Feasibility Frontier for Ergodic Exploration

Hunter Eggleston Christian Hughes

Motivation

Understanding underwater environments requires exploring them.

Spending equal time in all areas of the ocean is wasteful because not all areas have equal value.

Previous work has shown that ergodic exploration guarantees coverage proportional to the assigned belief of the value in an environment, assuming the agent can physically move fast enough to reach the value [1].

Is there a feasibility frontier – a minimum agent speed – below which this guarantee breaks down?

Approach

To track a moving value belief distribution a robot (agent) needs to match its velocity. we show that robots can leverage the advective flow of their environment to maintain this velocity.

$$(u_{\max} + \|A(x_t,t)\|)\left\|\nabla_{x_t}\phi(x_t)\right\|_{\text{op}} > v_{\max}$$

The power saved can then be used for other applications such as additional sensors for collecting information.

---

u_{max} = motor speed $\|A(x_t,t)\|$ = current speed

$u_{max} + \|A(x_t,t)\|$ = total speed $\left\|\nabla_{x_t}\phi(x_t)\right\|_{op}$ = efficiency

v_{max} = target speed

Ergodic Exploration

Goal: Time spent in a region is directly proportional to the value of the region

We use a flow adaptive version of a statistical tool called ergodic maximum mean discrepancy which shows the dissimilarity of two probability distributions [2].

$$\text{MMD}^2_\zeta = \left\|\hat{\mu}_{\rho,t,\zeta} - \hat{\mu}_{\mathcal{I},t,\zeta}\right\|^2_{\mathcal{H}}$$

This allows the robot to compare its path with its target and spend more time in areas with greater similarity.

~~MMD^2_ζ = match error~~ ~~$\hat{\mu}_{\rho,t,\zeta}$ = robot coverage~~

$\hat{\mu}_{\mathcal{I},t,\zeta}$ = belief coverage ζ = kernel scale

t = time $\|\cdot\|^2_{\mathcal{H}}$ = difference size

Feasibility Frontier

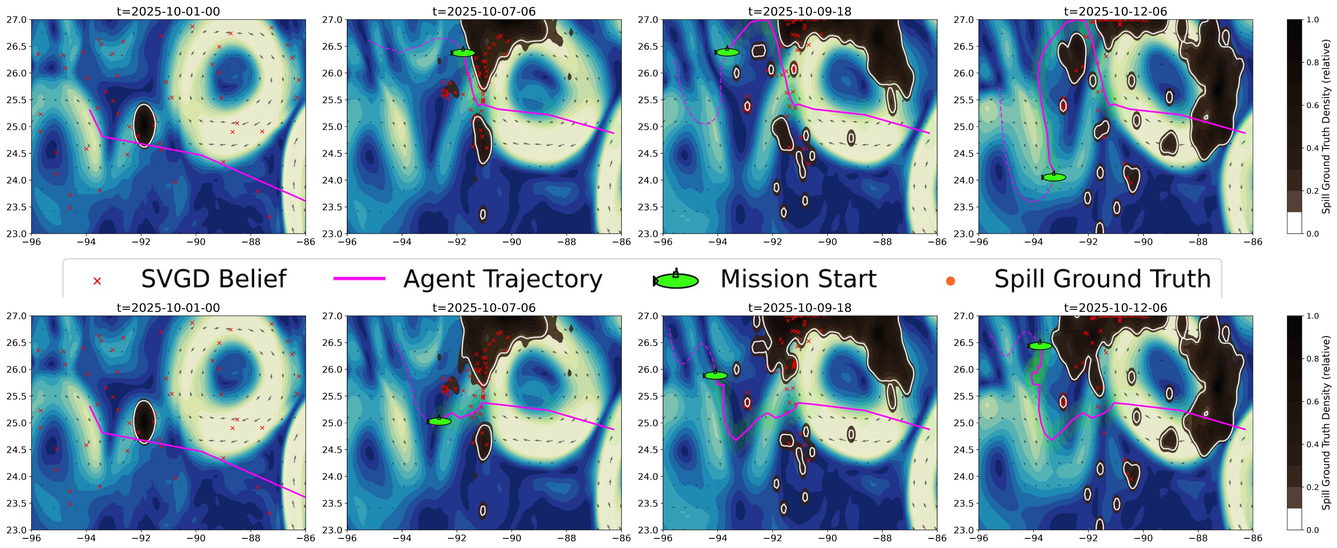


Fig. 1. Ergodic oil-spill cleanup over real ocean flows. The agent's trajectory (magenta) converges on the true spill (right color bar) and removes it. The agent uses the current to increase its total speed and achieve greater coverage. When the agent has increased actuation (motor speed) it can complete more of the mission — **Top** Agent with increased actuation **Bottom** Agent with decreased actuation —

