

GUTS: Generalized Uncertainty-Aware Thompson Sampling for Multi-Agent Active Search

Nikhil Angad Bakshi¹, Tejus Gupta¹, Ramina Ghods¹ and Jeff Schneider¹

Abstract—Robotic solutions for quick disaster response are essential to ensure minimal loss of life, especially when the search area is too dangerous or too vast for human rescuers. We model this problem as an asynchronous multi-agent active-search task where each robot aims to efficiently seek objects of interest (OOIs) in an unknown environment. This formulation addresses the requirement that search missions should focus on quick recovery of OOIs rather than full coverage of the search region. Previous approaches fail to accurately model sensing uncertainty, account for occlusions due to foliage or terrain, or consider the requirement for heterogeneous search teams and robustness to hardware and communication failures. We present the Generalized Uncertainty-aware Thompson Sampling (GUTS) algorithm, which addresses these issues and is suitable for deployment on heterogeneous multi-robot systems for active search in large unstructured environments. We show through simulation experiments that GUTS consistently outperforms existing methods such as parallelized Thompson Sampling and exhaustive search, recovering all OOIs in 80% of all runs. In contrast, existing approaches recover all OOIs in less than 40% of all runs. We conduct field tests using our multi-robot system in an unstructured environment with a search area of $\approx 75,000 m^2$. Our system demonstrates robustness to various failure modes, achieving full recovery of OOIs (where feasible) in every field run, and significantly outperforming our baseline.

I. INTRODUCTION

Multi-robot systems have been increasingly applied to large-scale search and rescue operations [2], [3], [4], [5]. Such emergency operations are time-critical, hence search efficiency is crucial. While human operators can remotely control a small number of vehicles, they cannot efficiently coordinate larger teams. This motivates the requirement for decentralized multi-robot systems capable of efficient active search in large unstructured terrestrial environments.

Active search is the problem of making efficient sequential data-collection decisions to identify sparsely located OOIs, while adapting to new sensing information [1], [7]. The basic framework for active search consists of maintaining a posterior probability distribution over locations of OOIs and optimizing for information-seeking actions. These actions must balance surveying unseen areas (exploration) and confirming suspected OOIs (exploitation) [21], [22], [23].

Even though this basic framework is straightforward, several additional challenges need to be addressed for deploying these algorithms on a multi-robot system in realistic search and rescue scenarios. OOIs may be occluded to overhead or lateral views; hence, a combination of ground and aerial

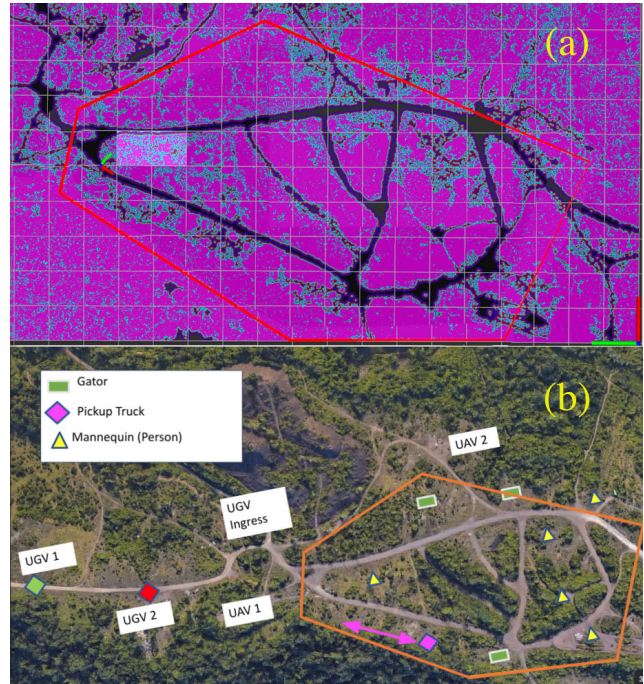


Fig. 1: Test-field in Pittsburgh, PA. (a) shows the traversal costmap and the search polygon (purple is high cost and black is low cost), and (b) shows an overhead image with the most common OOI locations and launching areas of the robots. The area inside the polygon is roughly 75,000 sq meters.

search vehicles is essential to canvass such a search region thoroughly. This requires the search algorithm to reason through occlusions, viewshed constraints, as well as each robot's sensing model and navigational constraints to plan their search actions. Similarly, though multi-robot teams can theoretically partition the search region for exploration efficiency [25], generating such a partitioning is challenging in a decentralized system. Though we would like the robots to share as much sensing information as possible for effective coordination, we cannot rely on consistent communication. Therefore, we require our system to be asynchronous and capable of making intelligent decisions with partial information. Finally, the number of OOIs in the search region is usually unknown and sensing actions are noisy [30], [31], [32]. Previous search methods make strong simplifying assumptions [19], [6] and don't deal with these challenges fully, making it intractable to deploy them on realistic search settings.

