

Fast and Scalable Signal Inference for Active Robotic Source Seeking

Christopher E. Denniston^{1,2} Oriana Peltzer³ Joshua Ott³ Sangwoo Moon¹ Sung-Kyun Kim¹
Gaurav S. Sukhatme² Mykel J. Kochenderfer³ Mac Schwager³ Ali-akbar Agha-mohammadi¹

Abstract—In active source seeking, a robot takes repeated measurements in order to locate a signal source in a cluttered and unknown environment. A key component of an active source seeking robot planner is a model that can produce estimates of the signal at unknown locations with uncertainty quantification. This model allows the robot to plan for future measurements in the environment. Traditionally, this model has been in the form of a Gaussian process, which has difficulty scaling and cannot represent obstacles. We propose a global and local factor graph model for active source seeking, which allows the model to scale to a large number of measurements and represent unknown obstacles in the environment. We combine this model with extensions to a highly scalable planner to form a system for large-scale active source seeking. We demonstrate that our approach outperforms baseline methods in both simulated and real robot experiments.

I. INTRODUCTION

Prompt and accurate situational awareness is crucial for first responders to emergencies, such as natural disasters or accidents. Visual information is often unavailable or insufficient to detect people needing assistance in disaster-stricken environments. Wireless signals, such as radio, Wi-Fi, or Bluetooth from mobile devices, can be important in search and rescue missions. Autonomous exploration and signal source seeking by a robotic system can significantly improve the situational awareness of the human response team operating in hazardous and extreme environments. In active source seeking, a robot not only attempts to get closer to the predicted source location but also gathers more information to actively reduce the source localization uncertainty.

Signal source seeking in large, unknown, obstacle-rich environments is an important but challenging task. Source seeking systems typically consist of two components: a model of the environment that provides a belief over the signal strength at unknown locations, and a planner which generates a policy to improve the model and locate the source. The model of the environment, as well as the planner, must scale to a large area and a large number of collected measurements. Models for signal inference typically do not scale well to large numbers of measurements, relying on either approximation or having exponential growth in the time

¹ NASA Jet Propulsion Laboratory, Caltech, ² University of Southern California, ³ Stanford University. Corresponding Author: cdennist@usc.edu. Sukhatme holds concurrent appointments as a Professor at USC and as an Amazon Scholar. This paper describes work not associated with Amazon.

The work is partially supported by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004), and Defense Advanced Research Projects Agency (DARPA). ©2022. All rights reserved.

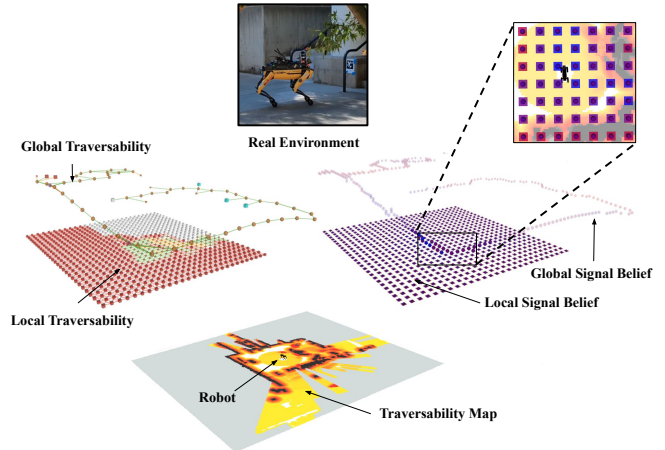


Fig. 1: **Overview of approach.** Both the information roadmap (which models the traversability of the environment) and the signal inference system maintain global and local representations of the environment. The signal inference system topology is built from the information roadmap. The environment is a paved outdoor walkway which can be seen in the image as well as the traversability map. The local signal representation maintains both the mean (shown as colored spheres) as well as the variance (flat squares), both are shown from blue (low) to red (high).

required for inference [1]–[3]. In the presence of unknown obstacles, the robot must continually re-plan at a high rate and adapt the signal model to account for newly discovered obstacles. The presence of unknown obstacles has complex impacts on signal propagation and requires the robot have the ability to re-visit un-visited corridors which potentially lead to the signal source.

In this work, we formulate the problem of finding the highest signal strength as a problem of finding the location with the maximum radios signal to noise ratio (SNR). To find this maximum, the robot takes actions which increase its understanding of the signal while the planner is creating the next policy for the robot to execute. We formulate this as an Bayesian informative path planning problem in which the robot must balance the exploration-exploitation trade off [4]. In order to find the signal source, the robot must explore to better the model of the signal in the environment, while being driven to the maxima of the signal strength.

