enters the goal radius before $T_{\max}$. For failed runs we report closest approach distance and the final distance to goal to quantify graceful degradation. Since effort can be identically zero when thrusters are disabled, we use value normalised by mission time/cost as the primary normalised performance metric; see Section V.

III. Method: FGDWI

FGDWI is an online receding-horizon waypoint-selection policy that maximises observation value subject to fast forecast-feasibility checks. At each replanning step, candidate waypoints are sampled, evaluated with short forecast rollouts, feasibility-gated against the remaining mission budget, and scored by value-per-cost. The policy therefore selects admissible detours that preserve predicted terminal reachability while trading observation value against travel time and effort.

A. Candidate Waypoint Generation

At replanning time t, we sample N candidates in a forward corridor aligned with the goal direction (forward range $d \in [d_{\min}, d_{\max}]$, lateral offset $\ell \in [-w, w]$). In the reported experiments, candidate sets are augmented with blob-centre points (uncertainty-value peaks), then clipped to domain bounds and deduplicated.

B. Forecast Rollouts and Feasibility Gate

For each candidate waypoint w, two forecast rollouts are computed: current position $\rightarrow w$ and $w \rightarrow$ goal region. A candidate is feasible if

$$\hat{T}_{\text{to-wp}} + \hat{T}_{\text{wp-to-goal}} \leq T_{\max} - t. \tag{3}$$

At each replan t_k, the policy generates N candidates (Sec. III-A), runs forecast rollouts under nominal forcing $(\hat{u}, \hat{w})$ and dynamics (1), applies the feasibility gate (3), and selects $w^\star$ by Sec. III-C; if no candidate is feasible, the planner falls back to direct-to-goal heading.

C. Value-Per-Cost Selection

At each replanning step, FGDWI selects $w^\star = \arg\max_{w \in \mathcal{W}_f(t)} S(w)$, where $\mathcal{W}_f(t)$ is the subset of candidates passing the gate of Section III-B and the score is

$$S(w) = \frac{V(w, \hat{t}_{\text{arr}})}{\alpha\, \hat{T}_{\text{to-wp}} + \beta\, \hat{E}_{\text{to-wp}} + \varepsilon}. \tag{4}$$

Here, $V(w, \hat{t}_{\text{arr}})$ is the value at the *predicted arrival time* to waypoint w from the forecast rollout, $\hat{T}_{\text{to-wp}}$ is forecast ETA, and $\hat{E}_{\text{to-wp}}$ denotes integrated propulsive work from the polar over the forecast rollout; with thrusters disabled in our experiments, $\hat{E}_{\text{to-wp}} \equiv 0$ and the score reduces to value-per-time. The constant ε (set to 10^{-9}) prevents division by zero. We use $\alpha = 3.0$ and $\beta = 0.2$. Multi-waypoint plans arise implicitly across the receding-horizon sequence rather than from a joint optimisation, decoupling per-step cost from horizon length.

For clarity, the value-only baseline differs: it selects $\arg\max_w V(w, t)$, i.e., value at the *current decision time* without rollout-time correction and without feasibility gating.

D. Unreachable Check and Close-In Mode

At a configured check time, a one-shot forecast reachability test to goal is performed. If deadline-feasible arrival is no longer predicted, FGDWI switches to a close-in mode that chooses the waypoint minimising predicted distance-to-goal over a short horizon. We report closest approach and terminal distance for graceful-degradation analysis.

E. Baselines

We compare against:

- Direct-to-goal: waypoint is always the terminal goal.
- Value-only: maximum value candidate, no feasibility check.
- Feasibility-only: minimum forecast ETA (to-waypoint + waypoint-to-goal), no value term.

IV. Simulation Study

We evaluate FGDWI and baselines over a broad grid of value-field layouts and flow conditions.

