

Sequential Bayesian Optimization for Adaptive Informative Path Planning with Multimodal Sensing

Joshua Ott¹, Edward Balaban², and Mykel J. Kochenderfer¹

Abstract—Adaptive Informative Path Planning with Multimodal Sensing (AIPPMS) considers the problem of an agent equipped with multiple sensors, each with different sensing accuracy and energy costs. The agent’s goal is to explore the environment and gather information subject to its resource constraints in unknown, partially observable environments. Previous work has focused on the less general Adaptive Informative Path Planning (AIPP) problem, which considers only the effect of the agent’s movement on received observations. The AIPPMS problem adds additional complexity by requiring that the agent reasons jointly about the effects of sensing and movement while balancing resource constraints with information objectives. We formulate the AIPPMS problem as a belief Markov decision process with Gaussian process beliefs and solve it using a sequential Bayesian optimization approach with online planning. Our approach consistently outperforms previous AIPPMS solutions by more than doubling the average reward received in almost every experiment while also reducing the root-mean-square error in the environment belief by 50%. We completely open-source our implementation to aid in further development and comparison.†

I. INTRODUCTION

In many real-world applications such as planetary exploration, search-and-rescue, and environmental monitoring, an autonomous agent must explore an unknown environment with partial observability. The agent is equipped with multiple sensors that allow it to take noisy observations of the environment. In addition, the agent has limited time and energy resources requiring that it balance exploring the environment and exploiting its knowledge of the environment based on its current belief. This problem is referred to as Adaptive Informative Path Planning with Multimodal Sensing (AIPPMS) [1]. The objective of the AIPPMS problem is to produce an adaptive decision making strategy to guide the agent to the goal location while balancing exploration with exploitation and respecting the resource constraints.

Many real-world scenarios involve agents equipped with multiple sensors that cannot be simultaneously active due to energy, sensor interference, or movement constraints. Consider an agent with multimodal sensing capabilities. In this context, multimodal refers to the existence of multiple sensors with potentially different levels of accuracy and cost associated with each sensor; for example, a rover equipped with a suite of sensors (e.g., CCD cameras, radar, mass spectrometers, and drills). Some of these sensors, such as the cameras, require small amounts of energy but produce less reliable measurements because they are unable




Fig. 1: The belief MDP framework where the belief-state is a Gaussian process. Upon taking an action, the agent receives a noisy observation of the environment, updates its belief-state, and receives a reward based on the belief-state and action taken. Sequential Bayesian optimization aims to find the best place to take samples of an unknown function accounting for the sequential constraints associated with the agent’s movement and energy costs. For clarity, we illustrate a one-dimensional Gaussian process, but our method is not restricted to one dimension.

to fully observe the subsurface environment. Conversely, other sensors, such as the drill, require very large amounts of energy but produce almost perfect measurements of the subsurface environment. The same trade-off between sensor efficiency and sensor accuracy can be found in many other applications such as ocean monitoring, mineral prospecting, aerial surveillance, and wildfire mapping.

The informative path planning (IPP) problem is NP hard [2] and involves an agent planning a path without accounting for noisy observations of the environment. The adaptive problem extends the IPP problem by seeking a policy that changes in response to received observations [3–7]. The AIPPMS problem further extends the adaptive problem by reasoning over multiple sensing modalities with different sensing costs corresponding to sensor accuracy and movement costs [1].

Both the adaptive informative path planning (AIPP) and non-adaptive IPP problems have been widely studied and a variety of approaches have been proposed [8–13]. Most of the previous work has focused on the case where the agent is equipped with a single sensor. In this work, we are interested in the case where the agent can choose between multiple sensing modalities. This extension adds additional complexity by requiring that the agent reasons about the cost-benefit trade-off associated with the different sensing modalities. Choudhury et al. considered multiple sensing modalities [1]. We build upon this work by formulating the problem as a belief Markov decision process (MDP) as shown in Fig. 1. We solve the AIPPMS problem through a sequential Bayesian optimization approach using Monte Carlo tree search with Double Progressive Widening (MCTS-

¹Department of Aeronautics & Astronautics, Stanford University ({joshuaott, mykel}@stanford.edu).

²NASA Ames Research Center (edward.balaban@nasa.gov).

†https://github.com/sisl/SBO_AIPPMS

DPW) and belief-dependent rewards [14–16]. We directly compare our results with those of Choudhury et al. and show a significant improvement in performance by taking sequential actions to maximize the approximate mutual information based on the current Gaussian process belief.

The key contributions of this work are:

- 1) The formulation of the AIPPMS problem as a belief MDP where the world belief-state is represented as a Gaussian process.
- 2) The extensive evaluation of our method in simulation, along with directly comparing to previous AIPPMS methods. We compare our results on the Information Search RockSample problem introduced by He et al. [17] and adapted by Choudhury et al. [1]. We also introduce a new benchmark for AIPPMS problems: the Rover Exploration problem.
- 3) The release of our implementation as an open-source software package for use and further development by the community.

The remainder of this paper is organized as follows. Section II discusses related work. Section III provides preliminary background on Gaussian processes and partially observable Markov decision processes (POMDP). Section IV formulates the AIPPMS problem as a belief MDP and describes our approach for solving the problem. Section V presents our results on two AIPPMS problems and compares our method to Choudhury et al. and other baseline methods [1]. Section VI concludes our work.