We propose a novel search algorithm, Generalized Uncertainty-aware Thompson Sampling (GUTS) that follows the parallelized Thompson Sampling framework [1], [7], [8]. We propose an improved reward function, realistically model

¹N. A. Bakshi, T. Gupta, R. Ghods and J. Schneider are with the Robotics Institute, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213 {nabakshi, tejusg, rghods, schneide}@cs.cmu.edu

sensing noise in UAVs and UGVs, and computationally optimize it to be run on large environments.

We design GUTS to be applicable to any unstructured search region, without assumptions on number of OOIs, and to be robust to communication and hardware failures. For instance, our search team continues to operate even if some robots fail and stop responding.

We demonstrate our search system consisting of UGVs and UAVs on field tests in a large unstructured natural environment in Pittsburgh, PA shown in Fig. 1 and on simulated tests. Our results show the robustness of our system, along with superior search performance compared to existing approaches. To the best of our knowledge, this is the first heterogeneous multi-robot system for asynchronous multi-agent active search built at this scale that is robust to communication and robot failures and operates without any human direction or manual subdivision of the search region.

II. RELATED WORK

Previous works on multi-robot search systems work in more specialized settings (e.g., homogeneous teams, single agent search) and make stronger assumptions (perfect communication, homoscedastic noise model). For instance, coverage-based search methods [10], [11], [12], [24] exhaustively navigate the search space. They assume that the search area is mapped and do not adapt to new information online. A similar criticism exists for frontier-based methods like [27].

Bayesian active-search methods [30], [31] [32] can accurately model observation uncertainty when searching for sparse signals with a single agent. However, the deterministic nature of their search policies makes extending these methods to the multi-agent setting difficult without centralization. Recent work has modelled the active-search problem as a POMDP and trained search policies in simulation using reinforcement learning [22], [21], [23], [38], [28]. It is usually difficult to deploy these methods on robotic systems because of their high sample complexity and unrealistic simulations.

Ghassemi et al. [6] employ an information-theoretic approach to asynchronous multi-agent active search, but work with a homogenous swarm assuming perfect communication and a homoscedastic noise model. Salman et al. [19] generate trajectories such that the time spent in an area corresponds to the likelihood of the presence of OOIs. However, their method does not incorporate new observations and relies on perfect communication and a strong unified prior on locations of OOIs.

Dec-MCTS [26] is a multi-agent active-perception algorithm that has shown promising simulated results. Dec-MCTS is a decentralized version of MCTS where all robots optimize their trajectories asynchronously but achieve coordination by sharing their partial search trees and do not model uncertainties in realistic perception pipelines.

Noise-Aware Thompson Sampling (NATS) [1] is an asynchronous multi-agent active-search algorithm that can deal with heteroscedastic noise models and has shown promising simulated results. NATS uses Myopic Posterior Sam-



Fig. 2: Left: RecBot UGV. Right: Lil Hexy UAV.

pling (MPS) [9] to ensure that the agents take diverse sensing actions without centralization. Thompson Sampling (i.e. MPS) [15] is an online optimization method that balances exploration and exploitation by maximizing the expected reward assuming that a sample from the posterior is the true state of the world [16]. This stochastic nature is the key attribute that makes it an excellent candidate for an asynchronous multi-agent setting with unreliable communication as it promotes some diversity in action selection.

III. CONTRIBUTIONS

- We propose the Generalised Uncertainty-aware Thompson Sampling (GUTS) algorithm, that uses an improved reward design, more accurately models sensing noise, and is computationally optimized over prior parallelized Thompson Sampling methods [1], [7].
- We present a working system of heterogeneous robots that is capable of asynchronous multi-agent active search in large unstructured natural environments without a central planner. We design our multi-agent search algorithm to be robust to loss of robots mid-search, noisy observations and communication breakdown without any human intervention.
- We experimentally show superior search success rate and search efficiency compared to existing search methods, namely NATS and coverage-based search, in field and simulated experiments.

IV. SYSTEM DESCRIPTION

This section describes our multi-robot system on which we evaluate our search algorithm. For the UGV, we use the RecBot, a John Deere E-Gator utility vehicle that has been retro-fitted with a drive-by-wire autonomy kit (Fig. 2-left). For the UAV, we use the Lil Hexy¹ air-frame paired with a PixHawk flight controller² as our base platform (Fig. 2-right). We added a custom camera payload and additional onboard compute, namely, NVIDIA Jetson AGX Xavier to this platform.

We summarize the components of our autonomy stack necessary to understand the work in this paper below.

- **Camera FOV:** The sensor pod on the UGV has a 37° vertical and 192° horizontal FOV (Fig. 3-left). The sensor module consists of five RGB cameras, each of 12

¹<https://vscl.tamu.edu/vehicles/little-hexy/>

²<https://pixhawk.org/>




Fig. 3: Left: Top-down view of UGV camera field-of-view. Span is 192° horizontally and 37° vertically. Right: UAV camera field-of-view. Span is $19^\circ \times 16^\circ$. At a flight height of 80m, this leads to approximately $27.8m \times 22.5m$ rectangle of visibility on the ground.

