Multi-Agent Vulcan: An Information-Driven Multi-Agent Path Finding Approach

Jake Olkin*, Viraj Parimi*, Brian Williams
Massachusetts Institute of Technology

SCAN ME

Enable multiple autonomous agents to search an unknown environment for phenomena of interest without redundant exploration in limited communication situations.

Motivation

- Carbon sequestration techniques store carbon in CO_2 pools under the ocean floor
- If one of these pools begins to leak, the stored CO_2 will be released into the atmosphere
- Autonomous ocean gliders are ideal vehicles to constantly monitor and search for potential leaks
- Additionally gliders need to be aware of communication loss

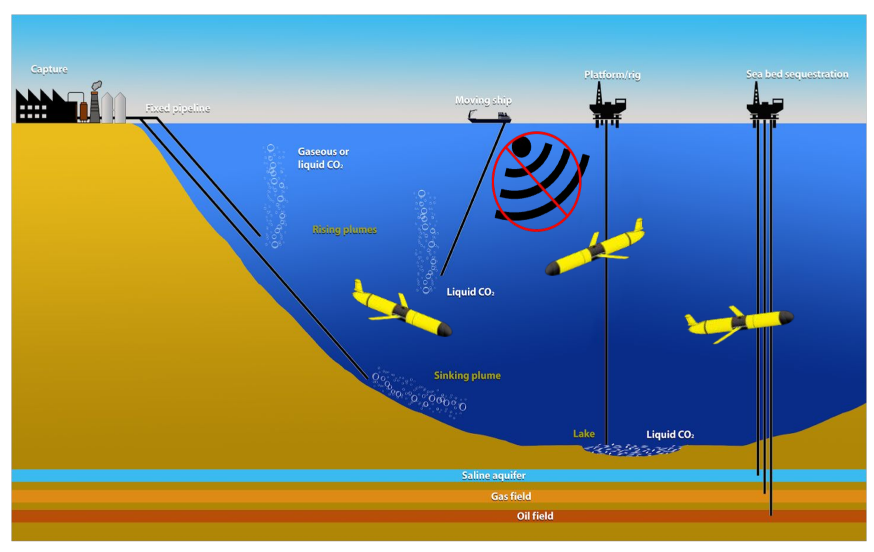

Goal: Locate as many potential CO_2 leaks as possible within a mission horizon

Problem

Autonomous vehicles tasked with searching for a phenomenon of interest will be more efficient if they make use of each other's past and potential future observation.

Problem Features

- Multiple agents & targets of interest
- Finite mission horizon
- Communication between agents is only reliable at short distances
- Spatial correlation between observed features and phenomena of interest

Key Insights

- Coordination is most important to minimize redundant observations
- Observations inform each other ↔ coupled reward function
- A decoupled reward function can act as an admissible heuristic

Approach

- Online planning procedure that updates agents' paths after each newly acquired observation
- When within communication range, agents form "*coordination bubbles*" that determine which agents will coordinate their next moves
- *Coordination bubbles* will share their past observation histories so an elected "*leader*" can plan coordinated paths for the entire bubble

Algorithm 1 High-level overview of the approach

Input Environment $\mathcal{E}$, Agents $A = \{a_1, \dots, a_k\}$
Mission Duration H, Communication Range r

while $t \leq H$ **do**
  for all $a_i \in A$ **do**
    $N_i \leftarrow \{a_j \mid d(a_i, a_j) \leq r\}$
  $\Lambda \leftarrow$ Extract minimal disjoint sets from $\{N_i \mid i \in \{1, \dots k\}\}$
  for all $\lambda_k \in \Lambda$ **do**
    $\Pi_{\lambda_k} \leftarrow$ MULTI-AGENT SEARCH$(\lambda_k, \mathcal{E})$
  for all $a_k \notin \Lambda$ **do**
    $\Pi_{a_k} \leftarrow$ SINGLE-AGENT SEARCH$(a_k, \mathcal{E})$
  for all $a_i \in A$ **do**
    Execute Π_{a_i} and collect observation ω_{a_i}
  $t \leftarrow t + 1$