To improve understanding of the signal strength, the robot must be able to estimate both the strength at unknown spatial locations, as well as the uncertainty in that estimate. This probabilistic signal model gives the robot the ability to plan

over future measurements to determine which locations to measure given the information about the signal strength.

To facilitate large-scale active source seeking both the planner and model of the environment must be highly scalable. To this end, we have developed a factor-graph based signal inference model and extended a large-scale multi-robot planner to handle these challenges. Both the planner and model maintain global and local representations of the environment, which allows for decomposition of the problem. This problem decomposition can be seen in fig. 1, which shows that both the information roadmap (which contains traversability information) and the signal inference are split into global and local components.

Our contributions are:

- A novel factor graph based model for belief over the signal strength which scales well with the number of measurements taken and models the signal propagation through obstacles,
- An extension to the PLGRIM planner [5] that allows it to take advantage of our factor-graph model and perform active source seeking, and
- Demonstration on challenging simulation and real-robot experiments with real-world signal strength data.

II. BACKGROUND / RELATED WORK

Active Source seeking is the process of using a robot to find a signal source in an unknown environment. Active source seeking has been used in locating radiation sources by representing the environment as voxels [6]. Active source seeking has also been used to map plumes of unknown hazardous agents [7]. Gaussian processes have been used as the model in active source seeking of radio signals by developing a control law that allows the robot to move towards the source, while updating the model in an online fashion [8]. Active source seeking has been extended to the multi-robot setting by using a particle filter based model [9].

Informative path planning (IPP) is a process in which a robot takes measurements of a concentration to maximize an information metric. Informative path planning can be used to minimize the uncertainty in the model, such as through an entropy metric [10]. Informative path planning can also be combined with Bayesian acquisition functions to form sequential Bayesian optimization [4], in which the goal is to find high concentration areas and can be used for active source seeking. Previously, this formulation has been used in finding high concentrations of chlorophyll for studying algae blooms [2], [11] or finding the quantiles of a chlorophyll distribution [12].

Gaussian Processes are non-parametric models with uncertainty quantification which are widely used for IPP [1], [4], [8], [12], [13]. They approximate an unknown function from its known outputs by computing the similarity between points from a kernel function, k , in our case the squared exponential kernel [3]. Function values y_* at any input location \mathbf{x}_* are approximated by a Gaussian distribution: $\mathbf{Y}_* | \mathbf{Y} \sim \mathcal{N}(K_* K^{-1} \mathbf{y}, K_{**} - K_* K^{-1} K_*^T)$ where \mathbf{Y} is training output, \mathbf{Y}_* is test output. The kernel matrices K ,

K_* , and K_{**} are computed by evaluating the kernel on \mathbf{X} , the training input, and \mathbf{X}_* , the test input, such that $K^{i,j} = k(x^i, x^j)$, $K_*^{i,j} = k(x^i, x_*^j)$, and $K_{**}^{i,j} = k(x_*^i, x_*^j)$.

Gaussian processes typically do not have a way to represent obstacles as the kernel function only relies only on the distance between the locations in the model. Gaussian processes tend to scale cubically in the time required for inference due to the need to compute K^{-1} , and scale linearly in the time required when adding new measurements as the kernel function k needs to be computed with the new measurement and all previous measurements [3].

Ergodic trajectory generation is rooted in the intuition that a path should spend time in a region proportional to the amount of expected information in that region [14], [15]. Ergodic trajectory design was originally presented in Mathew and Mezić [16] where they introduced a norm on the statistical distance between a trajectory and a reference distribution allowing problem to be framed as an optimization problem with the goal of achieving the lowest ergodic score. Miller [17] introduced a closed-loop ergodic control algorithm for active search problems and Dressel [18] extended this work to target source localization. Ergodic trajectory generation does not explicitly quantify the uncertainty in the surrounding environment but rather plans a trajectory assuming an information distribution is given. In this work, we use the uncertainty in the factor graph model to guide our exploration.

Factor Graphs are a common way to represent estimation problems for SLAM and other non-linear estimation problems. A factor graph is a graphical model involving factor nodes and variable nodes. The variables, or values, represent the unknown random variables in the estimation problem, such as robot poses. The factors are probabilistic information on the variables and represent measurements, such as of the signal strength, or constraints between values. Performing inference in a factor graph is done through optimization, making it suitable for large scale inference problems [19]. Factor graphs have been used for very large-scale SLAM problems [20] and for estimation of gas concentrations [21].