MP, to allow for accurate OOI detection up to several hundred meters away. A sensor pod with one RGB camera of 5 MP resolution is mounted on the UAV pointed downward with a 5° angle with the vertical, tilted in the direction of the travel of the UAV (Fig. 3-right). The $19^\circ \times 16^\circ$ FOV on this camera would cover an area of $27.8m \times 22.5m$ from a flight height of 80m.

- **Object Detector:** For the UGV, we use a customized version of the widely used YOLOv3 neural network architecture [13] with fast inference on the NVIDIA RTX 6000 GPU given very high-resolution input images. For the UAV detector, the computing power is limited; so we use a less computationally intensive model [14] and run it on the NVIDIA Jetson AGX Xavier.
- **Object Tracking:** The perception module uses a Kalman filter [34] to track possible OOIs. We describe our sensing model in detail in section VI. For example, the estimated OOI locations detected by the ground-vehicles have more uncertainty along the range to the object compared to the bearing direction. This is represented by the covariance associated with a track returned by the Kalman filter. Moreover, this uncertainty increases with object distance from the observer.
- **Navigation for UGV:** We utilise the SBPL planner³ with a local lookahead controller to chart safe paths to waypoints provided by the search module being presented in this paper. We update the offline global map with on-robot sensing to generate the traversal costmap.
- **Objects of Interest:** The OOIs are people (mannequins), a John Deere E-Gator, or a pickup truck. Fig. 4 shows example frames from the UGV and UAV camera systems with successful detections.

V. PROBLEM FORMULATION

We model the search region as a grid with a cell size of 30m x 30m. We have access to the costmap for ground robot traversability. Fig. 1 shows an overhead view and the costmap for the field testing site in Pittsburgh, PA. This costmap was generated from a drone flyover several months before field testing. Some parts of the map have changed since then, but our on-robot sensing and mapping system can recognize changes like blocked paths. Due to the ubiquity of satellite imagery, it is reasonable to assume such approximate map information of the search region to be available.

³<https://github.com/sbpl/sbpl>




Fig. 4: Object detector: qualitative results. We show images and OOI detections for the UGV (left) and the UAV (right)

As is evident from Fig. 1, several areas cannot be accessed or viewed by ground vehicles. In addition, large portions of the map are covered with dense foliage, which may obscure any OOIs hidden under it from the view of the air vehicles. This necessitates heterogenous coordination to explore the search region completely.

We model the active search problem as follows:

- We give the same search region to all robots. (see the red polygon in Fig. 1 for an example).
- The OOIs are sparsely placed and static. They need to be located as fast as possible with high certainty.
- Each robot must plan its next data collection action on-board, i.e., no central planner exists.
- The robots may communicate their locations and observations with each other; however, the algorithm should function anywhere in the spectrum of total absence of communication to perfect communication.

We note that the objective is *not to cover* the entire search space. Rather it is specifically to locate all OOIs as fast as possible.

More formally, we represent the locations of OOIs in our 2D grid representation using a sparse matrix $\mathbf{B} \in \mathbb{R}^{M_1 \times M_2}$. Let $\beta \in \mathbb{R}^M$ be the flattened version of matrix \mathbf{B} , where $M = M_1 M_2$. This vector is sparsely populated, with 1 at locations corresponding to OOIs and 0 elsewhere. We model the uncertainty in our observations by

$$y_t = \text{clip}(\mathbf{X}_t \beta \pm \mathbf{n}_t, 0, 1) \text{ with } \mathbf{n}_t \sim \mathcal{N}^+(0, \Sigma_t) \quad (1)$$

where $\mathbf{X}_t \in \mathbb{R}^{Q \times M}$ describes the sensing matrix such that each row in \mathbf{X}_t is a one-hot vector indicating one of the grid

cells in view of the robot at timestep t , and Q is the total number of grid cells the robot can view. $\mathbf{y}_t \in \mathbb{R}^{Q \times 1}$ is the resultant observation including the additive depth-aware noise vector $\mathbf{n}_t \in \mathbb{R}^{Q \times 1}$. This depth-aware noise encodes the intuition that observation uncertainty increases with distance to the robots, and we model it by having the diagonal elements of the noise covariance matrix Σ_t monotonically increase with the distance of the observed cell from the robot. The off-diagonal elements of Σ_t are zero. The noise is sampled from a positive half-Gaussian distribution $\mathcal{N}^+(0, \Sigma_t)$ and is added for cells without OOI and subtracted for cells with OOIs. The observation y_t is clipped to be within 0 and 1.

