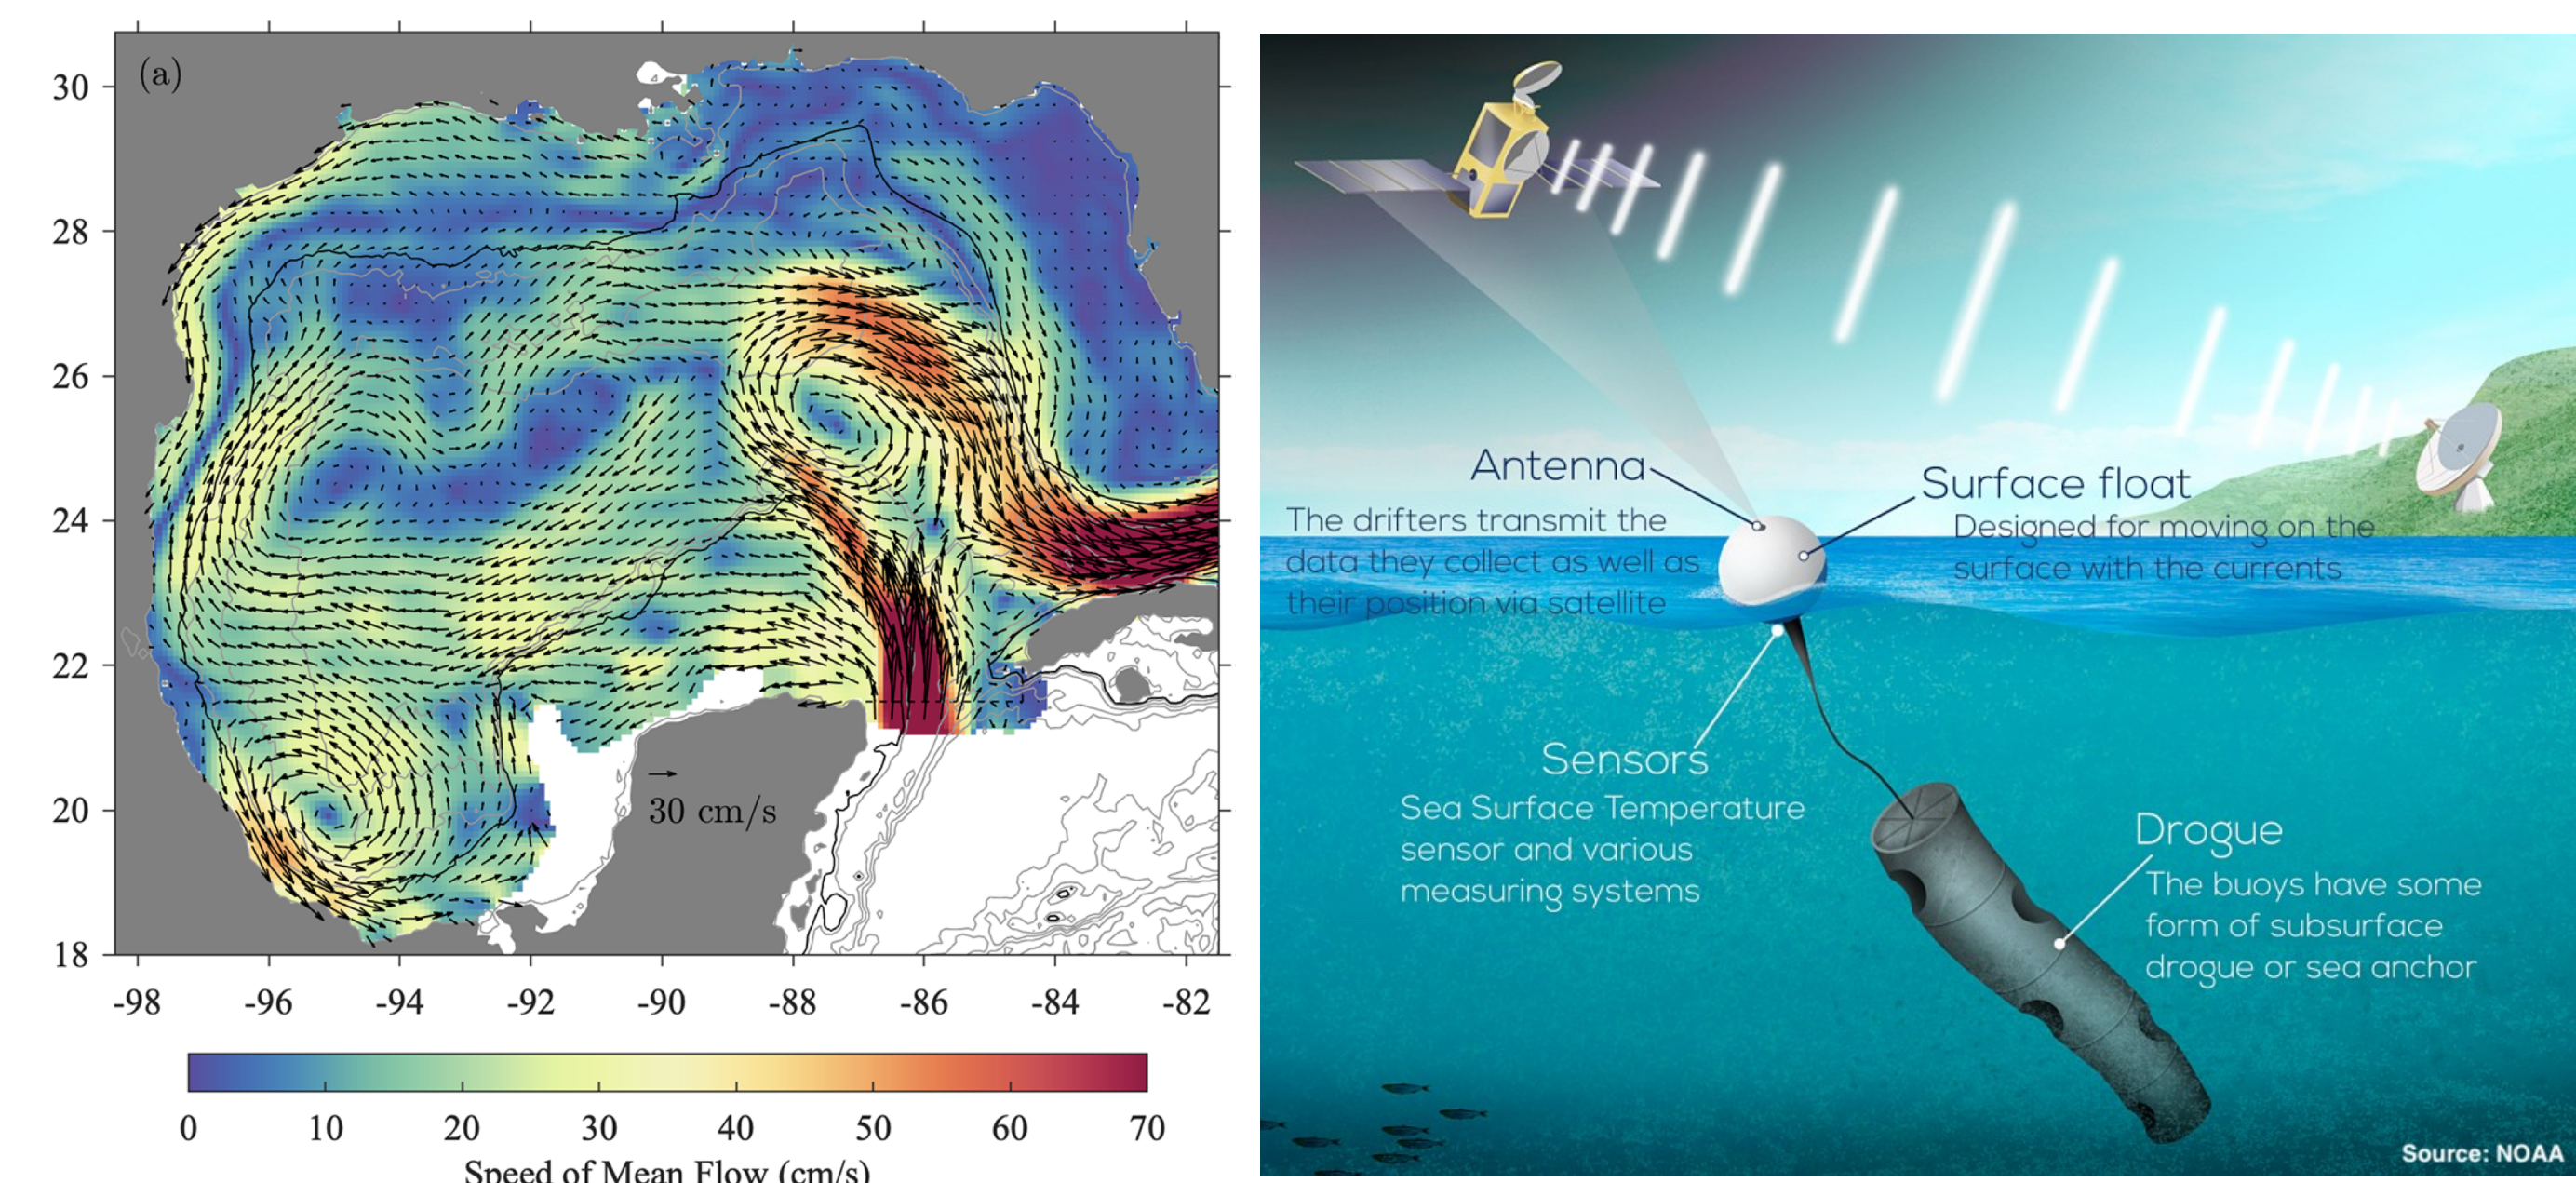


Ocean Currents and Drifters

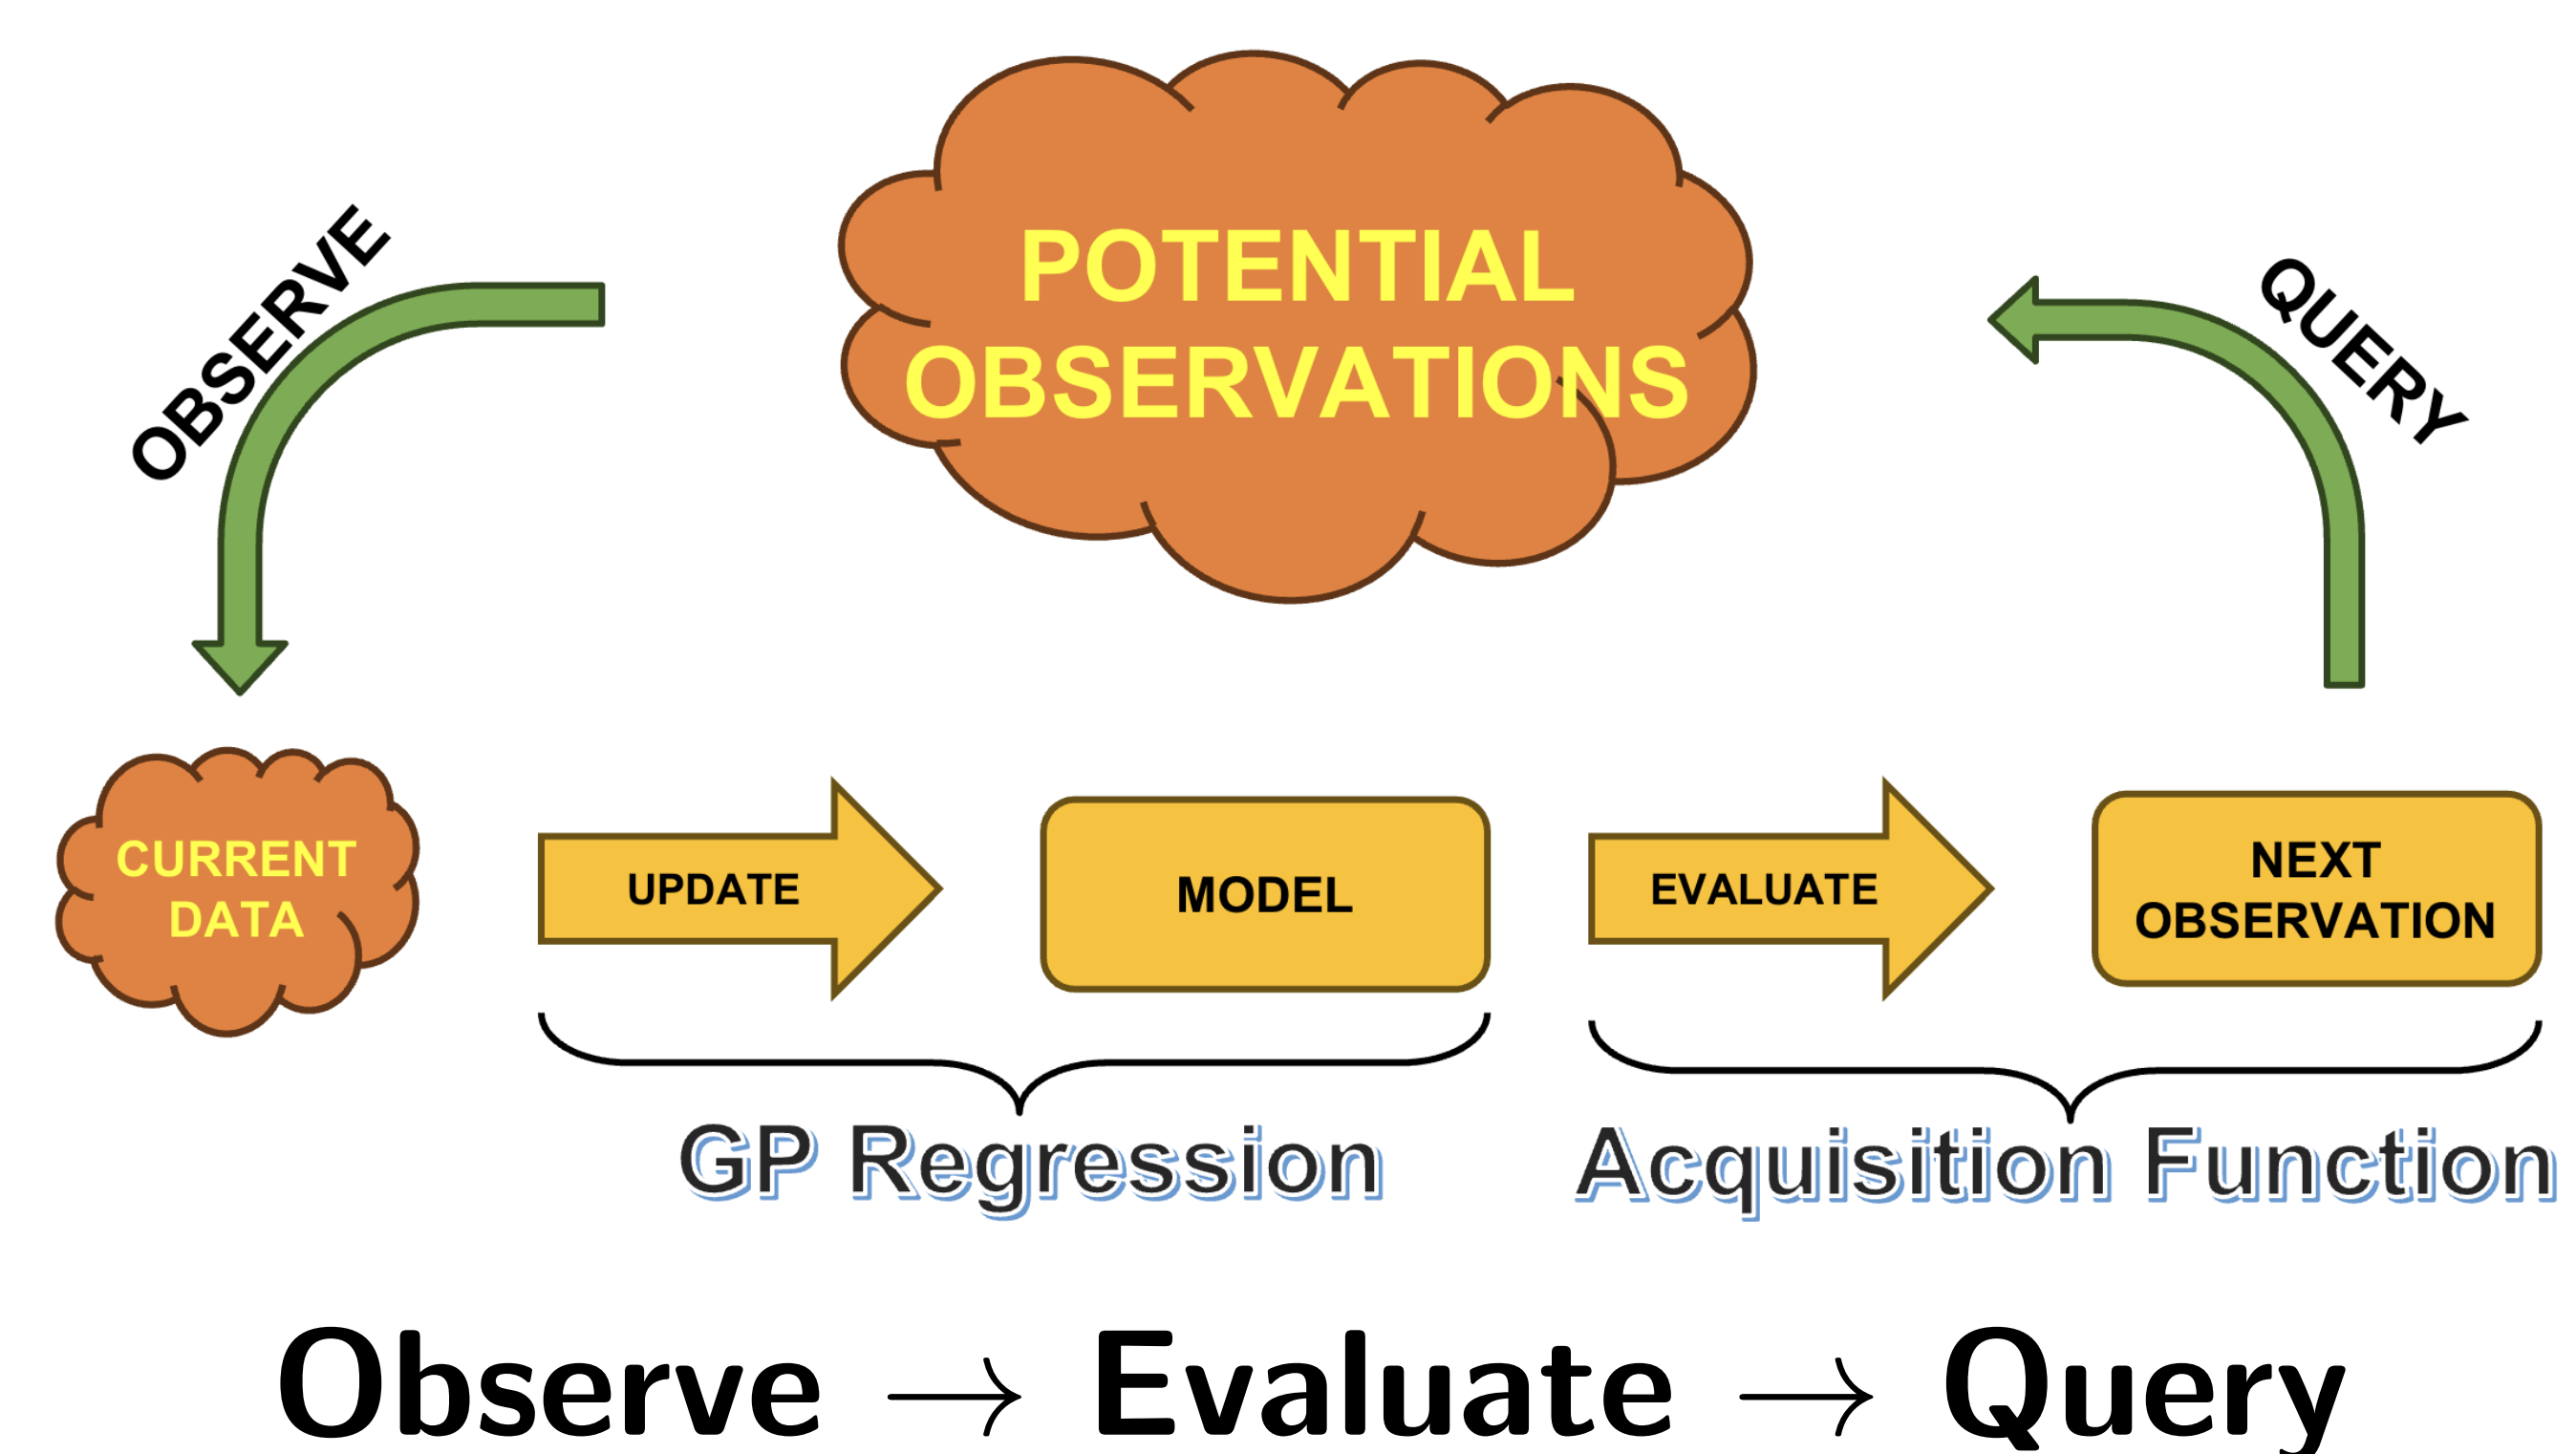


Ocean drifters measure currents while being advected by the flow, and are increasingly deployed due to their low costs.

Question

How to intelligently place drifters, using a fixed budget, to maximise knowledge about the vector field?

Active Learning



What Makes This Problem Hard?

Time-Varying Vector Field

How to emulate ocean currents?

Temporal Helmholtz GP

System-Observer Coupling

How to evaluate a deployment?