Factor Graph based Informative Path planning Factor graphs have been used as a model of a continuous concentration by modeling the problem as a Markov random field. This approach has been used to monitor time varying gas distributions in spaces with obstacles by defining a factor graph over the entire space with an unknown value at each cell [21]. It has also been used to jointly estimate the concentration of gas and the wind direction [22]. These approaches suffer from a scaling problem due to their ties to the geometry of the environment and do not handle unknown environments.

III. FORMULATION

In informative path planning for active source seeking, the robot is tasked to localize the source of a radio signal without any prior map of the environment and within limited time constraints. We define the state as $s = (q, W_s, W_r)$ where q is the robot state, W_s is the signal state (the signal

location) and W_r is the world traversal risk state. The task is to find a policy that maximizes an objective function S while minimizing its action cost. Typically, the objective S corresponds to the expected reduction in uncertainty in the belief over the signal source location W_s . We define the reward for taking action a at state s as

$$R(s, a) = f(S(q' | W_s), C(q, a | W_r)) \quad \text{s.t. } q' = T(q, a | W_r) \quad (1)$$

where $C(q, a | W_r)$ is the cost of taking action a from state s and T is the robot state transition function, and f is a weighted combination of the two terms defined in PLGRIM [5].

In order to maximize S , the robot must be able to infer signal strength at unseen locations, given the history of locations $\mathbf{X}_{0:t}$ and measurements $\mathbf{Y}_{0:t}$, as well as hypothetical future locations and measurements, predicted by the model. This multiple step look-ahead allows the robot to plan a policy which is non-myopic and looks ahead to the reward for complex, multiple step policies. We denote the belief at time t as $W_s^t = M(\mathbf{X}_{0:t}, \mathbf{Y}_{0:t})$ where M is a model of the belief over the signal strength in the environment. Under this model, we wish to find an optimal policy

$$\pi^* = \operatorname{argmax}_{\pi \in \Pi} \sum_{t'=t}^T \mathbb{E}[R(b_{t'}, \pi(b_{t'}))] \quad (2)$$

where t is the current planning time, $b_{t'}$ is the belief over state s at time t' , and T is the robot's time budget. As the robot progresses in the mission and updates its world belief, we re-solve the problem in a receding horizon fashion over the remaining time interval $[t, T]$.

IV. APPROACH

In this section, we introduce novel signal belief components W_s^g and W_s^l and propose a hierarchical planner for solving Eq. (2). We decompose the signal, traversability, and planning portions into global and local components to allow our system to scale easily through decomposition.

A. World Traversability Belief

Hierarchical planning frameworks address both space complexity (due to the large and increasing environment size) and model uncertainty challenges by breaking down the environment representation into components of different scales and resolutions. In PLGRIM [5], the world belief representation is composed of a meter-level resolution lattice representation of the robot's local surroundings, termed *Local Information Roadmap* (Local IRM), and a topological graph representing the explored space, called *Global Information Roadmap* (Global IRM). Together, the Local IRM W_r^l and Global IRM W_r^g compose the world traversability belief W_r . This can be seen in fig. 1 where the local and global information roadmaps are shown next to a traversability map of the environment.

B. Signal Belief

In order to facilitate large scale signal inference the model is split into global and local representations of the signal




Fig. 2: **The signal belief is represented by two factor graphs**, described in section IV-B. (a) The global factor graph represents the belief over the signal strength at locations the robot has received measurements, represented as the values g_1, \dots, g_n . (b) The local factor graph is centered at the robot using the values l_1, \dots, l_n to infer the signal locally around the robot.

Algorithm 1: Add To Global Graph

Add measurement y_i at location x_i , Given global graph G^g , Current Estimate V^g , Distance based signal variance function σ_{dist}^2 , Measurement variance σ_{meas}^2 , Minimum new global node distance d_{min} , Global values g_i, g_{i-1}

```

if  $\|x_i - x_{i-1}\| > d_{min}$  then
   $f_{link}^g \leftarrow \mathcal{N}(g_i - g_{i-1} \mid 0, \sigma_{dist}^2(x_i, x_{i-1}))$ 
   $f_{meas}^g \leftarrow \mathcal{N}(g_i \mid y_i, \sigma_{meas}^2)$ 
   $G^g \leftarrow G^g \cup \{f_{link}^g, f_{meas}^g\}$ 
   $V^g \leftarrow SOLVE(G^g, V^g)$ 

```

over the environment. The local representation in the planner and model represent local beliefs and plans about the area immediately surrounding the robot. The global representation describes the belief about the signal strength at locations the robot has received measurements. The posterior distributions over the global signal representation are used as priors in the local signal representation.