Decoupled Heuristic

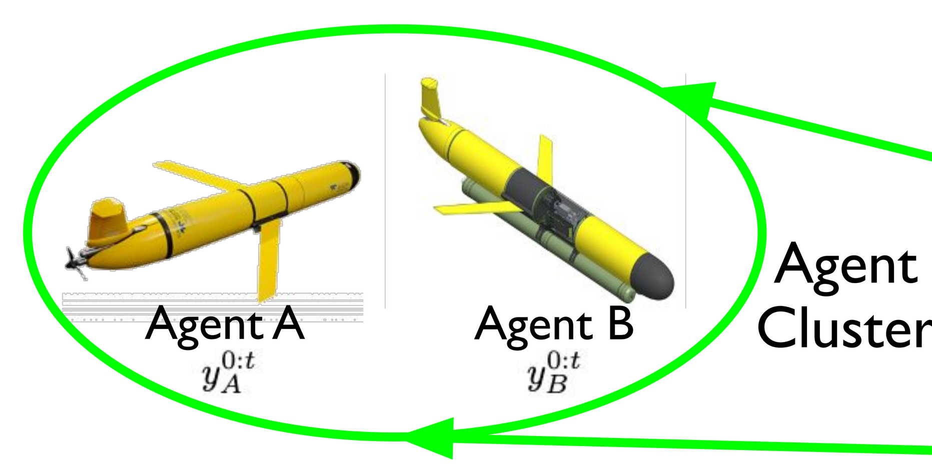

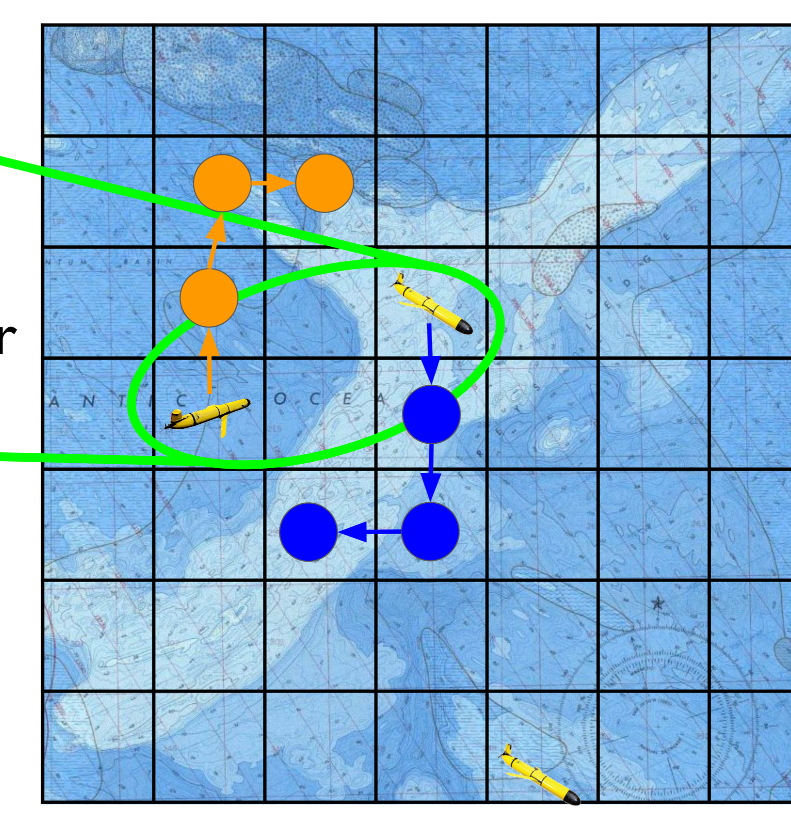

Coupled Reward Function

$$I(X_i; (Y_A^{t+1:t+h}, Y_B^{t+1:t+h}) \mid \tilde{y}^{0:t})$$

Decoupled Heuristic

$$I(X_i; Y_A^{t+1:t+h} \mid \tilde{y}^{0:t}) + I(X_i; Y_B^{t+1:t+h} \mid \tilde{y}^{0:t})$$

- Model the act of finding target phenomena as Adaptive Sampling problem using Multi-Agent Partially Observed Markov Decision Process (MA-POMDP) formulation
- Solve using receding-horizon search by maximizing the mutual information between the observations and a target phenomenon at a given location
- Space of possible future paths is exponential in the number of agents
- Reduce computation of coupled reward function
- Achieve scalability by combining information gain from individual agents' observations as a decoupled admissible heuristic reducing the number of expanded nodes