A. Evaluation Design

We evaluate robustness over a grid of environments comprising multiple value-field layouts, flow regimes, and forecast-error levels. Specifically, we evaluate $B = 11$ blob-map presets and $K = 6$ flow scenarios, with forecast current-noise sweep $\sigma_u \in \{0, 0.1, 0.2, 0.3, 0.4, 0.5\}$ and $S = 30$ Monte Carlo seeds per σ_u (shared across methods). This yields $B \times K \times |\sigma_u| \times S \times 4 = 47{,}520$ missions.

Fairness and reproducibility. All methods use identical vehicle dynamics, replanning cadence, success criteria, and random seeds. At each replanning step, candidate waypoint sets are drawn from the same distribution for all methods (and when blob-centre augmentation is enabled, the same augmented candidate set is provided to all methods). Differences in outcomes therefore arise only from the waypoint-selection logic.

B. Scenario Presets

We evaluate 11 value-field layouts and 6 flow regimes spanning aligned, off-axis, edge-concentrated, jet-dominated, eddy-dominated, and high-shear cases. Internal preset identifiers are used only to index the sweep; descriptive regime names are used in the discussion below. We additionally sweep forecast current-noise levels $\sigma_u \in \{0, 0.1, 0.2, 0.3, 0.4, 0.5\}$ with $S = 30$ Monte Carlo seeds per level, shared across methods.

C. Flow/Planner Settings

The reported evaluation uses lightweight planning parameters for tractable large sweeps: replanning every 6 h, execution step $\Delta t = 600$ s, rollout step $\Delta t_r = 3600$ s, and $N = 8$ sampled candidates per replan (plus blob-centre augmentation).

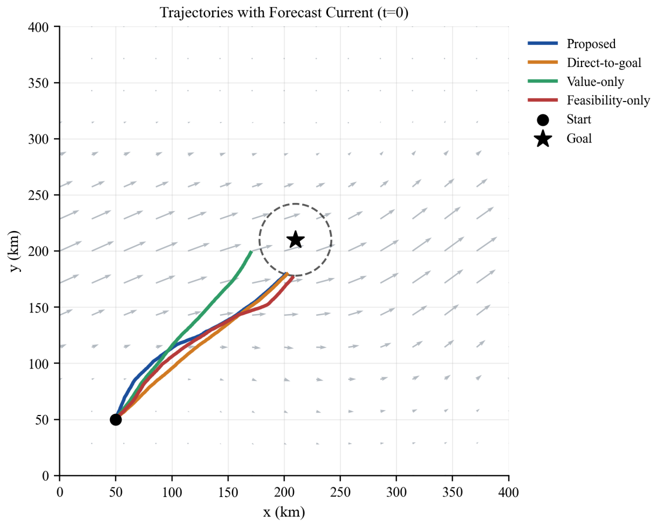

Fig. 2. Representative trajectories over forecast currents and value blobs for one off-axis value layout.

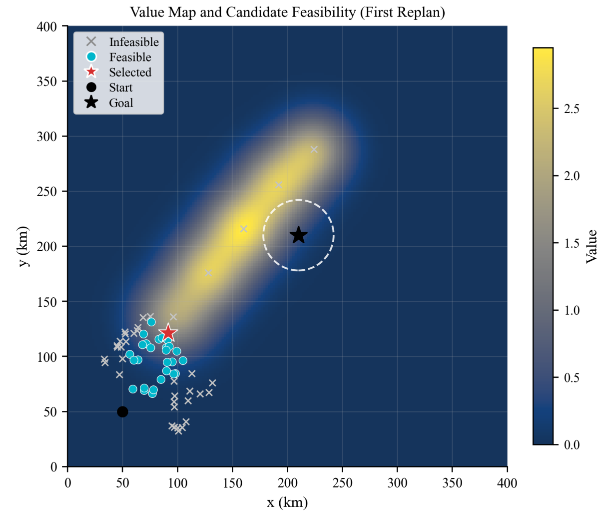

Fig. 3. Single-step feasibility-gated waypoint selection example. Feasible candidates form the admissible set $\mathcal{W}_f$, and $c^\star$ is the selected waypoint.