Let \mathbf{D}_t^j be the set of observations available to robot j at timestep t . \mathbf{D}_t^j comprises of $(\mathbf{X}_t, \mathbf{y}_t)$ pairs collected by robot j as well as those communicated to robot j by other robots. Let the total number of sensing actions by all agents be T . Our main objective is to sequentially optimize the next sensing action \mathbf{X}_{t+1} based on \mathbf{D}_t^j at each timestep t to recover the sparse signal β with as few measurements T as possible. Each robot optimizes this objective based on its own partial dataset \mathbf{D}_t^j in a decentralized manner.

VI. GENERALISED UNCERTAINTY-AWARE THOMPSON SAMPLING (GUTS)

This section presents the GUTS algorithm. The framework for decision-making follows [1]. Each robot j asynchronously estimates the posterior distribution over OOI locations based on its partial dataset \mathbf{D}_t^j . During the action selection stage, each robot generates a sample from this posterior and optimizes a reward function for this sampled set of OOI locations. Sec. VI-A and Sec. VI-B describe both these steps for the Noise-Aware Thompson Sampling (NATS) algorithm presented in [1]. We utilize the posterior computation step in Sec. VI-A for GUTS.

In Sec. VI-C and onward, we describe the new reward function to improve search performance, the depth-aware noise models for our UGVs and UAVs, and speed-up techniques to optimise the code to run on large environments.

A. Calculating Posterior

Each robot assumes a zero-mean gaussian prior per entry of the vector β s.t. $p_0(\beta_m) = \mathcal{N}(0, \gamma_m)$. The variances $\Gamma = \text{diag}([\gamma_1 \dots \gamma_M])$ are hidden variables which are estimated using data. We follow [17], [18] and use a conjugate inverse gamma prior on γ_m to enforce sparsity s.t. $p(\gamma_m) = IG(a_m, b_m) = \frac{b_m^{a_m}}{\Gamma(a_m)} \gamma_m^{-(a_m+1)} e^{-(b_m/\gamma_m)} \forall m \in \{1 \dots M\}$. We estimate the posterior distribution on β given data \mathbf{D}_t^j for robot j using Expectation Maximisation [40].

We can write analytic expressions for the E-step (estimating $p(\beta|\mathbf{D}_t^j, \Gamma) = \mathcal{N}(\mu, \mathbf{V})$) and M-step (computing $\max_{\Gamma} p(\mathbf{D}_t^j|\beta, \Gamma)$) respectively:

$$\mathbf{V} = (\Gamma^{-1} + \mathbf{X}^T \Sigma \mathbf{X})^{-1}; \mu = \mathbf{V} \mathbf{X}^T \Sigma \mathbf{y} \quad (2)$$

$$\gamma_m = ([\mathbf{V}]_{mm} + [\mu]_m^2 + 2b_m) / (1 + 2a_m) \quad (3)$$

where \mathbf{X} and \mathbf{y} are created by vertically stacking all measurements $(\mathbf{X}_t, \mathbf{y}_t)$ in \mathbf{D}_t^j , and Σ is a diagonal matrix composed of their corresponding depth-aware noise variances.

Each robot estimates $p(\beta|\mathbf{D}_t^j) = \mathcal{N}(\mu, \mathbf{V})$ on-board using its partial dataset \mathbf{D}_t^j . We drop the index j from equations (2) and (3) for clarity. We set the values of $a_m = 0.1$ and $b_m = 1$ as these were found to be effective in [1]. Finally, agent j samples from the posterior $\tilde{\beta} \sim p(\beta|\mathbf{D}_t^j)$.

B. Choosing Next Sensing Action

Each robot chooses the next sensing action \mathbf{X}_{t+1} by assuming that the sampled set of OOI locations $\tilde{\beta}$ is correct. Specifically, let $\hat{\beta}(\mathbf{D}_t^j \cup (\mathbf{X}_{t+1}, \mathbf{y}_{t+1}))$ be our expected estimate of the parameter β using all available measurements \mathbf{D}_t^j and the next candidate measurement $(\mathbf{X}_{t+1}, \mathbf{y}_{t+1})$. Then the reward function to evaluate each possible sensing action is defined as:

$$\mathcal{R}(\tilde{\beta}, \mathbf{D}_t^j, \mathbf{X}_{t+1}) = -\mathbb{E}_{\mathbf{y}_{t+1}|\mathbf{X}_{t+1}, \tilde{\beta}} [\|\tilde{\beta} - \hat{\beta}(\mathbf{D}_t^j \cup (\mathbf{X}_{t+1}, \mathbf{y}_{t+1}))\|_2^2] \quad (4)$$

We select the sensing action \mathbf{X}_{t+1} that maximizes this reward. The reward function is stochastic due to the sampling of $\tilde{\beta}$ and this ensures that the search actions selected by the robots are diverse.

C. GUTS Reward function

During action-selection, each robot optimizes a sensing location by maximizing the reward function in equation (4). We observed that this reward function leads to poor search performance at the start of the run. At the start of the episode, the posterior belief over OOI locations is very uncertain and all possible next waypoints have almost-zero probability of finding an object. Hence, this reward function leads to excessively explorative behavior.