Global Factor Graph (G^g, V^g) is seen in fig. 2a where each value $g_n \in V^g$ is connected to at least one measurement factor, $f_{meas} \in G^g$. These measurement factors are unary factors which represent the real world measurements taken by the robot at the location x_n . Each global value is connected to the previous and next global value through a link factor, $f_{link}^g \in G^g$. The link factor connects the previous and current global values, based on distance traveled, while the measurement factor connects the global graph to real world signal measurements. The global graph does not consider obstacles as the robot only moves a small distance between global node generation. This link factor, used in both the global and local factor graphs to connect values which have a spatial relationship, serves a similar purpose to the kernel distance in Gaussian processes. This link factor has an expected mean of 0, with an uncertainty that grows in proportion to the distance the location corresponding to the values are apart $\sigma_{dist}^2(x_i, x_j) = \epsilon + \alpha_{dist} \|x_i - x_j\|$ where ϵ is a small positive number and α_{dist} is an experimentally determined constant. In order to add new measurements to the global factor graph, alg. 1 is used, in which a new global value is created only when the robot has moved sufficiently from the current location. If the robot has not moved far

Algorithm 2: Create Local Graph

Produces a new local graph G^l and local values V^l , given new local information roadmap IRM^l , global graph G^g , global values V^g , global k -nearest neighbors function KNN^g , number of global nodes to link to k_g , and robot location x_t

```

 $V^l, G^l \leftarrow \emptyset, \emptyset$ 
for  $x_i \in \text{nodes}(IRM^l)$  do
   $V^l \leftarrow V^l \cup \{v_i\}$ 
  for  $x_j \in \text{neighbor}(x_i, IRM^l)$  do
     $\sigma_{occ}^2 \leftarrow \sigma_{occ}^2(x_i) + \sigma_{occ}^2(x_j)$ 
     $f_{link}^l \leftarrow \mathcal{N}(l_i - l_j \mid 0, \sigma_{dist}^2(x_i, x_j) + \sigma_{occ}^2)$ 
     $G^l \leftarrow G^l \cup \{f_{link}^l\}$ 
for  $x^g \in KNN^g(x_t, k_g)$  do
   $l \leftarrow CLOSEST(x^g, V^l)$ 
   $\mu(x^g), \sigma^2(x^g) \leftarrow V^g(x^g)$ 
   $f_{global}^l \leftarrow \mathcal{N}(l \mid \mu(x^g), \sigma^2(x^g) + \sigma_{dist}^2(l, x^g))$ 
   $G^l \leftarrow G^l \cup \{f_{global}^l\}$ 

```

enough from the current location, the another measurement factor is added to the previous value.

Local Factor Graph (G^l, V^l) is built from the local information roadmap over the area centered at the robot. The local information roadmap is an $n \times m$ grid which describes the area the robot can traverse in the local policy, described in section IV-A. The local factor graph (fig. 2b) models the local IRM topology by adding link factors between values $l \in V_l$ in the local information roadmap which are connected. The local link factors, $f_{link}^l \in G^l$, are similar to the global link factors as uncertainty increases with the distance between the spatial locations, but also has the addition of occupancy information. We assume that if a location corresponding to a local value is not traversable for the robot, the model estimate of the signal will have higher uncertainty in that area due to potential attenuation or reflection in the case of solid obstacles. To model this we increase the uncertainty proportional to the occupancy probability of that location, according to $\sigma_{occ}^2(i) = \alpha_{occ} p_{occ}(i)$, where α_{occ} is an experimentally chosen constant, set to 0.5 in this work. In order to incorporate the real world measurements into the local factor graph, we find the closest k_g locations corresponding to values in the global factor graph to the robot location. For each location that is close to the robot location, we add a factor, $f_{global}^l \in G^l$, to the closest local value which uses the posterior estimate from the global value as well as the distance between the local and global locations. We assume the distance is small and do not consider obstacles between the global and local nodes. The algorithm for this can be seen in alg. 2, which describes how both the local graph and local values are constructed.