II. RELATED WORK

A common approach for solving the AIPP problem is to represent the world belief as a Gaussian process. Gaussian processes are distributions over functions and are able to represent spatial correlation among observations while quantifying uncertainty in the underlying function [18]. Marchant et al. formulate the AIPP problem as a POMDP and solve it using sequential Bayesian optimization through MCTS Upper Confidence bound for Trees (MCTS-UCT) [19]. Extensions of this work include modifying the reward function to achieve a monitoring behavior that exploits high-gradient areas [20] and further reasoning over continuous action spaces through Bayesian optimization [12]. Rayas Fernandez et al. use Partially Observable Monte-Carlo Planning (POMCP) with Gaussian process beliefs to estimate quantiles of the underlying world state [21]. Their approach suggests sample locations for a team of scientists to visit once it has completed its exploration of the environment.

More recent methods have used Gaussian Markov random fields (GMRF) to model the world belief. GMRFs exploit the spatial Markov property to enforce conditional independence of non-neighboring lattice nodes in the world belief [10]. This approach avoids the $\mathcal{O}(n^3)$ inference requirement that is common among Gaussian processes, but introduces additional assumptions through the conditional independence of non-neighboring nodes. Rückin et al. combines tree search with an offline-learned neural network predicting informative sensing actions [13].

In these works, the only decision is the next location to visit, and energy is only depleted by movement. Choudhury et al. introduced the AIPPMS problem where the agent must make decisions about where to visit and what sensors to use [1]. Greater energy cost is incurred for using more accurate sensors. Choudhury et al. formulate the AIPPMS problem as a POMDP and use POMCP to produce a solution.

We extend this work by formulating the AIPPMS problem as a belief MDP and use belief-dependent rewards to penalize uncertainty in the belief through sequential Bayesian optimization-based rollouts with progressive widening.

III. BACKGROUND

A. Gaussian Processes

A Gaussian process is a distribution over functions that can be used to predict an underlying function f given some previously observed noisy measurements at location \mathbf{x} . That is, $y = f(\mathbf{x}) + z$ where f is deterministic but z is zero-mean Gaussian noise, $z \sim \mathcal{N}(0, \nu)$. The distribution of functions conditioned on the previously observed measurements $\hat{\mathbf{y}} \mid \mathbf{y}, \nu \sim \mathcal{N}(\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*)$ is given by:

$$\begin{aligned} \boldsymbol{\mu}^* &= \mathbf{m}(X^*) + \mathbf{K}(X^*, X)(\mathbf{K}(X, X) + \nu \mathbf{I})^{-1}(\mathbf{y} - \mathbf{m}(X)) \\ \boldsymbol{\Sigma}^* &= \mathbf{K}(X^*, X^*) - \mathbf{K}(X^*, X)(\mathbf{K}(X, X) + \nu \mathbf{I})^{-1}\mathbf{K}(X, X^*) \end{aligned} \quad (1)$$

where X is the set of previously measured locations, X^* is the set of locations we wish to predict the values $\hat{\mathbf{y}}$, $\mathbf{m}(X)$ is the mean function and $\mathbf{K}(X, X')$ is the covariance matrix constructed with the kernel $k(\mathbf{x}, \mathbf{x}')$ [18]. It is important to note the roll of ν in multimodal sensing problems. Larger values of ν correspond to less reliable, and less costly, measurements while smaller values of ν correspond to more reliable, and more costly, measurements.

B. POMDPs

A POMDP is a problem formulation for sequential decision making under uncertainty and is represented as a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{O}, T, \mathcal{Z}, R, \gamma)$, where \mathcal{S} , \mathcal{A} , and \mathcal{O} denote the state, action, and observation spaces, respectively. At each time step, the agent takes an action $a \in \mathcal{A}$ to move from state $s \in \mathcal{S}$ to $s' \in \mathcal{S}$ and then receives an observation $o \in \mathcal{O}$. The transition function $T(s, a, s') = p(s' \mid s, a)$ specifies the probability of transitioning to s' given that action a was taken from state s . Similarly, the observation function $Z(s', a, o) = p(o \mid s', a)$ specifies the probability of observing o given that the agent took action a and transitioned to state s' . The belief over system states, $b(s)$, is a probability distribution over \mathcal{S} . The agent receives a reward specified by $R(s, a)$ for taking action a in state s .

C. Belief MDP

Any POMDP can be viewed as an MDP that uses beliefs as states, also known as a belief MDP [14]. The state space of a belief MDP is the set of all beliefs and is therefore continuous. The action space is identical to that of the POMDP. The reward function for a belief MDP is dependent on the belief and action taken. Formulating the problem as a belief MDP allows one to directly access the belief-state in the reward function which is useful in information-gathering tasks.

IV. METHODS

A. Belief MDP Formulation

We formulate and solve the AIPPMS problem as a belief MDP in order to jointly reason about the effect of movement and multimodal sensing actions. The world environment is represented by a location graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where the nodes \mathcal{V} correspond to the set of all locations the agent can visit. The agent starts at v_s and needs to reach a specified goal v_g while remaining within its total cost budget B . The agent is equipped with a set of sensors with each sensor having a different cost and level of accuracy. We use notation consistent with previous work on the AIPP problem [1, 3].