We modify the reward function to

$$\mathcal{R}(\tilde{\beta}, \mathbf{D}_t^j, \mathbf{X}_{t+1}) = -\mathbb{E}_{\mathbf{y}_{t+1}|\mathbf{X}_{t+1}, \tilde{\beta}} [\|\tilde{\beta} - \hat{\beta}(\mathbf{D}_t^j \cup (\mathbf{X}_{t+1}, \mathbf{y}_{t+1}))\|_2^2] - \lambda \times I(\tilde{\beta}, \hat{\beta}) \quad (5)$$

Where $\lambda (= 0.01)$ is a hyperparameter that reduces the reward for a search location if the estimated $\tilde{\beta}$ does not have high likelihood entries in common with the sample at the current step $\hat{\beta}$. This change improves the discriminatory power of the reward and encourages exploitative actions early in a search run. Formally, let \hat{k} and \tilde{k} be the number of non-zero entries in $\hat{\beta}$ and $\tilde{\beta}$, then the indicator function $I(\cdot)$ is defined as:

$$I(\tilde{\beta}, \hat{\beta}) = \begin{cases} 0, & \text{if any matches between top } \frac{\hat{k}}{2} \text{ entries in } \hat{\beta}, \\ & \text{and top } \frac{\tilde{k}}{2} \text{ entries in } \tilde{\beta} \\ 1, & \text{otherwise} \end{cases}$$

Our experiments show that this modified reward function achieves better search efficiency.

D. UAV Sensing Action Model

The flight height of each UAV is fixed between 70m – 90m, with an approximate FOV of 30m × 30m on the ground below it. We consider flight paths that are straight lines and integrate over cells that will be flown over en-route to the final waypoint in the reward computation.

E. UGV Sensing Action model

We model the FOV of the UGV as a $30m \times 60m$ space directly in front of it. In the GUTS grid representation, this would be the two cells along the bearing of the UGV. We approximate the FOV conservatively as opposed to the full FOV depicted in Fig. 3. Overestimating the FOV would be disastrous since GUTS would then erroneously assume some parts of search region have been searched and never find OOIs in that region. Note that the perception module still reports positive detections even beyond the modelled FOV, hence there is no loss of information.

We do not consider possible OOIs en route to a candidate sensing action \mathbf{X}_t for computing the reward function for UGVs because computing and integrating over trajectories for each possible sensing action \mathbf{X}_t is expensive. Each robot still logs and shares observations along the way to a waypoint even though they weren't explicitly accounted for. The problem of integrating trajectory-level information gain for UGVs will be tackled in future work.

F. Modified Depth-Aware Noise Modelling

We describe our heteroscedastic noise-model for OOI observations in this section. NATS [1] models this noise using a diagonal covariance matrix where the diagonal elements are proportional to the distance between the observed grid cell and the robot. Intuitively, this represents our confidence in existence of a detected OOI at the grid cell in consideration. However, this is not a quantity that is actually available for learning-based perception systems.

We combine the detection confidence of our object detector with the location uncertainty obtained through the Kalman Filter to estimate this confidence. We employ the following heuristic to approximate the existence uncertainty for positive observations:

$$\sigma^2 = \min\left(0.5, \frac{\text{Vol. of OOI loc uncertainty ellipsoid (m}^3\text{)}}{\text{OOI detection confidence} \times 1000}\right)$$

This formula is based on the intuition that the location uncertainty obtained by the Kalman filter implicitly encodes existence uncertainty as well. We re-use the original depth-aware noise-model used in [1] for negative observations.

G. Accelerating GUTS

We describe a simple example to make the scale of the speedup clear. For an area of 0.7 sq.km, the original NATS implementation [1] took over 60 minutes to compute the first waypoint. After the following three changes, we cut the waypoint selection time to under two minutes:

- GUTS needs to compute \mathbf{V} in (2) repeatedly for candidate sensing action \mathbf{X}_{t+1} . The matrix inversion step is expensive. We leverage the diagonal nature of $(\Gamma^{-1} + \mathbf{X}^T \Sigma \mathbf{X})^{-1}$ to speed up inversion by simply inverting the diagonal elements.
- Additionally, the $\mathbf{X}^T \Sigma \mathbf{X}$ term in (2) can be calculated additively as follows:

$$\begin{bmatrix} \mathbf{X}_1 \\ \mathbf{X}_2 \end{bmatrix}^T \begin{bmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{bmatrix} \begin{bmatrix} \mathbf{X}_1 \\ \mathbf{X}_2 \end{bmatrix} = \mathbf{X}_1^T \Sigma_1 \mathbf{X}_1 + \mathbf{X}_2^T \Sigma_2 \mathbf{X}_2$$

This leverages the independence of sensing noise in different grid cells and results in significant-speedup later in the run when \mathbf{X} is high-dimensional.

- For μ in Eqn. 2, we exploit the sparse nature of the sensing matrix X to speed up computation [37].