D. Metrics

We report success rate, ETA, value, value-per-time, value-per-cost, unreachable-check rate, and failure degradation (minimum/final distance to goal). Normalised metrics are:

$$\text{value\_per\_time} = \frac{\text{value}}{\min(t_{\text{arrival}}, T_{\max})/3600}, \tag{5}$$

$$\text{value\_per\_cost} = \frac{\text{value}}{\alpha \min(t_{\text{arrival}}, T_{\max}) + \beta\, \text{effort} + \varepsilon}. \tag{6}$$

Where interval estimates are reported, we use 95% normal-approximation confidence intervals.

V. Results

A. Representative Behaviour

Figure 2 shows representative trajectories in an off-axis value layout. The methods exhibit distinct route-selection behaviour: direct-to-goal remains corridor-focused, while value-seeking policies deviate to exploit high-value regions.

B. Feasibility-Gate Decision Snapshot

Figure 3 shows one replanning snapshot: infeasible candidates are pruned by the gate, then $c^\star$ is selected from the feasible set by value-per-cost ranking.

C. Aggregate Trade-Off Across All Missions

Table I summarises the aggregate trade-off across the full 47,520-mission sweep. Because FGDWI is designed to exploit reachable off-corridor value rather than maximise arrival reliability in every regime, the table should be interpreted as an operating-point summary rather than a global ranking of methods.

Value denotes the accumulated uncertainty-proxy signal along executed trajectories. Value/Time is in units per hour; Value/Cost uses the waypoint-score denominator, $\text{value}/(\alpha \min(t_{\text{arrival}}, T_{\max}) + \beta\, \text{effort} + \varepsilon)$. Because thruster effort is zero in this setup, Value/Time is the primary normalised metric.

TABLE I
TOP-LEVEL MISSION-WEIGHTED MEANS ACROSS THE FULL 47,520-MISSION SWEEP.

Method	Success (%)	Value/Time	Value/Cost
Direct-to-goal	78.9	1265.4	0.117
Feasibility-only	73.3	1493.2	0.138
FGDWI (proposed)	52.4	1366.9	0.127
Value-only	18.8	1993.2	0.185

D. Overall Comparison Against Direct-to-Goal

FGDWI targets a different operating point than direct-to-goal, trading reliability for higher value collection in regimes where reachable high-value structure lies away from the time-efficient corridor.

Mission-weighted (Table I). Aggregated over all missions, FGDWI increases value-per-time from 1265.4 to 1366.9 (approximately +8%) but reduces success from 78.9% to 52.4% (a 26.5 percentage-point reduction).

Case-averaged across blob-map$\times$scenario combinations. At the case level over $B \times K = 66$ blob-map$\times$scenario combinations, FGDWI beats direct-to-goal on value-per-time in $42/66$ cases and on both metrics jointly in $8/66$. The mean proposed–direct deltas are +40.1% in value-per-time and -26.5 percentage points in success (median: +28.5% and -15.0 percentage points).

E. Comparison Against Feasibility-Only

Feasibility-only provides a strong baseline in this sweep, achieving higher mission-weighted success and value-rate than FGDWI (Table I). This occurs because in many layouts the highest-value structure lies near time-efficient corridors, so a

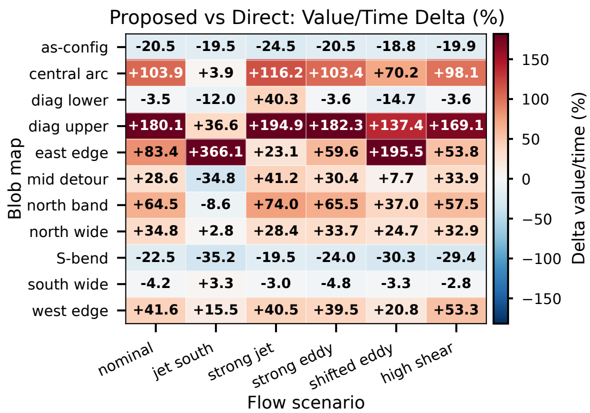

Fig. 4. Proposed vs direct value-per-time delta (%) across blob maps and flow scenarios. Positive values indicate regimes in which FGDWI improves value collection rate relative to direct-to-goal guidance.

