Multi-Agent Active Search and Rescue

Ramina Ghods[1], Arundhati Banerjee[1], William J. Durkin[2], Jeff Schneider[1]

[1] Carnegie Mellon University
[2] Ohio State University

NEURAL INFORMATION PROCESSING SYSTEMS
3rd Robot Learning Workshop

Active Search and Rescue

- Consider the gridded area in this figure to be an area of interest for the active search mission where the marks X show the location of the rescue subjects.
- We are interested in multi-agent setting where multiple ground robots (agents) are sent to the environment to locate said rescue subjects as fast as possible.

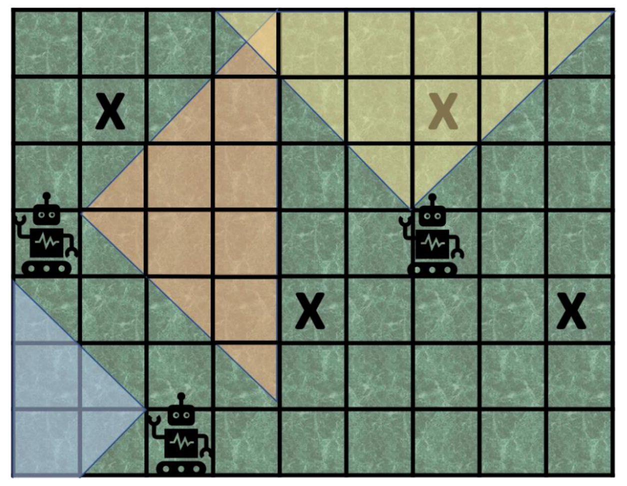

Active search and rescue defines the problem of efficiently locating rescue mission targets in an unknown environment by interactively collecting data.

Multi-Agent Search: Challenges

❖ There are multiple challenges involved with developing a **practical** multi-agent algorithm applicable to ground robots:

1. **Centralized planning is not practical.**
   - A central coordinator of all agents that expects synchronicity is not feasible as any communication or agent failure could disrupt the entire process [4].
2. **A real autonomous robot has sensing impairments.**
   - *It reports object detections probabilistically, and its precision-recall curves degrade with distance between the object and the detector.*
   - *Its performance is further constrained by the field of view as well as occlusions created by terrain or other obstacles in the scene.*

Unfortunately…

☒ Existing active search algorithms are in general only amenable to a single agent, or if they extend to multi-agent they require a central control system to coordinate the actions of all agents [1].

☒ To the best of our knowledge, no existing active search method has modelled object detection uncertainty in terms of distance [2].

Our Proposed Method (1)

❖ **To consider sensor impairments:**

- Describe the mission environment with a sparse vector β with k non-zero elements at the location of targets.

$$y_t = X_t\beta + n_t, \qquad n_t \sim \mathcal{N}(0, \Sigma_t)$$

- X_t describes the sensing matrix at time t (colorful triangles representing the robot's field of view).
- y_t is the observation vector modelling the output of an imperfect object detector as an accuracy measure.

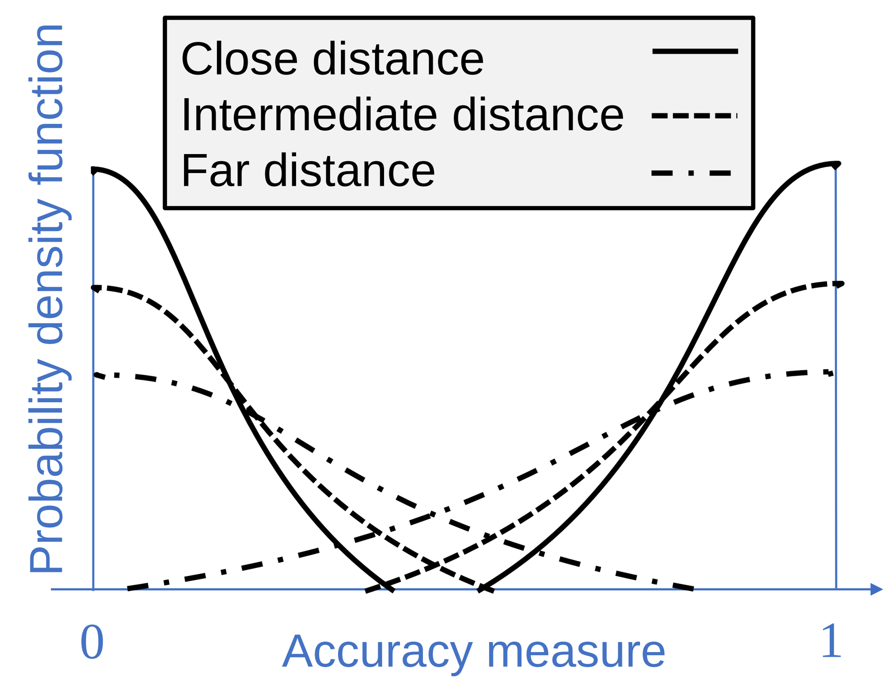

We model accuracy measure as a one-sided Gaussian distribution with its variance increasing with distance.

Our Proposed Method (2)

❖ **To develop a multi agent algorithm without the need for a central planner, we propose Thompson Sampling (TS).**

- TS is an exploration-exploitation algorithm applicable to active learning.
- We can parallelize TS to develop an active learning method where each agent can make independent and intelligent decision on the next sensing action given the available measurements.