Inference In order for the planner to do active source seeking, the signal belief must be able to efficiently infer a distribution over the signal at unknown spatial locations. The model must also support the ability to condition the model on locations and measurements the robot plans to take, but has not taken yet, termed hypothetical locations \mathbf{X}_h and hypothetical values \mathbf{Y}_h . This allows the planner to make non-myopic plans in which the future values take into

Algorithm 3: Inference at locations

Produces a belief over signal at location x_q given hypothetical future locations \mathbf{X}_h and measurements \mathbf{Y}_h , initially empty local value cache C , local graph G^l and local values V^l .

```

if  $\mathbf{X}_h \in C$  then
  return  $C(\mathbf{X}_h, x_q)$ 
 $G_h^l = G^l$ 
for  $(x, y) \in (\mathbf{X}_h, \mathbf{Y}_h)$  do
   $l \leftarrow CLOSEST(x, V^l)$ 
   $\mu(l), \sigma^2(l) \leftarrow V^l(l)$ 
   $G_h^l \leftarrow G_h^l \cup \{\mathcal{N}(l \mid y, \sigma_i^2)\}$ 
 $V_h^l \leftarrow SOLVE(G_h^l, V^l)$ 
 $C(\mathbf{X}_h) = V_h^l$ 
return  $V_h^l(x_q)$ 

```

account measurements that will be taken in the policy [4]. To this end, alg. 3 describes the algorithm to infer the posterior belief at a spatial location x_q . Typically, the graph cannot be re-solved with every query so the algorithm determines if it has already computed the local values given the current hypothetical locations \mathbf{X}_h by checking if it is in the cache C . The cache C allows the re-use of previously solved values as the posterior distribution over x_q has already been computed when computing a different query location. C is emptied when a new local IRM is received due to the robot moving, or a new measurement is added to the global graph. If algorithm cannot use the cached values, the algorithm adds unary factors representing the conditioning on the hypothetical measurements, using the previous prior uncertainty for the unary factor. The algorithm then re-solves the factor graph for the entire local information roadmap and adds the newly solved factors to the cache. This cache is emptied each time a new local factor graph is constructed.

To produce a belief over locations that are not in the local information roadmap, query locations are added to the global roadmap with distance based link factors and the resulting graph is solved.

C. Hierarchical Solver

Our approach is to compute a local planning policy that solves eq. (2) on W_r^l and W_s^l over a receding horizon t_l , which we refer to as *Local Planning*, while simultaneously solving eq. (2) over the full approximate world representation W_r^g and W_s^g , which we refer to as *Global Planning*.

Local Policy In Receding Horizon Planning (RHP), the objective function in Eq. (2) is modified:

$$\pi_{t:t+T}^*(b) = \operatorname{argmax}_{\pi} \mathbb{E} \left[\sum_{t'=t}^{t+t_l} \gamma^{t'-t} R(b_{t'}, \pi_{t'}(b_{t'})) \right] \quad (3)$$

where t_l is a finite planning horizon for a planning episode at time t . The expected value is taken over future measurements and the reward is discounted by γ [23].

For the local policy, the reward is based on the current signal belief and uses the upper confidence bound (UCB) acquisition function, which is commonly used for informative path planning for Bayesian optimization [4]. The UCB equation is $S^l(x_i) = \mu(x_i) + \beta \sigma^2(x_i)$ where $\mu(x)$ and $\sigma^2(x)$ are the current mean and variance of W_s at location x_i and

β is an empirically determined weight. We use this objective function for local exploration as the robot should investigate areas that the model is uncertain about, but should still prefer high concentration areas. In our experiments we set $\beta = 3$ to encourage exploration.

Global Policy Boundaries in between explored and unexplored space, termed *frontiers*, are encoded in the topological graph of the environment (Global IRM) and represent goal waypoints. By visiting frontiers, the robot takes new signal measurements and uncovers swaths of the unexplored environment. The robot must plan a sequence of frontiers p that it can reach within the time budget that maximizes its expected reward S^g . Equation (2) becomes

$$\pi^* = \operatorname{argmax}_p \sum_{n \in p} F(t_p(n); W_r^g) \cdot S^g(n; W_s^g) \quad (4)$$

where $t_p(n)$ is the time required to reach frontier n through p , F is the *front-loading* function proposed in [24], and where p should verify the time budget constraint. $F(t_p(n); W_r^g)$ is a greedy incentive that encourages gathering reward earlier in time in order to trade off long-term finite-horizon planning with immediate information gain.

We wish to maximize the expected improvement (EI) for signal measurement taken at a frontier n_i (corresponding to location x_i), given current signal belief W_s^g . EI favors actions that offer the best improvement over the current maximal value, with an added exploration term ξ to encourage diverse exploration. The EI objective function is defined according to $S^g(x_i) = I\Phi(Z) + \sigma(x_i)\phi(Z)$ where $Z = \frac{I}{\sigma^2(x_i)}$, $I = \mu(x_i) - \max(\mu(\mathbf{X}_{0:t})) - \xi$, and Φ and ϕ are the CDF and PDF of the normal distribution, respectively [25], [26]. Expected improvement is chosen as the global objective function because the robot should prefer frontiers which have a gain in the expected signal strength, while during local exploration the robot should seek out information rich locations rather than only seeking the maxima.