Since compute is limited on the UAVs, we have an additional sampling parameter that subsamples the possible sensing actions the drone could take to dial in the planning time as needed. Typically we tested within the 1% to 10% subsampling range.

VII. EXPERIMENTS AND RESULTS

Our experiments aim to demonstrate the improved search efficiency of our proposed algorithm GUTS compared to existing search methods: NATS and coverage-based search. We run a computationally optimised (similar to Sec. VI-G) version of NATS for our experiments since the original implementation [1] is too slow for the scale of our experiments. The coverage-based baseline myopically chooses the next waypoint in an unvisited part of the search region. We evaluate these methods on a realistic multi-robot Gazebo [39] simulation. Our simulation uses a pre-mapped costmap of our field test site and accurately simulates the robots' sensing, trajectory planning, and traversal.

We also seek to validate that GUTS can run in real-time on a heterogenous multi-robot team and can function robustly in the face of communication breakdown, hardware failures, and noisy detections. We conduct field tests with a combination of ground and aerial vehicles in a natural unstructured environment with a search area of $\sim 75,000m^2$.

A. Testing Setup

Fig. 1b shows the launch locations for the UGVs and UAVs, the search region for each robot, and the most common OOI locations. We start evaluating each robot's search performance once it enters the search polygon.

B. Evaluation Metric

Our primary evaluation metric is the search success rate, which is defined as the fraction of runs in which the search method locates all OOIs within the prescribed time budget.

For our simulated experiments, we measure search efficiency by plotting recall versus runtime for each algorithm. For our field tests, we compute the number of sensing actions per OOI found. We prefer to isolate each search algorithm's performance from physical factors such as terrain conditions in field tests and measure the recall in terms of the number of decisions rather than the total time taken. That said, we also report wall clock time for each run, and our total search time per run is always under 30 mins.

C. Simulated Results

We compare the search efficiency of GUTS with NATS and coverage-based search in Fig. 5. We plot results for a single UGV (center plot) and two UGVs (right plot) with five OOIs and a 1000 second (≈ 17 mins) runtime budget averaged across 20 runs and 10 runs, respectively.

TABLE I: Field Testing Results: We report the search parameters (team size, total search area, total OOIs, algorithm used) with evaluation metrics (runtime, OOIs found, and search efficiency T/C) for each run. We observe that GUTS outperforms the coverage-based planner in terms of search efficiency (T/C). We also note runs with communication breakdown (CB), hardware failures (HF), and issues with detectors (ND = noisy detector).

Run	Team Size (J)	Search Area (m^2)	Runtime (s)	OOIs found (C)	Total OOIs (M)	T	T/C	Algorithm	Notes
1	1 UAV	$\sim 75,000$	1600	4	7	27	6.75	Coverage	-
2	1 UAV	$\sim 75,000$	1500	5	8	28	5.6	Coverage	-
3	1 UAV	$\sim 75,000$	1000	2	8	8	4	5% GUTS	-
4	1 UAV	$\sim 75,000$	1600	4	8	16	4	5% GUTS	-
5	1 UAV	$\sim 75,000$	500	3	4	5	1.67	GUTS	-
6	2 UGVs	$\sim 75,000$	1000 (avg)	3	3	3	1	GUTS	ND
7	2 UGVs	$\sim 75,000$	600 (avg)	6	6	6	1	GUTS	CB, ND
8	2 UGVs 1 UAV	$\sim 75,000$	1200 (avg)	6	6	13	2.167	GUTS	CB, HF, ND




Fig. 5: Comparison of Search Efficiency. (Left) We visualize the search efficiency for each run in Table I. We plot the recall versus the number of sensing actions taken per robot. This graph shows that GUTS outperforms the coverage baseline on our system. We also observe that subsampling the total set of waypoints achieves a good balance of computational efficiency and performance, and slightly outperforms the coverage baselines on UAVs. (Middle and Right) We compare the recall rate of different search algorithms vs runtime in a single-robot and multi-robot simulation. We again observe that GUTS outperforms NATS and coverage-based search.

We can see that GUTS clearly outperforms our baselines: NATS and coverage-based search, and is the only search algorithm to consistently recover all the OOIs within our time budget. GUTS has a success rate of 80% in the multi-robot setting (i.e. it recovers all OOIs in 80% of the runs), compared to 40% and 30% success rate for NATS and coverage-based search respectively.

GUTS outperforms NATS due to the modification in the reward function described in VI-C where confirming suspected OOIs is prioritized over exploratory behaviour. This leads to a slower start but GUTS overtakes NATS towards the end of the run. Coverage-based search performs poorly because it navigates the search region exhaustively rather than optimizing its sensing actions for OOI identification.

D. Field-Testing Results

We report detailed run parameters and evaluation metrics for our field tests in Table I. We report results for 6 runs of GUTS and 2 runs of our coverage-based baseline. We plot the number of sensing actions required per agent (T/J) against the recall for each run in Fig. 5-left. We observe that the GUTS algorithm shows markedly more efficient search behaviour compared to our coverage based baseline.