Challenge:
Implementing TS given the sparse number of targets leads to poor performance because sparsity limits exploration capability of TS.

Solution:
We propose making an assumption that the vector β is block sparse with dynamic block length. By gradually reducing the length of the blocks from a large value, we gently trade exploration with exploitation capability over time.

Synthetic Results

❖ We now compare the performance of our proposed method **SPATS (Sparse Parallel Asynchronous Thompson Sampling)** against the information-theoretic approach called RSI proposed in [3].

- We estimate a 5-sparse signal β with size 8 × 16. Here, β is generated with a randomly uniform sparse vector. We set the noise variance to 1.
- We plot the mean and standard error of the full recovery rate versus search time over 50 random trials.
- The full recovery rate is defined as the rate at which an algorithm correctly recovers the entire vector β over random trials.

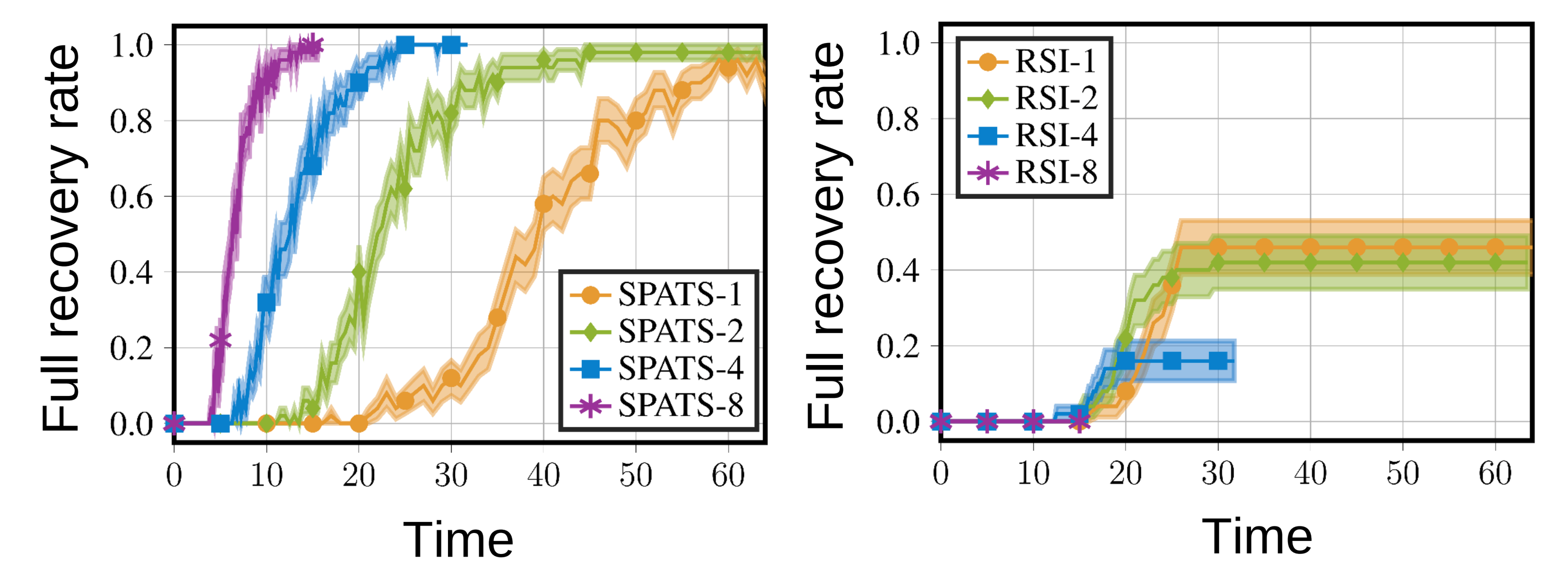

- By increasing the number of agents g, SPATS become g times faster.
- SPATS significantly outperforms RSI. (SPATS has a well-informed randomness in its reward allowing agents to make independent and intelligent decisions.)

We additionally provide theoretical analysis in [1] showing that for a 1-sparse true parameter β and reward function $R(X, \beta) = X\beta^2$:

1. block sparse prior and varying block length significantly reduces regret bounds.
2. our asynchronous multi-agent algorithm with g agents performs g times faster than a single-agent method.

Realistic Unreal Engine Environment

- We demonstrate the real-world viability of our method using a realistic environment we created in the Unreal Engine 4 game development platform with AirSim plugin.
- The QR code provides **a link to the video demonstration** of our algorithm in this environment.

[1] R. Ghods, A. Banerjee, and J. Schneider. ``Asynchronous multi agent active search.'' arXiv:2006.14718, 2020.
[2] R. Ghods, W. Durkin, and J. Schneider. ``Multi-Agent Active Search using Realistic Depth-Aware Noise Model''.arXiv: 2011.04825v1
[3] Y. Ma, R. Garnett, and J. Schneider, "Active search for sparse signals with region sensing," in Thirty-First AAAI Conference on Artificial Intelligence, 2017.
[4] Z. Yan, N. Jouandeau, and A. A. Cherif. A survey and analysis of multi-robot coordination. International Journal of Advanced Robotic Systems, 10(12):399, 2013.