Policy Selection The local planning policy π^l provides immediate signal reward, and is computed over the local space immediately surrounding the robot, while the global policy π^g is computed over a sparse representation of the entire environment the robot has explored so far. Therefore, the global policy has the potential to bring the robot to valuable locations it has not yet explored, but may also require more travel time. As in [27], we select the planning policy π according to its utility $U(b_t; \pi)$ computed from eq. (2) and the probability $\hat{P}(\pi)$ that the plan will be successfully executed:

$$\pi_t^* = \operatorname{arg} \max_{\pi \in \{\pi_{t:t+T}^g, \pi_{t:t+t_l}^l\}} \hat{P}(\pi)U(b_t; \pi). \quad (5)$$

V. EXPERIMENTS

We demonstrate our system on two simulation environments as well as a hardware trial. The simulated environments were collected in a parking garage with a robot with an attached radio and a radio placed in the environment. In the simulation environments, we first run the robot in a lawnmower path to collect signal strength data and traversability data in the environment. We position the




Fig. 3: Maps of the Signal to Noise Ratio collected in a real parking garage. SNR is colored from blue (low) to yellow (high). Areas in white are not traversable by the robot.

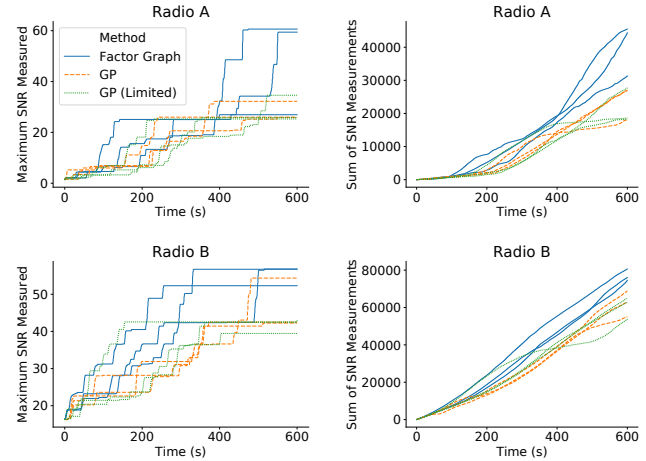


Fig. 4: Simulation experiment results over two real-world environments. Left: the maximal signal reading up until time t . Right: the sum of the measurements taken until time t .

radio in two different locations (fig. 3). Location A is less accessible, so that stronger readings are only seen in direct line of sight, in location B the radio is placed so that it is within line of sight of more areas of the environment. The robot starts in the same location in both environments, and the traversability is the same in both environment.

In the experiments we compare our approach (labeled Factor graph) with a standard Gaussian process using a squared exponential kernel. The kernel hyper-parameters are set to a length scale of 3 and an observation noise of 0.01. This Gaussian process also only adds new measurements when the robot has moved a minimum distance (d_{min}), which was set to 0.25m in our experiments. We also compare against a Gaussian process (GP) which is limited during inference to the same number of measurements (k_g) during inference as the Factor graph, which we term Gaussian process (limited), (GP (limited)). This limited Gaussian process is added to show that the performance of our system is not merely due to limiting the number of measurements included in the model. For the factor graph, we set the measurement noise to 0.01, α_{occ} to 0.5, and α_{dist} to 0.1.

A. Simulation Experiments

We compare our novel Factor graph approach with the two baseline approaches (Gaussian Process and Gaussian process

(limited)) in both environments (location A and location B). To compare against the baselines, we show two metrics for each environment (fig. 4). The first shows the maximal SNR measurement recorded up until time t . It compares how closely the robot approaches the signal source, and how quickly it approaches the signal source. The second metric is the cumulative sum of the SNR measurements the robot has taken until time t . This metric, inspired by the cumulative reward metric [4], shows how the robot makes decisions to find maximal areas over time. For each environment, we run the robot three times and report each trial.

As can be seen in fig. 4, the factor graph model allows the robot to find a much higher signal strength in 2 out of 3 trials, and the performance is more consistent. Limiting the Gaussian process to only 200 measurements does not have a big impact on the performance compared to the regular Gaussian process, but it still does not perform as well as the factor graph model. In the comparison of the cumulative signal measurement the factor graph outperforms the other methods in all 3 trials, with the other two methods performing about as well. This implies the factor graph consistently allows the planner to make better decisions through the entire experiment.