Belief State Space \mathcal{B} : The belief-state at time t is defined as $b_t = (v_t, \Delta e_t, GP(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')))$ where v_t is the agent's current location and $\Delta e_t = B - \sum_{i=1}^{t-1} C(b_i, a_i)$ is the remainder of the cost budget. The cost associated with taking action a_i from belief-state b_i is $C(b_i, a_i)$. Different costs are incurred for movement and sensing actions. Similarly, selecting a sensing action with greater accuracy will incur more cost than a sensing action with less accuracy. Note that the cost budget is different from the reward and is meant to track the time and energy resources of the agent. The problem terminates when the agent has too few resources left to take any actions, which occurs when $\Delta e < \min_a C(b_t, a)$. If the agent is not at the goal location v_g when this happens it receives $-\infty$ reward. The belief over the world state $\mathcal{X}_{\mathcal{V}}$ is represented as a Gaussian process where $\mathcal{X}_{\mathcal{V}} \sim GP(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$.

Action Space \mathcal{A} : The action space is $\text{neigh}(v_t) \cup S$ where $\text{neigh}(v_t)$ is all of the neighboring locations to v_t and S is the space of sensing actions.

Transition Model T : The transition function is fully deterministic for v_t and Δe_t . If a_t is a movement action to a neighboring node, then $v_{t+1} = a_t$ and $\Delta e_{t+1} = \Delta e_t - C(b_t, a_{t+1})$. If a_t is a sensing action, then $v_{t+1} = v_t$ and the remaining cost budget will change accordingly. The transition for the world belief involves an update of the current belief-state based on the result of any sampling actions taken. If a sampling action is taken, a measurement y_t will be received at location v_t and the Gaussian process update equations provided in Eq. (1) can be used to obtain the posterior distribution $\hat{\mathbf{y}} \mid \mathbf{y}_t$ where v_t is appended to the set of previously measured locations M_t , and y_t is appended to the set of previously received measurements \mathbf{y}_t .

Reward Function R : The reward function is made up of two key elements. The first is the expected reward of interacting with the environment based on our current belief: $\sum_s R(s, a)b(s)$. The second is based on the goal of minimizing the uncertainty in the world belief. To do so, we use the mutual information. We have that $\mathcal{X}_{\mathcal{V}}$ refers to the set of random variables at all locations in the environment and M_t contains all of the locations that have been observed by the agent at time t . The mutual information between the locations M_t and the rest of the world $\mathcal{X}_{\mathcal{V}}$ given previous measurement locations M_{t-1} is:

$$\begin{aligned} I(\mathcal{X}_{M_t}; \mathcal{X}_{\mathcal{V} \setminus M_{t-1}}) &= \mathcal{H}(\mathcal{X}_{\mathcal{V} \setminus M_{t-1}}) - \mathcal{H}(\mathcal{X}_{\mathcal{V} \setminus M_{t-1}} \mid \mathcal{X}_{M_t}) \quad (2) \\ &= \mathcal{H}(\mathcal{X}_{\mathcal{V}} \mid \mathcal{X}_{M_{t-1}}) - \mathcal{H}(\mathcal{X}_{\mathcal{V}} \mid \mathcal{X}_{M_t}) \quad (3) \end{aligned}$$

where $\mathcal{H}(\mathcal{X}_{\mathcal{V}} \mid \mathcal{X}_{M_t})$ is the conditional entropy. For a Gaussian process with D observations, the conditional entropy is

$$\mathcal{H}(\mathcal{X}_{\mathcal{V}} \mid \mathcal{X}_{M_t}) = \frac{1}{2} \log |\Sigma_{M_t}^*| + \frac{D}{2} (1 + \log(2\pi)). \quad (4)$$

Therefore, Eq. (3) can be rewritten as

$$I(\mathcal{X}_{M_t}; \mathcal{X}_{\mathcal{V} \setminus M_{t-1}}) = \frac{1}{2} \log |\Sigma_{M_{t-1}}^*| - \frac{1}{2} \log |\Sigma_{M_t}^*| \quad (5)$$

where we drop the constant term from Eq. (4). Mutual information was proposed as a sampling criterion by Caselton and Zidek [22]. It seeks to find sensor placements that are most informative about unsensed locations. This optimization criterion directly measures the effect of sensor placements on the posterior uncertainty of the Gaussian process. To avoid computing the determinant at every step during the rollout, we approximate Eq. (5) with

$$I(\mathcal{X}_{M_t}; \mathcal{X}_{\mathcal{V} \setminus M_{t-1}}) \approx \text{Tr}(\Sigma_{M_{t-1}}^*) - \text{Tr}(\Sigma_{M_t}^*), \quad (6)$$

where the trace of the covariance matrix corresponds to the total variance in the world belief-state. As shown by Krause et al., minimizing the conditional entropy in Eq. (4) is equivalent to sampling at locations of highest entropy. For a Gaussian process, this is equivalent to sampling at locations with the greatest predictive variance [23]. Similarly, sampling at locations with the greatest predictive variance is equivalent to sampling at locations that decrease the predictive variance the most. The predictive variance at each location in $\mathcal{X}_{\mathcal{V}}$ is given by the diagonals of the posterior covariance matrix Σ^* . Therefore, we seek to take the sequence of actions that maximally reduce the predictive variance of the world belief-state, and thereby approximately maximize the mutual information between the new sample locations and the rest of the world belief-state given the previous measurement locations. The reward function is then given by:

$$R(b, a, b') = \sum_s R(s, a)b(s) + \lambda I(\mathcal{X}_{M_t}; \mathcal{X}_{\mathcal{V} \setminus M_{t-1}}) \quad (7)$$

where λ weights the mutual information uncertainty reduction against the expected reward.