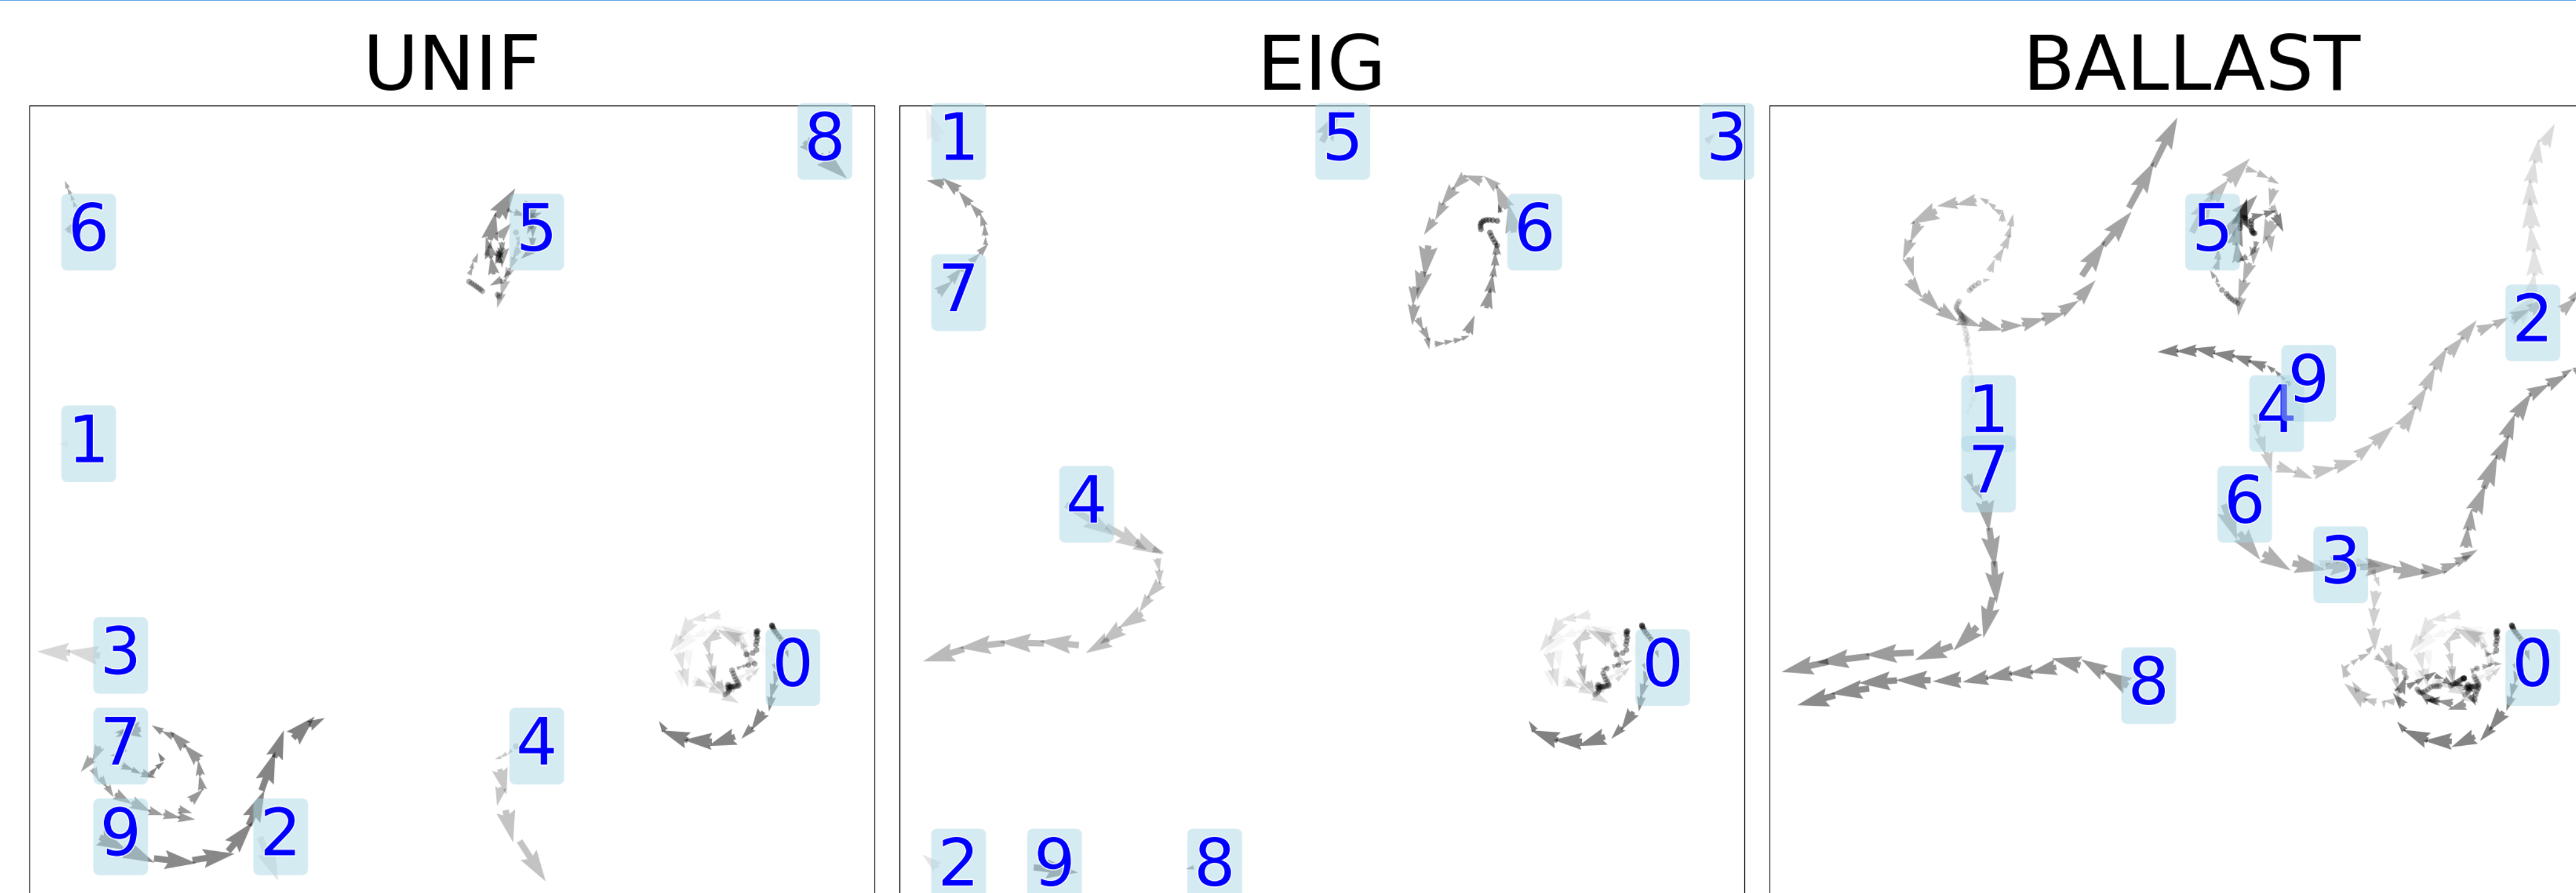
Trajectory look-aheads

Computational Overhead

How to sample posteriors efficiently?