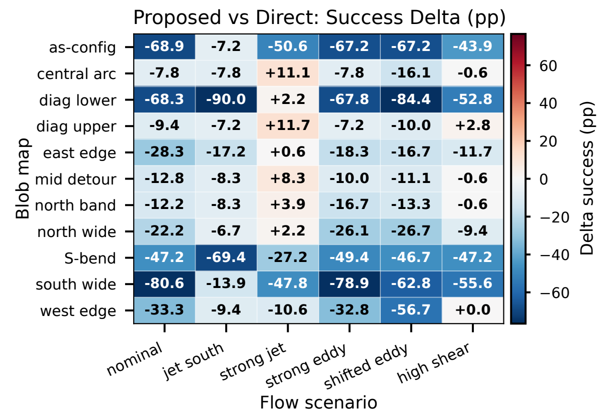

Fig. 5. Proposed vs direct success-rate delta (percentage points) across blob maps and flow scenarios. Negative values indicate the reliability cost associated with value-seeking detours.

policy that aggressively minimises forecast arrival time can also collect value incidentally.

At the case level, FGDWI beats feasibility-only on value-per-time in $34/66$ combinations and on success rate in only $1/66$. FGDWI value-rate wins are concentrated in off-axis but still reachable value layouts (e.g., diagonal/edge-style maps under shifted-eddy or high-shear flows), while success-rate wins are rare (occurring in only one south-wide layout under a south-shifted jet regime).

F. Blob-Map and Scenario Dependence

Performance is strongly dependent on the spatial layout of the value field, summarised across all blob-map×scenario combinations in Figs. 4 and 5. Under nominal flow, FGDWI outperforms direct-to-goal on value-per-time in layouts where valuable structure is offset but still reachable. This is evident in the north-band layout (+64.5%, 1933.4 vs 1175.2) and the mid-detour layout (+28.6%, 664.0 vs 516.3), albeit with lower success rates (0.706 vs 0.828 and 0.700 vs 0.828, respectively). By contrast, FGDWI underperforms in layouts where the direct corridor already captures substantial value or where detours are less favourable, such as the reference configured layout (-20.5%, 1902.9 vs 2394.6; success 0.139 vs 0.828) and the S-bend layout (-22.5%, 1534.6 vs 1980.5; success 0.356 vs 0.828).

Across the 11 value-field layouts, FGDWI beats direct-to-goal on value-per-time in 6–8 layouts per flow scenario, but usually with lower success. The main exception is the strong-jet regime, where FGDWI beats direct on both metrics in 7/11 layouts.

G. Regime Identification and Policy Selection

The sweep results suggest that method choice should be regime-dependent rather than universal. FGDWI is most attractive in missions where high-value structure lies away from the direct corridor, feasible detours still exist under the forecast, and some reduction in terminal reliability is acceptable. In these settings, the feasibility gate prevents clearly inadmissible diversions while the value-per-cost score converts remaining slack into increased value collection.

Direct-to-goal remains preferable for arrival-critical missions, especially when forecast reachability is already marginal and preserving terminal success dominates any potential value gain. Feasibility-only emerges as a strong intermediate policy when high-value structure lies near time-efficient corridors, since it can retain much of the attainable value without explicitly pursuing off-corridor detours.

These results therefore support a practical interpretation of FGDWI as a mission-dependent operating-point policy rather than a universal planner replacement. A natural next step is to derive a higher-level policy-selection rule, using forecast-derived features such as estimated reachability margin, corridor alignment of high-value structure, and predicted detour cost, to decide when direct guidance, feasibility-only guidance, or FGDWI is most appropriate.

VI. Discussion

Direct-to-goal remains the strongest reliability baseline overall, while FGDWI often converts feasible detours into higher value-per-time. FGDWI is therefore a tunable operating point on the reliability–value curve, valuable where some terminal risk is acceptable in exchange for observational gain.