B. Online Sequential Bayesian Optimization Solution

As opposed to classical Bayesian optimization methods that greedily reason about the next sample location, we use a tree search approach to reason about the agent's actions several steps into the future. To solve the belief MDP, we use MCTS-DPW, which accommodates continuous state and action spaces. MCTS-DPW with belief-states is very similar to the PFT-DPW algorithm described by Sunberg and Kochenderfer [24]. The main difference being they use a particle representation of the belief and apply particle filter belief updates in the simulate function of the algorithm. Marchant et al. outlined a sequential Bayesian optimization

Algorithm 1 Monte Carlo Tree Search with Double Progressive Widening

```

function PLAN( $b$ )
  for  $i \in 1 : n$ 
    SIMULATE( $b, d_{\max}$ )
  return  $\underset{a}{\operatorname{argmax}} Q(ba)$ 

function ACTIONPROGWIDEN( $b$ )
  if  $|C(b)| \leq k_a N(b)^{\alpha_a}$ 
     $a \leftarrow \operatorname{NEXTACTION}(b)$ 
     $C(b) \leftarrow C(b) \cup \{a\}$ 
    return  $\underset{a \in C(b)}{\operatorname{argmax}} Q(ba) + c \sqrt{\frac{\log N(b)}{N(ba)}}$ 
    ▷ feasible actions only
  else
     $(b', r) \leftarrow \text{sample uniformly from } C(ba)$ 
     $q \leftarrow r + \gamma \operatorname{SIMULATE}(b', d - 1)$ 
     $N(b) \leftarrow N(b) + 1$ 
     $N(ba) \leftarrow N(ba) + 1$ 
     $Q(ba) \leftarrow Q(ba) + \frac{q - Q(ba)}{N(ba)}$ 
  return  $q$ 

```

approach using MCTS but did not make use of progressive widening for continuous state and action spaces [19].

MCTS is a widely studied algorithm for online decision making [25]. MCTS constructs a tree of state and action nodes using a generative model G and estimates the state-action value function $Q(ba)$ at each of the action nodes. Nodes are expanded using the upper confidence bound $UCB(b, a) = Q(ba) + c \sqrt{\frac{\log N(b)}{N(ba)}}$ where $N(ba)$ is the number of times the action node has been visited, $N(b) = \sum_{a \in \mathcal{A}} N(ba)$, and c is a parameter that determines the weight given to exploration in the tree.

When the state and action spaces are large or continuous, MCTS will produce trees that are very shallow. With MCTS-DPW, the number of new children sampled from any node in the search tree is limited by DPW using the parameters k_s , α_s , k_a , and α_a . The number of children of a node in the tree is limited to kN^α where N is the number of times the node has been visited and k and α are hyperparameters [26]. When the number of state nodes is greater than kN^α , one of the previously generated states is chosen instead of simulating a new state transition. The belief-state implementation of MCTS-DPW is provided in Algorithm 1 following the notation of Sunberg and Kochenderfer, where G_{GP} denotes a Gaussian process belief update [24].

Similar to Choudhury et al., we apply action pruning so that we only consider feasible actions during the tree search. A feasible action is one that transitions the agent to a feasible state where a feasible state is a state in which the agent is still able to reach the goal within the specified resource constraints. That is, belief-state b_t is feasible if $\Delta e_t > C_G(v_t, v_g)$ where $C_G(v_t, v_g)$ is the cost along the shortest path from the agents current location to the goal location. We use the POMDPs.jl framework to implement the preceding algorithms in our experiments [27].

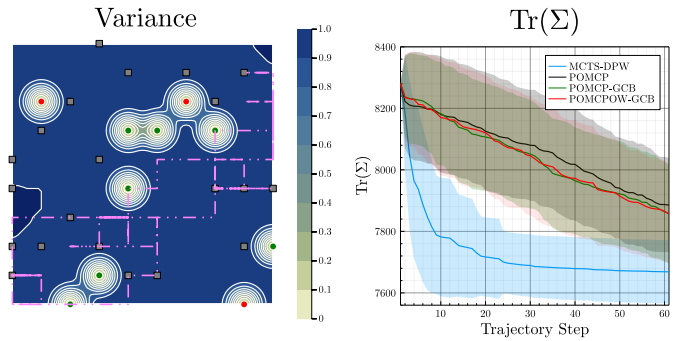


Fig. 2: Results from the ISRS problem. The left image shows the posterior variance of the Gaussian process belief using MCTS-DPW. Bad rocks are shown in red, good rocks in green, and beacons in gray. The agent’s trajectory is shown in pink. The right plot shows $\operatorname{Tr}(\Sigma)$ along the agent’s trajectory. Note that for even comparison we evaluated $\operatorname{Tr}(\Sigma)$ using the same Gaussian process setup, even though MCTS-DPW was the only method to use Gaussian process beliefs online.