VaSE

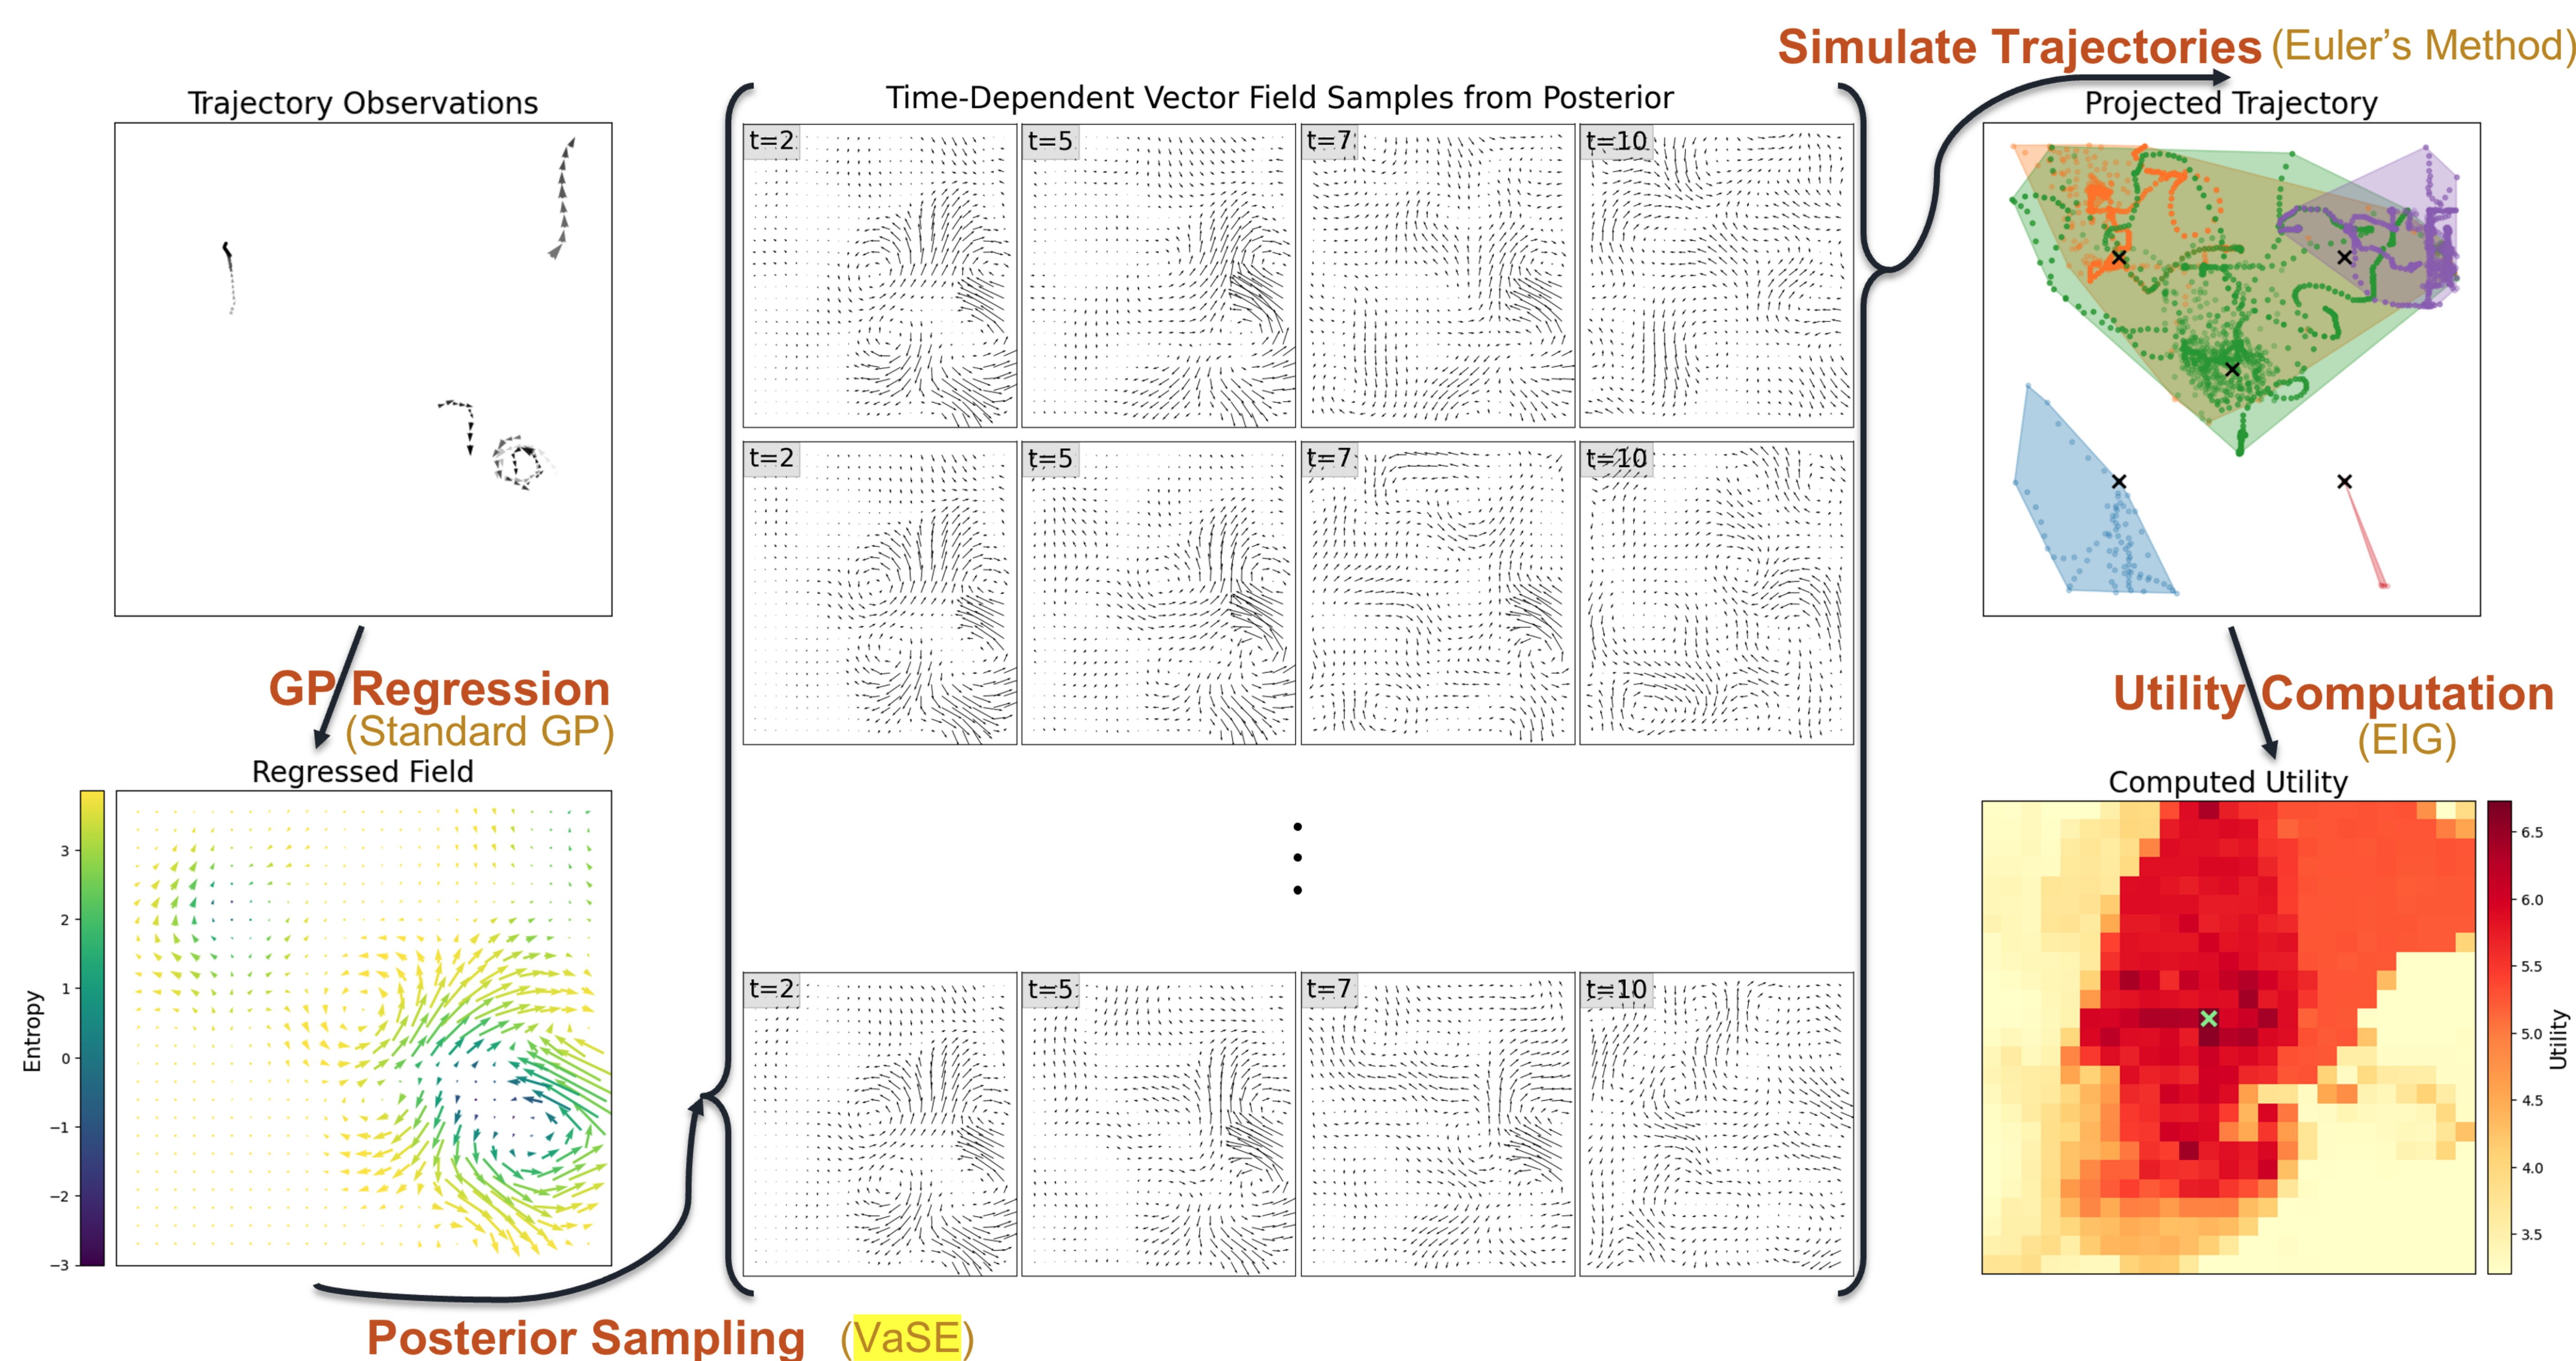
Pitfalls of Standard Active Learning



Standard policy scores only the **initial placement** but a drifter produces a **trajectory** of observations. Failure to account for subsequent observations can lead to sub-optimal placement selections.


How BALLAST Works

Sample Fields → Simulate Trajectories → Aggregate Utilities



Ablation: How Many Samples?