Environment and Reward Function

- X_i : RV modelling presence of target phenomenon using Gaussian Processes
- Y_i : RV modelling noisy observation function

$$R = I(\{X_i\}_{i=1}^n; Y_j^\tau \mid y_j^{0:t}) \approx \sum_{i=1}^{n} I(X_i; Y_j^\tau \mid y_j^{0:t})$$

$$I(X_i; Y_j^\tau \mid y_j^{0:t}) = \mathbb{E}_{Y_j^\tau}\left[D_{\mathrm{KL}}(p_{X_i \mid Y_j^\tau, y_j^{0:t}} \| p_{X_i \mid y_j^{0:t}})\right]$$

$$p(X_i = 1 \mid y_j^{0:t}) = \frac{P_1}{2}\left(1 - \mathrm{erf}\left(\frac{\tilde{u} - \mu}{\sqrt{2\Sigma}}\right)\right) + \frac{P_2}{2}\left(1 + \mathrm{erf}\left(\frac{\tilde{u} - \mu}{\sqrt{2\Sigma}}\right)\right)$$

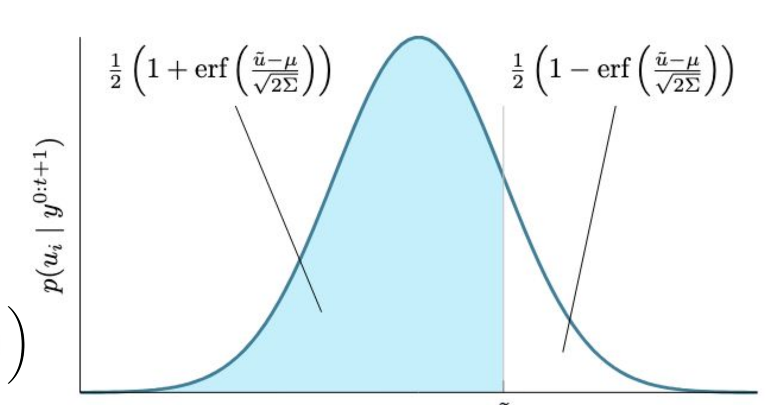

Experiments

MAPF Maps

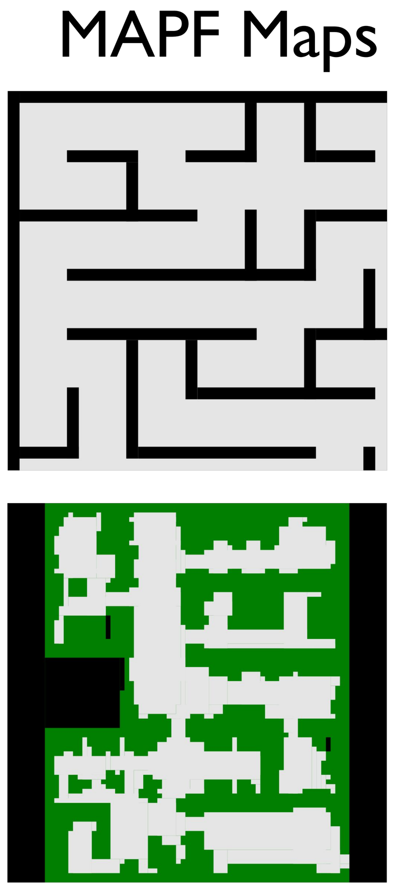

Real-Bathymetric Dataset

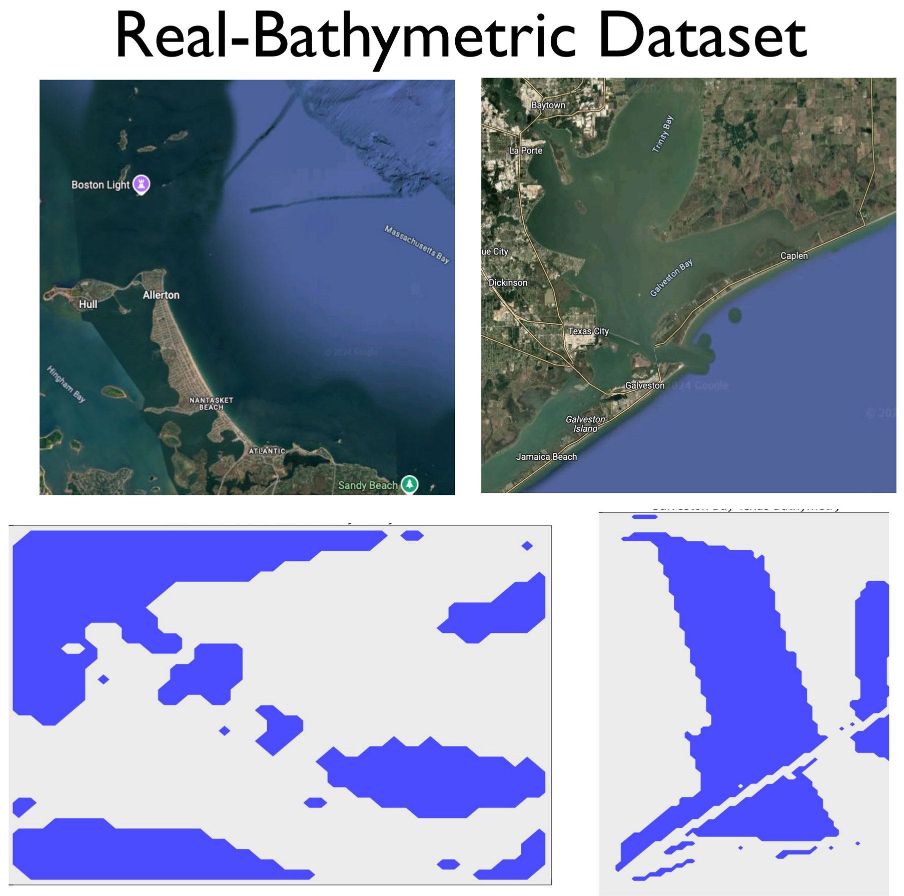

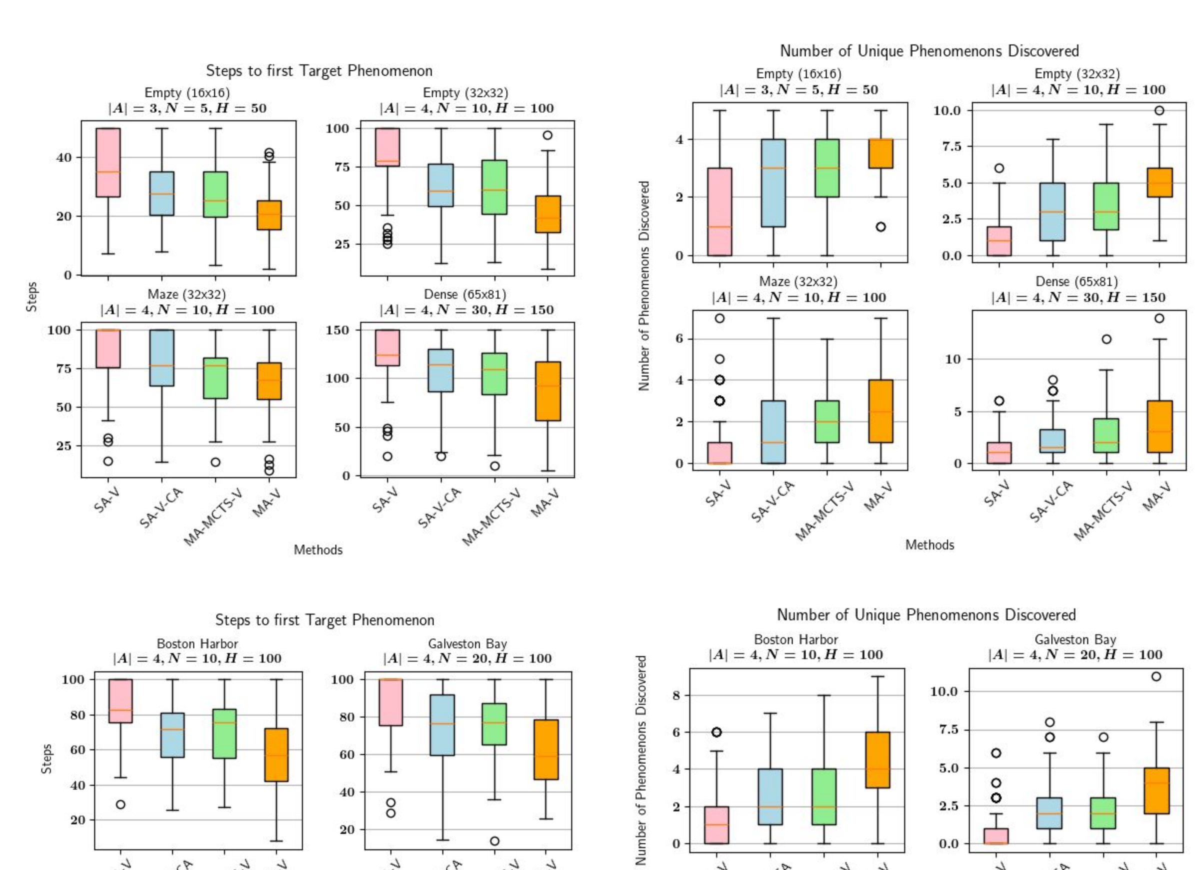

Hardware Demonstrations

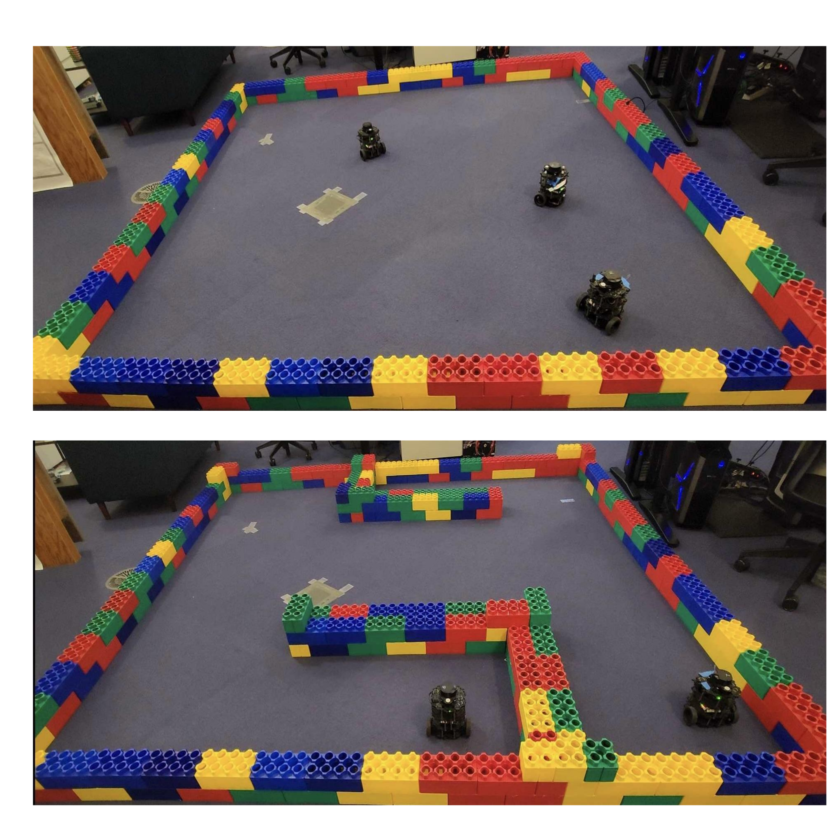

References

[1] B. Ayton, "Risk-bounded autonomous information gathering for localization of phenomena in hazardous environments," Master's thesis, Massachusetts Institute of Technology, September 2017.
[2] S. Bone, L. Bartolomei, F. Kennel-Maushart, and M. Chli, "Decentralised multi-robot exploration using monte carlo tree search," in 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 2023, pp. 7354–7361.
[3] K. H. Low, J. Dolan, and P. Khosla, "Information-theoretic approach to efficient adaptive path planning for mobile robotic environmental sensing," Proceedings of the International Conference on Automated Planning and Scheduling, vol. 19, 05 2013.

Acknowledgements

This work was supported by the British Petroleum Company (BP). Any opinions, findings and conclusions or recommendations in this material are those of the author(s) and do not necessarily reflect the views of BP.