V. RESULTS

A. Information Search RockSample

We first evaluate our method on the Information Search RockSample (ISRS) problem introduced by He et al. [17] and adapted by Choudhury et al. [1]. ISRS is a variation of the classic RockSample problem. The agent must move through an environment represented as an $n \times n$ grid. Scattered throughout the environment are k rocks with at most one rock in each grid cell. Only some of the rocks are considered to be ‘good,’ meaning that they have some scientific value. The agent receives +10 reward for visiting a good rock and −10 reward for visiting a bad rock. Once a good rock is visited, it becomes bad. The positions of the agent and rocks are known apriori, but visiting a rock is the only way to reveals its true state.

There are also b beacons scattered throughout the environment. Upon reaching a beacon location, the agent has the option to take a sensing action where it will receive observations about the state of the nearby rocks. The fidelity of the observation decreases with increasing distance from the beacon location. There are multiple sensing modalities available to the agent with a higher cost for choosing the more accurate sensing modality. Moving between adjacent cells also expends energy cost. The agent’s goal is to visit as many good rocks and as few bad rocks as possible while returning to the origin without exceeding its resource budget.

Similar to the approach taken by Choudhury et al., we focus on the relationship between the number of rocks k , the number of beacons b , and the relative proportion of good rocks, based on the independent Bernoulli probability p of a rock being good. We vary these three parameters in our experiments while keeping the others fixed. We directly compare our approach with the POMCP and POMCP-GCB approach used by Choudhury et al. [1]. In addition, we compare against the POMCPOW algorithm introduced by Sunberg and Kochenderfer [24], where we use the Generalized Cost Benefit (GCB) algorithm in the rollouts. In the remainder of the paper, we refer to this algorithm as POMCPOW-GCB. Choudhury et al. formulates the problem as a POMDP with discrete beliefs over whether a rock is good or bad and updates the belief as observations are

		$p = 0.5$				$p = 0.75$				$p = 1.0$			
k	b	POMCP	POMCP-GCB	POMCPOW-GCB	MCTS-DPW	POMCP	POMCP-GCB	POMCPOW-GCB	MCTS-DPW	POMCP	POMCP-GCB	POMCPOW-GCB	MCTS-DPW
10	10	-24.0	-18.0	-20.2	5.8	-1.6	4.0	11.0	13.0	31.4	35.8	42.8	26.2
10	25	-26.0	-16.4	-20.4	12.2	-1.4	5.8	7.4	21.6	32.8	39.2	38.0	33.6
25	10	-60.0	-54.2	-61.0	15.2	9.8	17.8	13.6	30.4	80.6	95.0	77.2	61.4
25	25	-61.4	-40.4	-47.8	18.0	8.4	15.0	9.2	41.2	93.4	100.8	89.0	73.4

TABLE I: Comparison of the three different methods as the distribution of good rocks, number of rocks, and number of beacons are varied. Each reported value is the average of 50 different simulation runs. Note that we report different values from Choudhury et al. since we have included a -10 reward for visiting a bad rock whereas they do not.

received. They solve the problem using POMCP with random rollouts and a variation of POMCP with a GCB rollout policy. The main difference is that we formulate the problem as a belief MDP, represent the belief over the world state as a Gaussian process, and use the mutual information $I(\mathcal{X}_{M_t}; \mathcal{X}_{V|M_{t-1}})$ in our MCTS-DPW solver. We refer the reader to Choudhury et al. for further details [1].

Fig. 2 shows the variance of the world belief after the agent has finished exploring the environment. The figure also shows $\text{Tr}(\Sigma)$ as the agent moves along its trajectory. Table I compares the results as the k , b , and p parameters are varied across 50 different simulation runs. We see that the belief MDP formulation with Gaussian process beliefs and MCTS-DPW is able to significantly outperform POMCP, POMCP-GCB, and POMCPOW-GCB in situations where there are both good and bad rocks present. Intuitively, this makes sense because MCTS-DPW is able to take information-gathering actions to reduce the uncertainty in the world belief-state. This encourages the agent to take sensing actions to become more certain about the true value of the rocks before sampling them. POMCP does not explicitly reason about information-gathering actions. Instead, as the random rollouts are performed and more sensing actions are taken, the agent will become more certain about the state of the environment, thereby indirectly benefiting from sensing. With the GCB rollout policy, sensing actions are incentivized to maximize in expectation the mode of the belief-state [1]. This incentive leads the agent to take actions that have a greater impact on the difference of the belief mode; however, it is not as robust as reasoning about the mutual information through sequential Bayesian optimization as shown in Table 1. Other work by Fischer and Tas has observed similar performance when information-gathering actions are a key to successful execution [28].

When only good rocks are present ($p = 1.0$), MCTS-DPW does not outperform the other methods. This is because MCTS-DPW still initially prioritizes sensing actions to reduce its uncertainty in the environment belief, whereas the other methods greedily head toward the rocks with fewer sensing actions. Intuitively, this makes sense because sensing actions are not as valuable in homogeneous environments where POMCP, POMCP-GCB, and POMCPOW-GCB learn that they will receive reward for visiting any rock location.