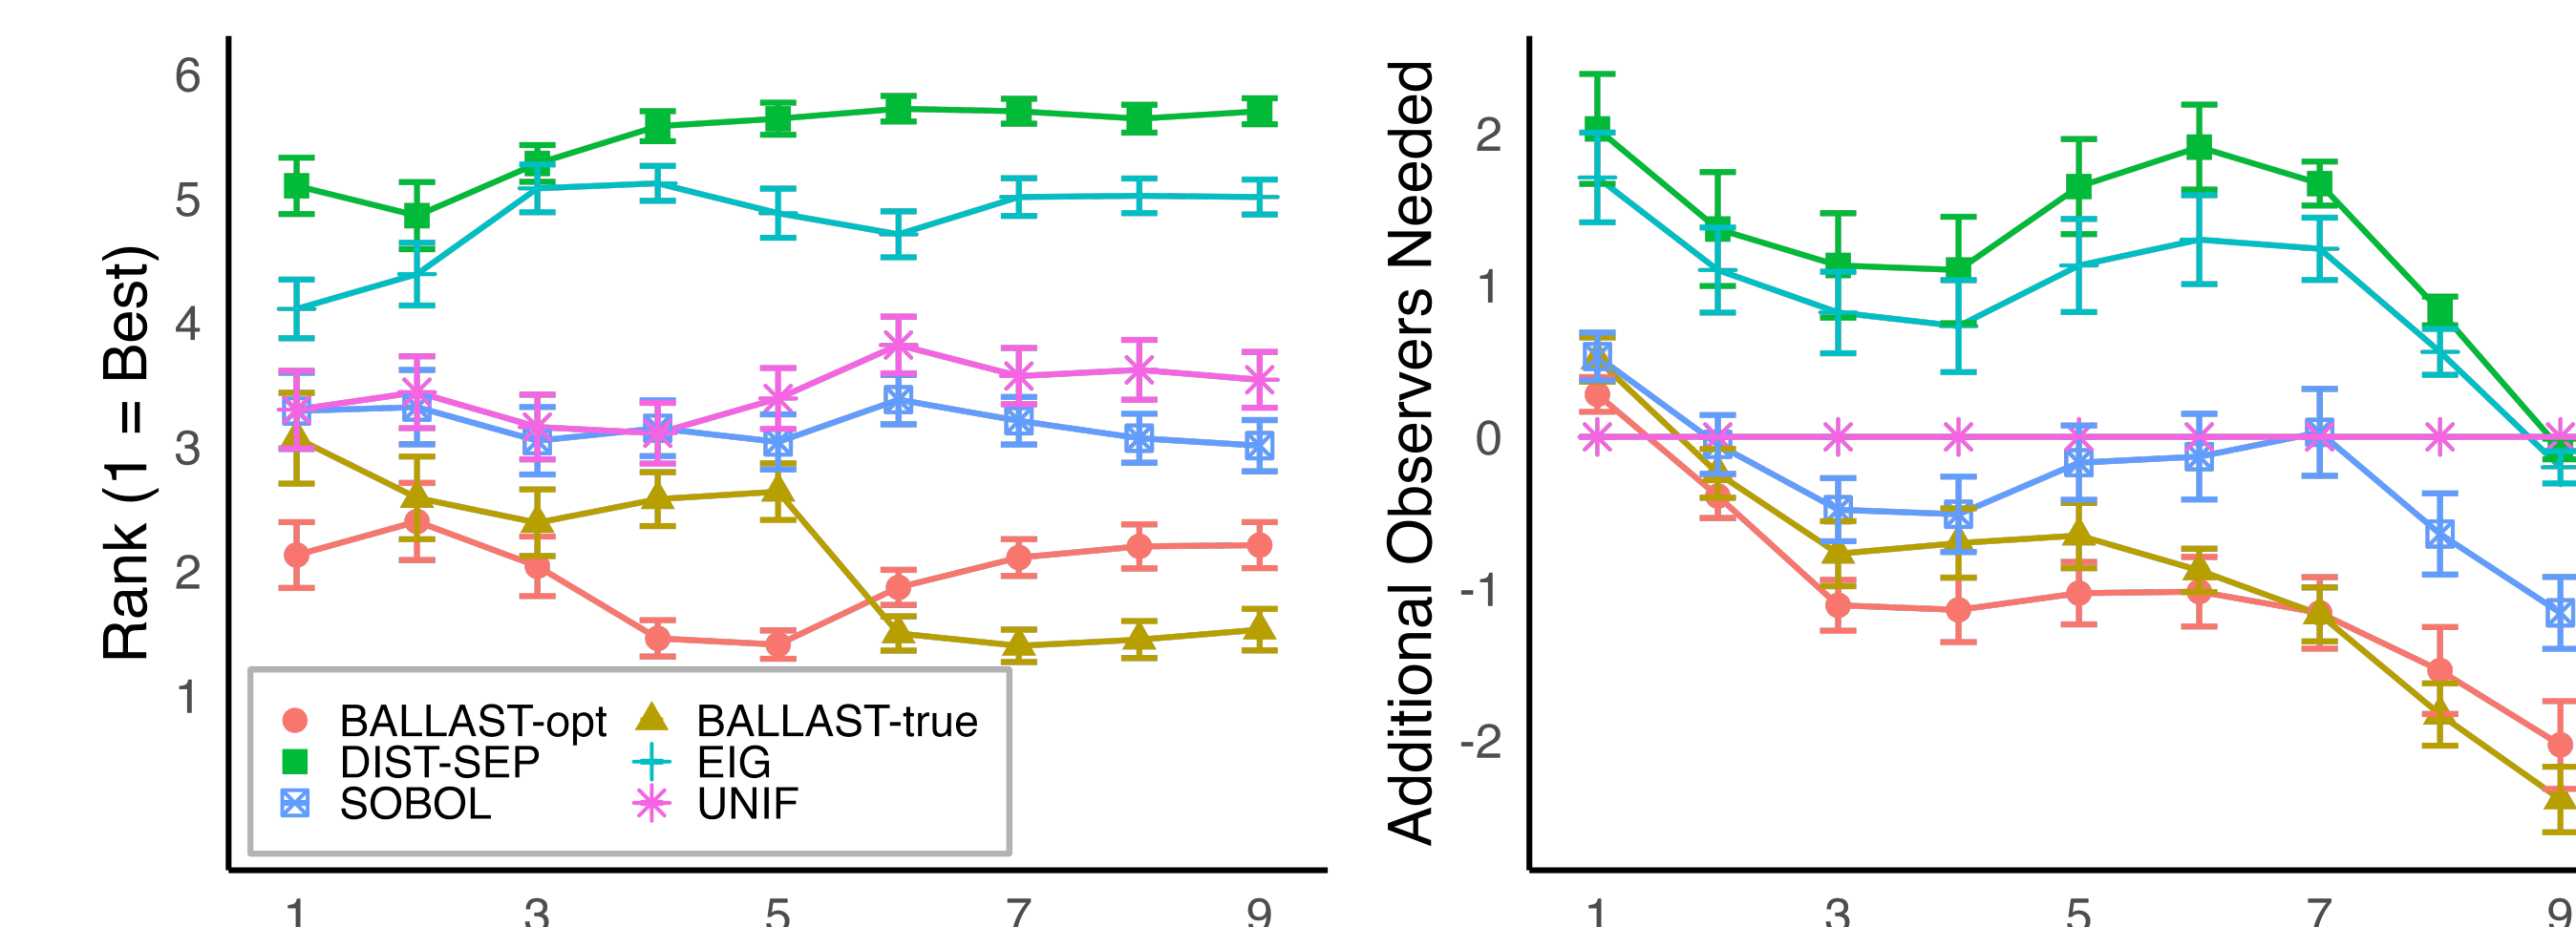
Optimality gap of deployment choice at different decision times with varying sample numbers.



20 samples are enough to reach 1% optimality gap.

Experiment: SUNTANS Model

On SUNTANS ocean currents, BALLAST consistently outperforms uniform, Sobol, EIG, and heuristic deployments.



Around 22% fewer drifters needed than uniform.

Key Contributions

- Showed that standard policies are inadequate and proposed **BALLAST**, a novel trajectory-aware acquisition strategy using look-ahead simulations.
- Developed VaSE, an efficient exact GP sampling method for non-gridded observations.
- Demonstrated consistent gains on synthetic and high-fidelity ocean current models.