We also observe that 5% GUTS, which is run on a 5% random subsample of all possible sensing actions, slightly outperforms coverage as well. We note the runs where we observed communication breakdown or hardware failure show no degradation in performance. The search team sometimes fails to recover all the OOIs in runs with only UAVs, because some OOIs are only recoverable from the ground-vehicles.

This demonstrates the need to use both UAVs and UGVs to thoroughly canvass the search region.

VIII. CONCLUSION

We present a novel algorithm, Generalized Uncertainty-aware Thompson Sampling (GUTS), for efficient multi-agent active search. GUTS uses an improved reward function, realistically models sensing noise, and is computationally optimized to be run on large environments. We show that GUTS outperforms state-of-the-art methods to more quickly and consistently recover all objects in our simulated experiments. We also validated the search efficiency and robustness of GUTS through field tests in a large unstructured environment subject to hardware and communication failures.

Our future work would focus on modelling dynamic targets, as well as integrating path-level information-gain and topographical information for selecting sensing actions to improve search performance further for the UGVs.

We believe that greater adoption of multi-robot teams in dangerous large-scale search and rescue missions has the potential to save human lives and effort. We hope our work encourages research in more realistic search settings and quicker adoption in real search missions.

IX. ACKNOWLEDGEMENT

Authors would like to acknowledge the contributions of Arundhati Banerjee, Herman Herman, Jesse Holdaway, Prasanna Kannappan, Luis Ernesto Navarro-Serment, Matthew Schnur, Wennie Tabib and Vamsavardan Vemuru.