B. Rover Exploration

We now introduce a new AIPPMS benchmark problem known as the Rover Exploration problem, which is inspired by multiple planetary rover exploration missions [29–31]. The rover begins at a specified starting location and has a fixed amount of energy available to explore the environment

and reach the goal location. The rover is equipped with a spectrometer and a drill. Drilling reveals the true state of the environment at the location the drill sample was taken and is a more costly action to take from a resource budget perspective. Conversely, the spectrometer provides a noisy observation of the environment and uses less of the resource budget. At each step, the rover can decide whether or not it wants to drill. The rover’s goal is to collect as many unique samples as it can while respecting its energy constraints. The rover receives $+1$ reward for drilling a sample that it has not yet seen and -1 reward for drilling a sample that it has already previously collected. This AIPPMS problem highlights the importance of taking sensing actions to decrease the uncertainty in the belief-state before deciding to drill.

For the Rover Exploration problem, we focus on the interplay between the resource budget allotted to the rover and the sensing quality of the spectrometer, where σ_s denotes the standard deviation of a Gaussian sensor model. We also include a raster policy that attempts to fully sweep the environment in a raster pattern and evenly distribute its drilling actions along the way. However, the raster policy may not always be able to make it to the goal location in the specified budget in which case the rover receives $-\infty$ reward for a mission failure.

The environment is modeled as an $n \times n$ grid with β unique measurement types with values between 0 and 1. To construct a spatially-correlated environment, we start by sampling each grid cell value from a uniform distribution of the β unique measurement types. Then, each cell takes on the average value of its neighboring cells with probability p_g . This process creates environments with spatial correlation as well as some noise to simulate a realistic geospatial environment where the amount of spatial correlation is controlled by p_g .

To implement the approach from Choudhury et al. on this benchmark problem, we maintain a discrete belief at each of the grid cells corresponding to the β unique measurement types. When a measurement is received, the weights corresponding to the discrete values are updated based on a Gaussian sensor model.

Fig. 3 shows an example of a rover trajectory with the posterior mean and variance of the Gaussian process, as well as a comparison of $\text{Tr}(\Sigma)$ and the RMSE along the trajectory for the different methods. Table II summarizes the results as the resource budget and spectrometer sensing quality are varied. We highlight three key observations from the results in Fig. 3 and Table II.

First, as expected, MCTS-DPW is able to find more unique drill locations as the resource budget is increased. As the spectrometer noise is increased, less reward is received due

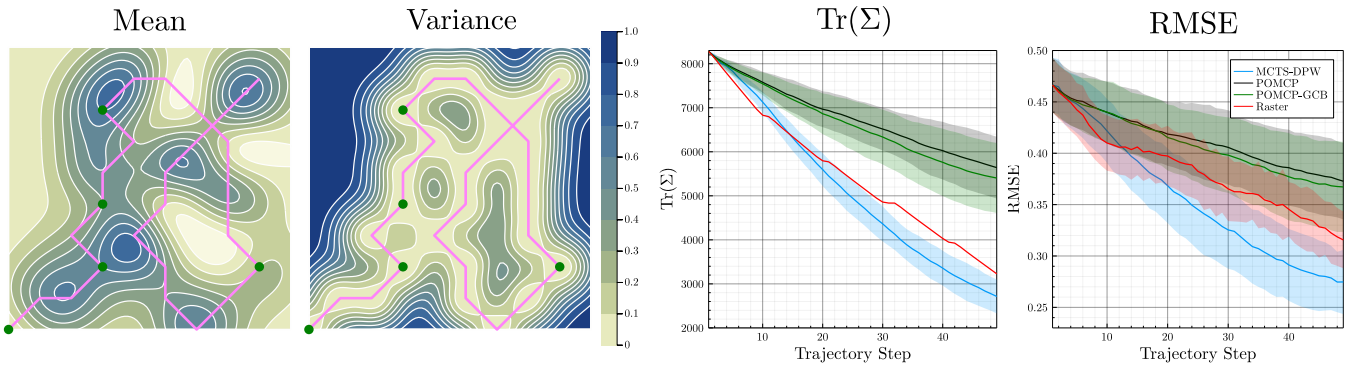


Fig. 3: From left to right: contour plot showing the posterior mean and variance of the Gaussian process belief after the agent has reached the goal state. The agent starts in the bottom left corner and the resource budget was set to 50. The pink line indicates the path the agent took and the green dots represent successful drill sites. The two plots on the right show the trace of the Gaussian process covariance matrix and the RMSE of the beliefs with the true map. Note that for even comparison we evaluated $\text{Tr}(\Sigma)$ and the RMSE all using the same Gaussian process setup, even though MCTS-DPW was the only method to use Gaussian process beliefs online. The plots show the average and standard deviation from 50 different simulation runs with the resource budget set to 100.