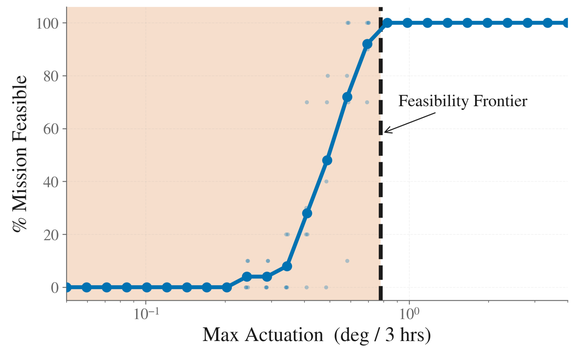
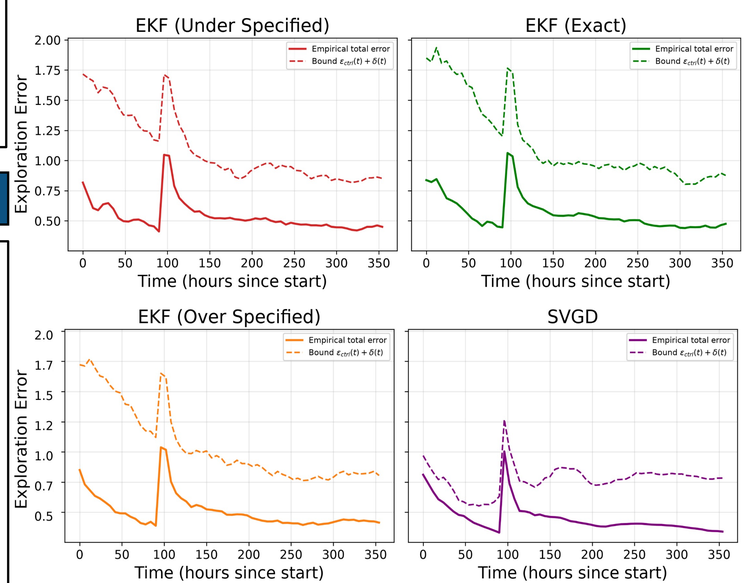


Fig. 2. Feasibility Frontier Line Plot. The agent can achieve more of its mission when it has increased speed. When the agent falls below 0.8 deg / 3 hrs, the agent is still able to achieve partial coverage but will be unable to complete the entire mission.

Mission feasibility can be predicted by speed alone!

Belief Model



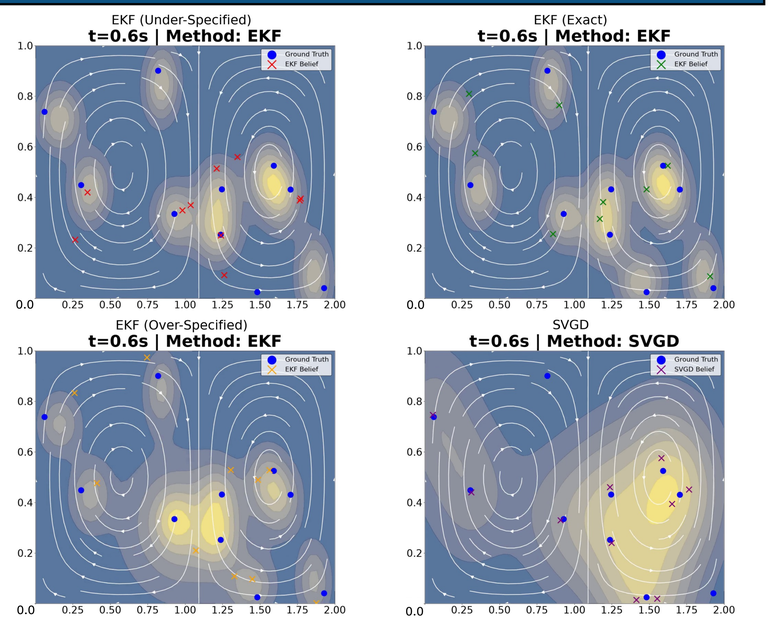


Fig. 3. Ergodic coverage using arbitrary belief models on the double gyre flow field. Stein Variational Gradient Decent (SVGD) achieves both a lower empirical error and a tighter upper bound when compared to Extended Kalman Filters (EKF). Kalman filters require an assumption about the number of targets being tracked. SVGD does not require a fixed assumption about the targets being tracked making it better suited to problems with unknown or varying number of targets

Ergodic Exploration can guarantee coverage regardless of the belief model used.

Acknowledgements

This work was supported by the REVU program at Yale University. Special thanks to my research mentor, Christian Hughes, for insightful guidance, and to Program Director Dr. Marla Geha and Program Manager Dr. Jeremy Bradford for their leadership and support.

References

[1] Hughes, C., & Abraham, I. (2026). Infinite-horizon ergodic control via kernel mean embeddings. arXiv preprint arXiv:2604.01023.

[2] Hughes, C., Liu, Y., Lahrach, Y., Engdahl, J., Warren, H., Lee, D., Ramos, F., Miles, T., & Abraham, I. (2026). Asymptotically optimal ergodic coverage on generalized motion fields. arXiv preprint arXiv:2605.13442.