We also present a study of the scalability of the approaches, shown in fig. 5. We find that the Gaussian process has the typical exponential scaling with in the time required for inference as the number of measurements in the model grows. We also find that even limiting the Gaussian process to a fixed number of measurements per inference is still more expensive than the factor graph on average. The factor graph has a bimodal distribution in its performance, when the graph has to be re-solved it is slower than the Gaussian process methods, but due to its ability to cache the model posterior, most of the inference is a hash table lookup.

We believe the consistent performance of the factor graph is due to two reasons. The first is that the factor graph is generally faster for inference than Gaussian process methods, allowing the planner to plan at a higher frequency. The second reason is the ability of the factor graph to represent uncertainty in the signal due to obstacles. This allows the factor graph model to infer that signal strengths on the opposite side of obstacles and around corners are more interesting as the signal uncertainty increases through the obstacle.

B. Hardware Experiments

We demonstrate our system on a real robot locating a radio in an unknown environment. The robot, a Boston Dynamics Spot [28] equipped with custom sensing and computing systems [29]–[31], is initially started at one end of a parking garage. The radio is deployed at another end of the hallway after two lefthand turns. We use the factor graph model as with k_g equal to 100 due to the fact that the computer payload on the robot is limited. As can be seen in fig. 6, the robot explores an open passage before going directly to the source of the radio. We find that our system is able to be easily tuned to the computational limits of the robot and efficiently




Fig. 5: Comparison of time required to perform inference for the three models over the number of measurements in the model. Left: the data binned data with the standard deviation, Right: the raw data.




Fig. 6: Experiment of our system on a real Spot robot. The robot explores a man made structure to find the radio. The collected measurements, colored from blue (low) to yellow (high), as well as the robot's trajectory are shown. Black dots show the non-traversable areas such as walls.

finds the signal source. As with fig. 5 we found that the robot does not increase its planning time drastically when more measurements are added to the model. An extended hardware experiment can be seen in [32].

VI. CONCLUSION

Large scale active source seeking using a robot is a complex and challenging task which requires the ability to scale the planner and model. In this work we have demonstrated the ability of our factor graph based model to decompose the problem into local and global tasks which allows the model to scale easily with a large number of measurements. We have also demonstrated the model's ability to handle the propagation of signals through unknown obstacles by modeling the explicit uncertainty due to these obstacles. Through simulation experiments we have shown that the ability to plan rapidly due to the inference speed of the model as well as the ability to represent signal uncertainty due to obstacles allow our overall system to perform active source seeking better than the baseline methods. We have also shown the ability of our model to be scaled to the computational needs of a real life robot and demonstrated our system in a real life unknown environment.