Budget	$\sigma_s = 0.1$				$\sigma_s = 0.5$				$\sigma_s = 1.0$			
	POMCP	POMCP-GCB	MCTS-DPW	Raster	POMCP	POMCP-GCB	MCTS-DPW	Raster	POMCP	POMCP-GCB	MCTS-DPW	Raster
30	0.80	0.96	3.42	–	0.62	0.42	1.88	–	0.64	0.70	1.54	–
60	–0.04	1.36	4.84	–	0.22	0.06	2.56	–	0.18	0.48	1.78	–
100	–0.50	1.12	5.60	2.24	0.32	–1.50	2.80	2.24	–0.68	–0.06	1.90	2.24

TABLE II: Comparison of the four methods as the resource budget and spectrometer measurement noise are varied. Each reported value is the average of 50 simulation runs.

to MCTS-DPW having less confidence in its belief about the environment. However, for POMCP and POMCP-GCB this is not the case. As the resource budget increases, we generally see that both POMCP and POMCP-GCB receive less reward since the agent has more opportunities to make mistakes by receiving a sample type it has already seen. MCTS-DPW is able to avoid this by effectively balancing exploration and exploitation by reasoning about its confidence in a particular region before drilling, leading it to make more risk-aware decisions. As the agent has drilled in more locations it will become more conservative in drilling in future locations unless it is very certain that the sample type is unique. Additionally, the Gaussian process belief enables reasoning about spatial correlation in the environment whereas the discrete belief supports reasoning only about the specific cell it has observed, but not about the relationships with other neighboring cells. This is one of the key reasons Gaussian processes are used for geospatial modeling.

Second, under significant amounts of noise ($\sigma_s = 1.0$), the deterministic raster policy is able to outperform MCTS-DPW. The raster policy simply tries to fully cover the space and evenly distribute its drill samples instead of attempting to reason online about the measurements it has received. It is important to note that the β unique measurement types take on values between 0 and 1 so $\sigma_s = 1.0$ represents a very noisy signal. Therefore, in situations where the sensor is not reliable and the resource budget is sufficient, a deterministic raster policy may be the best option. However, often in real-world scenarios, the resource budget is not large enough to fully cover the area of interest and the onboard sensors have some degree of reliability.

Third, we see that MCTS-DPW is able to achieve the lowest $\text{Tr}(\Sigma)$ and RMSE along the trajectory, further emphasizing its ability to balance exploration and exploitation. MCTS-DPW consistently achieves the largest reward while

simultaneously decreasing its uncertainty in the world belief and predicting the true world state the most accurately.

Finally, it is important to note that POMCP and POMCP-GCB are able to plan much faster than MCTS-DPW because they do not have to conduct inference over the Gaussian process belief online. There is a wide body of literature devoted to scaling Gaussian processes efficiently [32–34] but in the scenarios where the agent is collecting < 1000 samples the scalability of the Gaussian process proved not to be an issue in our experiments. In all our experiments, the average planning time for POMCP, POMCP-GCB, POMCPOW-GCB, and MCTS-DPW were on the order of 0.01, 0.02, 0.01, and 0.5 seconds, respectively.

VI. CONCLUSION

This work focused on the variant of the AIPP problem where the agent is equipped with multiple sensors. We formulated the AIPPMS problem as a belief Markov decision process with Gaussian process beliefs. Using an approximation of the mutual information along with the expected reward, we solved the belief MDP using MCTS-DPW. We showed that our approach was able to outperform those of previous AIPPMS approaches as well as the POMCPOW-GCB algorithm by balancing information-gathering with exploitative reward-seeking actions. Future work includes integrating more efficient Gaussian process belief representations, such as adaptive query resolutions and expanding the belief representation to reason over Gaussian process parameters. These areas of future research would allow a similar formulation to be extended to a variety of real-world problems.

ACKNOWLEDGMENTS

We would like to thank Kyle Wray, Shushman Choudhury, Jordan Ott, and Lauren Ho-Tseung for their insightful discussions throughout the process of this work.