REFERENCES

- [1] R. Ghods, W. J. Durkin, and J. Schneider, "Multi-Agent Active Search using Realistic Depth-Aware Noise Model," IEEE International Conference on Robotics and Automation (ICRA), 2021
- [2] C. Mouradian, J. Sahoo, R. H. Glitho, M. J. Morrow and P. A. Polakos, "A coalition formation algorithm for Multi-Robot Task Allocation in large-scale natural disasters," 2017 13th International Wireless Communications and Mobile Computing Conference (IWCMC), 2017, pp. 1909-1914, doi: 10.1109/IWCMC.2017.7986575.
- [3] D. S. Drew, "Multi-agent systems for search and rescue applications," *Curr Robot Rep.*, vol. 2, no. 2, pp. 189–200, 2021.
- [4] K. Nagatani et al., "Multirobot exploration for search and rescue missions: A report on map building in RoboCupRescue 2009," *J. field robot.*, vol. 28, no. 3, pp. 373–387, 2011.
- [5] Baxter, J. et al., "Multi-Robot Search and Rescue: A Potential Field Based Approach," *Autonomous Robots and Agents*, pp. 9-16, 2007.
- [6] P. Ghassemi and S. Chowdhury, "Decentralized Informative Path Planning With Balanced Exploration-Exploitation for Swarm Robotic Search," vol. 1: 39th Computers and Information in Engineering Conference. 08 2019.
- [7] R. Ghods, A. Banerjee, and J. Schneider, "Decentralized multi-agent active search for sparse signals," *Uncertainty in Artificial Intelligence*, 696-706, 2021.
- [8] K. Kandasamy, A. Krishnamurthy, J. Schneider, and B. Póczos, "Parallelised Bayesian Optimisation via Thompson Sampling," in *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, 09–11 Apr 2018, vol. 84, pp. 133–142.
- [9] K. Kandasamy, W. Neiswanger, R. Zhang, A. Krishnamurthy, J. Schneider, and B. Póczos, "Myopic Posterior Sampling for adaptive goal oriented design of experiments," in *Proceedings of the 36th International Conference on Machine Learning*, 09–15 Jun 2019, vol. 97, pp. 3222–3232.
- [10] L. Lin and M. A. Goodrich, "UAV intelligent path planning for wilderness search and rescue," in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2009, pp. 709–714.
- [11] S.-Y. Chien, H. Wang, and M. Lewis, "Human vs. algorithmic path planning for search and rescue by robot teams," in *The Human Factors and Ergonomics Society Annual Meeting*, vol. 54, no. 4. SAGE, 2010.
- [12] A. Ryan and J. K. Hedrick, "A mode-switching path planner for UAV assisted search and rescue," in *Conference on Decision and Control*. IEEE, 2005, pp. 1471–1476.
- [13] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," *arXiv preprint arXiv:1804.02767*, 2018.
- [14] W. Ke, T. Zhang, Z. Huang, Q. Ye, J. Liu, and D. Huang, "Multiple anchor learning for visual object detection," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp.10 203–10 212.
- [15] W. R. Thompson, "On the likelihood that one unknown probability exceeds another in view of the evidence of two samples," *Biometrika*, vol. 25, no. 3/4, pp. 285–294, 1933.
- [16] D. J. Russo, B. Van Roy, A. Kazerouni, I. Osband, Z. Wen et al., "A tutorial on Thompson sampling," *Foundations and Trends® in Machine Learning*, vol. 11, no. 1, pp. 1–96, 2018.
- [17] M. E. Tipping, "Sparse Bayesian learning and the relevance vector machine," *Journal of machine learning research*, vol. 1, no. Jun, 2001.
- [18] D. P. Wipf and B. D. Rao, "Sparse bayesian learning for basis selection," *IEEE Transactions on Signal processing*, vol. 52, no. 8, 2004.
- [19] H. Salman, E. Ayvali, and H. Choset, "Multi-agent ergodic coverage with obstacle avoidance," *Proc. Int. Conf. Autom. Plan. Sched.*, vol. 27, pp. 242–249, 2017.
- [20] E. Ayvali, H. Salman, and H. Choset, "Ergodic coverage in constrained environments using stochastic trajectory optimization," *arXiv [cs.RO]*, 2017.
- [21] X. Ye, Z. Lin, H. Li, S. Zheng, and Y. Yang, "Active object perceiver: Recognition-guided policy learning for object searching on mobile robots," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 6857–6863.
- [22] Y. Zhu et al., "Target-driven visual navigation in indoor scenes using deep reinforcement learning," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 3357–3364.
- [23] N. Xu, "Active object searching on mobile robot using reinforcement learning," in *2021 2nd International Conference on Computing and Data Science (CDS)*, 2021, pp. 296–300.
- [24] A. Sadeghi and S. L. Smith, "Heterogeneous task allocation and sequencing via decentralized large neighborhood search," *Unmanned syst.*, vol. 05, no. 02, pp. 79–95, 2017.
- [25] D. W. Schuldt and J. A. Kurucur, "Efficient partitioning of space for multiple UAS search in an unobstructed environment," *Int. J. Intell. Robot. Appl.*, vol. 2, no. 1, pp. 98–109, 2018.
- [26] G. Best, O. M. Cliff, T. Patten, R. R. Mettu, and R. Fitch, "Dec-MCTS: Decentralized planning for multi-robot active perception," *Int. J. Rob. Res.*, vol. 38, no. 2–3, pp. 316–337, 2019.
- [27] P. Quin, D. D. K. Nguyen, T. L. Vu, A. Alempijevic, and G. Paul, "Approaches for efficiently detecting frontier cells in robotics exploration," *Front. Robot. AI*, vol. 8, p. 616470, 2021.
- [28] M. T. J. Spaan, G. J. Gordon, and N. Vlassis, "Decentralized planning under uncertainty for teams of communicating agents," in *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems - AAMAS '06*, 2006.
- [29] T. H. Chung, G. A. Hollinger, and V. Isler, "Search and pursuit-evasion in mobile robotics: A survey," *Auton. Robots*, vol. 31, no. 4, pp. 299–316, 2011.
- [30] Y. Ma, R. Garnett, and J. Schneider, "Active search for sparse signals with region sensing," *Proc. Conf. AAAI Artif. Intell.*, vol. 31, no. 1, 2017.
- [31] B. Jedynak, P. I. Frazier, and R. Sznitman, "Twenty questions with noise: Bayes optimal policies for entropy loss," *J. Appl. Probab.*, vol. 49, no. 1, pp. 114–136, 2012.
- [32] P. Rajan, W. Han, R. Sznitman, P. I. Frazier, and B. M. Jedynak, "Bayesian multiple target localization," in *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, 2015, pp. 1945–1953.
- [33] Z.-Q. Zhao, P. Zheng, S.-T. Xu, and X. Wu, "Object detection with deep learning: A review," *arXiv [cs.CV]*, 2018.
- [34] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems", *Transactions of the ASME–Journal of Basic Engineering*, vol 82, no Series D, bll 35–45, 1960.
- [35] Stanford Artificial Intelligence Laboratory et al., *Robotic Operating System*. 2018.
- [36] M. R. O'Shaughnessy, M. A. Davenport, and C. J. Rozell, "Sparse Bayesian learning with dynamic filtering for inference of time-varying sparse signals," *IEEE Trans. Signal Process.*, vol. 68, pp. 388–403, undefined 2020.
- [37] P. Virtanen et al., "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python", *Nature Methods*, vol 17, bll 261–272, 2020.
- [38] C. Igoe, R. Ghods, and J. Schneider, "Multi-agent active search: A reinforcement learning approach," *IEEE Robot. Autom. Lett.*, vol. 7, no. 2, pp. 754–761, 2022.
- [39] N. Koenig and A. Howard, "Design and use paradigms for Gazebo, an open-source multi-robot simulator," *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566)*, 2004, pp. 2149–2154 vol.3, doi: 10.1109/IROS.2004.1389727.
- [40] T. K. Moon, "The expectation-maximization algorithm," *IEEE Signal Process. Mag.*, vol. 13, no. 6, pp. 47–60, 1996, doi: 10.1109/79.543975.