Table II summarises the mission-dependent policy-selection guidance suggested by these results. In particular, for value-critical missions with tolerated risk, FGDWI can be preferable; for strict arrival-critical missions, direct guidance remains stronger in this calibration. The strong performance of feasibility-only further suggests that, in many environments, a substantial fraction of attainable value lies near time-efficient corridors, so explicit value-seeking is most justified only when value structure is sufficiently misaligned with those corridors.

We do not compare against full-horizon global-optimisation planners, since such methods assume access to the entire forecast trajectory and pay an offline cost incompatible with

TABLE II
PRACTICAL POLICY-SELECTION GUIDANCE SUGGESTED BY THE SWEEP RESULTS.

Mission / forecast regime	Preferred policy	Rationale
Arrival-critical, low reachability margin	Direct-to-goal	Maximises terminal reliability
Value aligned with efficient corridor	Feasibility-only	Retains value with better success
Off-corridor but reachable high-value structure, moderate risk tolerance	FGDWI	Converts feasible detours into higher value-rate

our online setting; they act as oracle upper bounds rather than directly comparable baselines.

The limitations of the current study should also shape how these results are interpreted. The flow fields are synthetic, observation value is represented by a blob-based proxy, and planner discretisation is deliberately lightweight to enable large Monte Carlo sweeps. Accordingly, the results are best viewed as a controlled characterisation of policy behaviour rather than as a final operational benchmark.

VII. Conclusion

We presented FGDWI, a lightweight online feasibility-gated waypoint insertion policy for value-aware ASV navigation in forecast flows. Across a broad Monte Carlo evaluation, the method was shown to improve value collection rate in regimes where high-value structure lies off the time-efficient corridor but remains reachable via feasible detours. However, it does not dominate simpler baselines universally, and its benefits typically come with a reliability cost relative to direct guidance. The contribution is therefore a practical policy that exposes and quantifies the reliability–value trade-off in flow-aware ASV missions, to be applied selectively by mission regime.

Several extensions follow naturally from these results. The operating point of FGDWI could be tuned more systematically through (α, β) selection, replanning cadence, and candidate-set design to target specific reliability/value requirements. The synthetic blob proxy could also be replaced with forecast- or assimilation-derived value functions so that pursued detours correspond more directly to operational scientific gain. Forecast uncertainty could be propagated through the feasibility gate, e.g., by admitting a candidate only if $\Pr[\hat{T}_{\text{to-wp}} + \hat{T}_{\text{wp-to-goal}} \leq T_{\max} - t] \geq 1 - \delta$ under an ensemble of forecasts. The regime dependence observed here suggests the possibility of a higher-level policy-selection layer, in which forecast-derived features such as goal reachability margin, corridor alignment of high-value structure, and predicted detour cost are used to decide whether direct guidance, feasibility-only guidance, or FGDWI is most appropriate for a given mission instance.

References

[1] R. G. Patterson, M. F. Cronin, S. Swart, J. Beja, J. M. Edholm, J. McKenna, J. B. Palter, A. Parker, C. I. Addey *et al.*, "Uncrewed surface vehicles in the global ocean observing system: a new frontier for observing and monitoring at the air-sea interface," *Frontiers in Marine Science*, vol. 12, p. 1523585, 2025. [Online]. Available: https://www.frontiersin.org/articles/10.3389/fmars.2025.1523585/full

[2] C. Whitt, J. Pearlman, B. Polagye *et al.*, "Future vision for autonomous ocean observations," *Frontiers in Marine Science*, vol. 7, p. 697, 2020. [Online]. Available: https://www.frontiersin.org/articles/10.3389/fmars.2020.00697/full

[3] M. Wiggert, M. Doshi, P. F. J. Lermusiaux, and C. J. Tomlin, "Navigating underactuated agents by hitchhiking forecast flows," in *Proceedings of the IEEE Conference on Decision and Control (CDC)*, 2022, pp. 2417–2424.