REFERENCES

- [1] S. Kemna, J. G. Rogers, C. Nieto-Granda, S. Young, and G. S. Sukhatme, "Multi-robot coordination through dynamic voronoi partitioning for informative adaptive sampling in communication-constrained environments," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 2124–2130.
- [2] S. McCammon, G. Marcon dos Santos, M. Frantz, et al., "Ocean front detection and tracking using a team of heterogeneous marine vehicles," *Journal of Field Robotics*, vol. 38, no. 6, pp. 854–881, 2021. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/rob.22014>.
- [3] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT press, 2006.
- [4] R. Marchant, F. Ramos, and S. Sanner, "Sequential Bayesian optimisation for spatial-temporal monitoring," in *Conference on Uncertainty in Artificial Intelligence (UAI)*, ser. UAI'14, Arlington, Virginia, USA: AUAI Press, Jul. 2014, pp. 553–562.
- [5] S.-K. Kim*, A. Bouman*, G. Salhotra, et al., "PLGRIM: Hierarchical value learning for large-scale exploration in unknown environments," in *International Conference on Automated Planning and Scheduling (ICAPS)*, vol. 31, 2021, pp. 652–662.
- [6] F. Mascarich, T. Wilson, C. Papachristos, and K. Alexis, "Radiation source localization in gps-denied environments using aerial robots," in *IEEE International Conference on Robotics and Automation (ICRA)*, Brisbane, Australia: IEEE Press, 2018, pp. 6537–6544.
- [7] Y. Sung and P. Tokekar, "A competitive algorithm for online multi-robot exploration of a translating plume," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 3391–3397.
- [8] J. Fink and V. Kumar, "Online methods for radio signal mapping with mobile robots," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2010, pp. 1940–1945.
- [9] B. Charrow, N. Michael, and V. R. Kumar, "Cooperative multi-robot estimation and control for radio source localization," *International Journal of Robotics Research*, vol. 33, pp. 569–580, 2012.
- [10] C. Guestrin, A. Krause, and A. P. Singh, "Near-Optimal Sensor Placements in Gaussian Processes," Tech. Rep., 2005.
- [11] T. O. Fossum, J. Eidsvik, I. Ellingsen, et al., "Information-driven robotic sampling in the coastal ocean," *Journal of Field Robotics*, vol. 35, no. 7, pp. 1101–1121, 2018. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/rob.21805>.
- [12] I. M. Rayas Fernández, C. E. Denniston, D. A. Caron, and G. S. Sukhatme, "Informative path planning to estimate quantiles for environmental analysis," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10 280–10 287, 2022.
- [13] G. Salhotra, C. E. Denniston, D. A. Caron, and G. S. Sukhatme, "Adaptive sampling using pomdps with domain-specific considerations," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 2385–2391.
- [14] L. Dressel and M. J. Kochenderfer, "Efficient decision-theoretic target localization," in *International Conference on Automated Planning and Scheduling (ICAPS)*, 2017.
- [15] L. K. Dressel and M. J. Kochenderfer, "On the optimality of ergodic trajectories for information gathering tasks," in *American Control Conference (ACC)*, IEEE, 2018, pp. 1855–1861.
- [16] G. Mathew and I. Mezić, "Metrics for ergodicity and design of ergodic dynamics for multi-agent systems," *Physica D: Nonlinear Phenomena*, vol. 240, no. 4-5, pp. 432–442, 2011.
- [17] L. Miller, "Optimal ergodic control for active search and information acquisition," Ph.D. dissertation, Northwestern University, 2015.
- [18] L. K. Dressel, *Efficient and Low-cost Localization of Radio Sources with an Autonomous Drone*. Stanford University, 2018.
- [19] F. Dellaert, "Factor Graphs and GTSAM: A Hands-on Introduction," Tech. Rep., 2012.
- [20] Y. Chang, K. Ebadi, C. E. Denniston, et al., "Lamp 2.0: A robust multi-robot slam system for operation in challenging large-scale underground environments," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9175–9182, 2022.
- [21] J. G. Monroy, J.-L. Blanco, and J. Gonzalez-Jimenez, "Time-variant gas distribution mapping with obstacle information," en, *Autonomous Robots*, vol. 40, no. 1, pp. 1–16, Jan. 2016.
- [22] A. Gongora, J. Monroy, and J. Gonzalez-Jimenez, "Joint estimation of gas and wind maps for fast-response applications," *Applied Mathematical Modelling*, vol. 87, pp. 655–674, 2020.
- [23] A. Bouman, J. Ott, S.-K. Kim, et al., "Adaptive coverage path planning for efficient exploration of unknown environments," in *IROS*, 2022.
- [24] O. Peltzer, A. Bouman, S.-K. Kim, et al., "FIG-OP: Exploring Large-Scale Unknown Environments on a Fixed Time Budget," in *IROS*, 2022.
- [25] D. R. Jones, M. Schonlau, and W. J. Welch, "Efficient Global Optimization of Expensive Black-Box Functions," *Journal of Global Optimization*, 1998.
- [26] C. Qin, D. Klabjan, and D. Russo, "Improving the expected improvement algorithm," in *NeurIPS*, Long Beach, California, USA, 2017.
- [27] J. Ott, S.-K. Kim, A. Bouman, et al., *Risk-aware meta-level decision making for exploration under uncertainty*, 2022. arXiv: 2209.05580.
- [28] *Spot@ - The Agile Mobile Robot*, en, <https://www.bostondynamics.com/products/spot>.
- [29] A. Agha, K. Otsu, B. Morrell, et al., "Nebula: Quest for robotic autonomy in challenging environments; team costar at the darpa subterranean challenge," *Journal of Field Robotics*, 2021.
- [30] K. Otsu, S. Tepsuporn, R. Thakker, et al., "Supervised autonomy for communication-degraded subterranean exploration by a robot team," in *IEEE Aerospace Conference*, 2020.
- [31] A. Bouman*, M. Ginting*, N. Alatur*, et al., "Autonomous Spot: Long-Range Autonomous Exploration of Extreme Environments with Legged Locomotion," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- [32] C. E. Denniston, O. Peltzer, J. Ott, et al., "Fast and scalable signal inference for active robotic source seeking," arXiv, 2023.