REFERENCES

- [1] S. Choudhury, N. Gruver, and M. J. Kochenderfer, "Adaptive informative path planning with multimodal sensing," in *International Conference on Automated Planning and Scheduling (ICAPS)*, vol. 30, 2020, pp. 57–65.
- [2] A. Meliou, A. Krause, C. Guestrin, and J. M. Hellerstein, "Nonmyopic informative path planning in spatio-temporal models," in *AAAI Conference on Artificial Intelligence (AAAI)*, vol. 10, 2007, pp. 16–7.
- [3] A. Singh, A. Krause, and W. J. Kaiser, "Nonmyopic adaptive informative path planning for multiple robots," in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2009.
- [4] G. A. Hollinger, B. Englot, F. S. Hover, U. Mitra, and G. S. Sukhatme, "Active planning for underwater inspection and the benefit of adaptivity," *The International Journal of Robotics Research*, vol. 32, no. 1, pp. 3–18, 2013.
- [5] G. Hitz, A. Gotovos, M.-É. Garneau, C. Pradalier, A. Krause, R. Y. Siegwart, *et al.*, "Fully autonomous focused exploration for robotic environmental monitoring," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 2658–2664.
- [6] Z. W. Lim, D. Hsu, and W. S. Lee, "Adaptive informative path planning in metric spaces," *The International Journal of Robotics Research*, vol. 35, no. 5, pp. 585–598, 2016.
- [7] Y. Girdhar and G. Dudek, "Modeling curiosity in a mobile robot for long-term autonomous exploration and monitoring," *Autonomous Robots*, vol. 40, no. 7, pp. 1267–1278, 2016.
- [8] M. Toussaint, "The Bayesian search game," in *Theory and Principled Methods for the Design of Metaheuristics*, Springer, 2014, pp. 129–144.
- [9] M. Popović, T. Vidal-Calleja, G. Hitz, J. J. Chung, I. Sa, R. Siegwart, and J. Nieto, "An informative path planning framework for UAV-based terrain monitoring," *Autonomous Robots*, vol. 44, no. 6, pp. 889–911, 2020.
- [10] D. A. Duecker, B. Mersch, R. C. Hochdahl, and E. Kreuzer, "Embedded stochastic field exploration with micro diving agents using Bayesian optimization-guided tree-search and GMRFs," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 8649–8656.
- [11] J. Ott, S.-K. Kim, A. Bouman, O. Peltzer, M. Sobue, H. Delecki, M. J. Kochenderfer, J. Burdick, and A.-a. Agha-mohammadi, *Risk-aware meta-level decision making for exploration under uncertainty*, 2022. arXiv: 2209.05580.
- [12] P. Morere, R. Marchant, and F. Ramos, "Continuous state-action-observation POMDPs for trajectory planning with Bayesian optimisation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 8779–8786.
- [13] J. Rückin, L. Jin, and M. Popović, "Adaptive informative path planning using deep reinforcement learning for UAV-based active sensing," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2022.
- [14] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, "Planning and acting in partially observable stochastic domains," *Artificial Intelligence*, vol. 101, no. 1-2, pp. 99–134, 1998.
- [15] A. Couëtoux, J.-B. Hoock, N. Sokolovska, O. Teytaud, and N. Bonnard, "Continuous upper confidence trees," in *Learning and Intelligent Optimization (LION)*, Springer, 2011, pp. 433–445.
- [16] M. J. Kochenderfer, T. A. Wheeler, and K. H. Wray, *Algorithms for Decision Making*. MIT Press, 2022.
- [17] R. He, E. Brunskill, and N. Roy, "Efficient planning under uncertainty with macro-actions," *Journal of Artificial Intelligence Research*, vol. 40, pp. 523–570, 2011.
- [18] M. J. Kochenderfer and T. A. Wheeler, *Algorithms for Optimization*. MIT Press, 2019.
- [19] R. Marchant, F. Ramos, S. Sanner, *et al.*, "Sequential Bayesian optimisation for spatial-temporal monitoring.," in *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2014, pp. 553–562.
- [20] P. Morere, R. Marchant, and F. Ramos, "Sequential Bayesian optimization as a POMDP for environment monitoring with UAVs," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 6381–6388.
- [21] I. M. R. 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*, 2022.
- [22] W. F. Caselton and J. V. Zidek, "Optimal monitoring network designs," *Statistics & Probability Letters*, vol. 2, no. 4, pp. 223–227, 1984.
- [23] A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies.," *Journal of Machine Learning Research*, vol. 9, no. 2, 2008.
- [24] Z. N. Sunberg and M. J. Kochenderfer, "Online algorithms for POMDPs with continuous state, action, and observation spaces," in *International Conference on Automated Planning and Scheduling (ICAPS)*, 2018.
- [25] C. B. Browne, E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton, "A survey of Monte Carlo tree search methods," *IEEE Transactions on Computational Intelligence and AI in games*, vol. 4, no. 1, pp. 1–43, 2012.
- [26] A. Couëtoux and H. Dohmen, "Adding double progressive widening to upper confidence trees to cope with uncertainty in planning problems," in *European Workshop on Reinforcement Learning (EWRL)*, 2011.

- [27] M. Egorov, Z. N. Sunberg, E. Balaban, T. A. Wheeler, J. K. Gupta, and M. J. Kochenderfer, “POMDPs.jl: A framework for sequential decision making under uncertainty,” *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 831–835, 2017.
- [28] J. Fischer and Ö. S. Tas, “Information particle filter tree: An online algorithm for POMDPs with belief-based rewards on continuous domains,” in *International Conference on Machine Learning (ICML)*, 2020, pp. 3177–3187.
- [29] R. Francis, T. Estlin, G. Doran, S. Johnstone, D. Gaines, V. Verma, M. Burl, J. Frydenvang, S. Montañó, R. C. Wiens, *et al.*, “Aegis autonomous targeting for chemcam on Mars science laboratory: Deployment and results of initial science team use,” *Science Robotics*, vol. 2, no. 7, eaan4582, 2017.
- [30] N. Mangold, S. Gupta, O. Gasnault, G. Dromart, J. Tarnas, S. Sholes, B. Horgan, C. Quantin-Nataf, A. Brown, S. Le Mouélic, *et al.*, “Perseverance rover reveals an ancient delta-lake system and flood deposits at Jezero crater, Mars,” *Science*, vol. 374, no. 6568, pp. 711–717, 2021.
- [31] J. L. Heldmann, A. Colaprete, R. C. Elphic, B. Bussey, A. McGovern, R. Beyer, D. Lees, and M. Deans, “Site selection and traverse planning to support a lunar polar rover mission: A case study at Haworth crater,” *Acta Astronautica*, vol. 127, pp. 308–320, 2016.
- [32] W. Bruinsma, E. Perim, W. Tebbutt, S. Hosking, A. Solin, and R. Turner, “Scalable exact inference in multi-output Gaussian processes,” in *International Conference on Machine Learning (ICML)*, 2020, pp. 1190–1201.
- [33] A. Damianou and N. D. Lawrence, “Deep Gaussian processes,” in *Artificial Intelligence and Statistics*, 2013, pp. 207–215.
- [34] E. Snelson and Z. Ghahramani, “Local and global sparse Gaussian process approximations,” in *Artificial Intelligence and Statistics*, 2007, pp. 524–531.