[4] Y. Wu, T. Wang, and S. Liu, "A review of path planning methods for marine autonomous surface vehicles," *Journal of Marine Science and Engineering*, vol. 12, no. 5, p. 833, 2024. [Online]. Available: https://www.mdpi.com/2077-1312/12/5/833

[5] T. Lolla, M. Ueckermann, T. Sapsis, and P. F. J. Lermusiaux, "Time-optimal path planning in dynamic flows," *Ocean Dynamics*, 2014. [Online]. Available: https://mseas.mit.edu/publications/PDF/Lolla_etal_path_plan_LSM_theory_scheme_OD2014.pdf

[6] C. S. Kulkarni and P. F. J. Lermusiaux, "Three-dimensional time-optimal path planning in the ocean," *Ocean Modelling*, vol. 152, p. 101644, 2020.

[7] A. A. Pereira, J. Binney, G. A. Hollinger, and G. S. Sukhatme, "Risk-aware path planning for autonomous underwater vehicles using predictive ocean models," *Journal of Field Robotics*, vol. 30, no. 5, pp. 741–762, 2013.

[8] S. J. Majumdar, "A review of targeted observations," *Bulletin of the American Meteorological Society*, vol. 97, no. 12, pp. 2287–2303, 2016.

[9] C. H. Bishop, B. J. Etherton, and S. J. Majumdar, "Adaptive sampling with the ensemble transform kalman filter," *Monthly Weather Review*, vol. 129, no. 3, 2001. [Online]. Available: https://journals.ametsoc.org/view/journals/mwre/129/3/1520-0493_2001_129_0420_aswtet_2.0.co_2.xml

[10] A. Singh, A. Krause, C. Guestrin, W. Kaiser, and M. Batalin, "Efficient informative paths for multiple robots," in *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2007. [Online]. Available: https://www.ijcai.org/Proceedings/07/Papers/355.pdf

[11] J. Binney, A. Krause, and G. Sukhatme, "Informative path planning for an autonomous underwater vehicle," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2010. [Online]. Available: https://fileadmin.cs.lth.se/ai/Proceedings/ICRA2010/MainConference/data/papers/1895.pdf

[12] Y. Fujii, E. Rémy, H. Zuo, P. Oke, G. Halliwell, F. Gasparin, M. Benkiran, N. Loose *et al.*, "Observing system evaluation based on ocean data assimilation and prediction systems: On-going challenges and a future vision for designing and supporting ocean observational networks," *Frontiers in Marine Science*, vol. 6, p. 417, 2019. [Online]. Available: https://www.frontiersin.org/articles/10.3389/fmars.2019.00417/full

[13] A. M. Moore, H. G. Arango *et al.*, "Synthesis of ocean observations with data assimilation," *Frontiers in Marine Science*, 2019. [Online]. Available: https://www.frontiersin.org/articles/10.3389/fmars.2019.00090/full

[14] E. Zermelo, "über das navigationsproblem bei ruhender oder veränderlicher windverteilung," *ZAMM - Zeitschrift für Angewandte Mathematik und Mechanik*, vol. 11, no. 2, pp. 114–124, 1931.

[15] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin, "A time-dependent hamilton–jacobi formulation of reachable sets for continuous dynamic games," *IEEE Transactions on Automatic Control*, vol. 50, no. 7, pp. 947–957, 2005.

[16] J. A. Sethian, "A fast marching level set method for monotonically advancing fronts," *Proceedings of the National Academy of Sciences*, vol. 93, no. 4, pp. 1591–1595, 1996.

[17] C. Petres *et al.*, "Path planning for auvs in strong currents using fast marching methods," in *IEEE OCEANS*, 2005. [Online]. Available: https://www.eece.hw.ac.uk/~ceeyrp/WWW/Research/Ocean05_Petres.pdf

[18] A. Carrassi, M. Bocquet, L. Bertino, and G. Evensen, "Data assimilation in the geosciences: An overview of methods, issues, and perspectives," *WIREs Climate Change*, vol. 9, no. 5, 2018.