

Adaptive Sampling using POMDPs with Domain-Specific Considerations

Gautam Salhotra*, Christopher E. Denniston*, David A. Caron, Gaurav S. Sukhatme

Abstract—We investigate improving Monte Carlo Tree Search based solvers for Partially Observable Markov Decision Processes (POMDPs), when applied to adaptive sampling problems. We propose improvements in rollout allocation, the action exploration algorithm, and plan commitment. The first allocates a different number of rollouts depending on how many actions the agent has taken in an episode. We find that rollouts are more valuable after some initial information is gained about the environment. Thus, a linear increase in the number of rollouts, i.e. allocating a fixed number at each step, is not appropriate for adaptive sampling tasks. The second alters which actions the agent chooses to explore when building the planning tree. We find that by using knowledge of the number of rollouts allocated, the agent can more effectively choose actions to explore. The third improvement is in determining how many actions the agent should take from one plan. Typically, an agent will plan to take the first action from the planning tree and then call the planner again from the new state. Using statistical techniques, we show that it is possible to greatly reduce the number of rollouts by increasing the number of actions taken from a single planning tree without affecting the agent’s final reward. Finally, we demonstrate experimentally, on simulated and real aquatic data from an underwater robot, that these improvements can be combined, leading to better adaptive sampling.

I. INTRODUCTION

Adaptive sampling is the process of intelligently sampling the environment by an agent, such as an underwater or aerial robot. The robot does this by creating an internal model and selecting sampling positions that improve the model [1]. Adaptive sampling is often preferred to full workspace coverage plans when either a. the robot cannot cover the entire workspace due to a constrained time or energy budget, or b. an approximate model of the workspace is acceptable. Adaptive sampling can also make use of domain-specific information when it is available. For example, researchers in the area of algal bloom monitoring in aquatic ecosystems find areas of high chlorophyll concentration more valuable to study. Such use-cases naturally lend themselves to the integration of Bayesian optimization in adaptive sampling [2].

Solving adaptive sampling problems exactly is known to be NP-hard [3]. However, these problems can be solved using exact solvers [4], sampling-based planners [5] or Monte Carlo tree search (MCTS)-based solvers which sample random trajectories from the final reward distribution [6], [7]. Here we focus on using MCTS-based solvers to effectively sample complex environments. These iterative solvers generally use rollouts to sample a reward value for a given state by following trajectories from that state. The process of sampling from a final reward distribution using a random

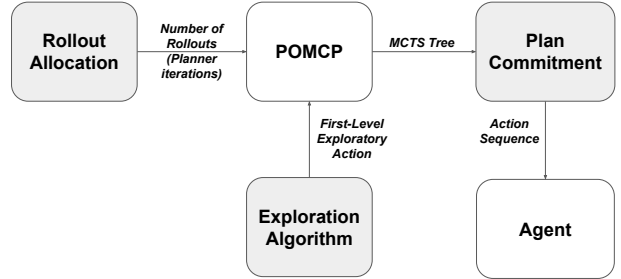



Fig. 1: Fig. a shows an agent’s trajectories on a hyperspectral orthomosaic collected in Clearlake, California. Blue is low value, red is high value. The baseline trajectory overlaps with itself (wasted time), goes outside the bounds of the orthomosaic (no reward outside workspace), and mostly samples near the starting position. The trajectory from our proposed method avoids such behavior and samples regions further away. Fig. b shows POMDP planning with the POMCP planner. Portions in grey are areas we study and improve in this work.

trajectory from a state at the leaf of the planning tree is called a rollout. Rollouts are used in local planners, such as POMCP [8], to sample discounted rewards over trajectories, from an unknown reward distribution.

In adaptive sampling, this reward distribution is defined by some objective function over samples.. Typically, rollouts are used to build an estimate of the mean reward for an action. By performing more rollouts, the planner improves its estimate of the expected reward for a particular sequence of actions.

Often, planning for adaptive sampling is done online. A fixed number of environment steps (typically one) are enacted after planning for a fixed number of iterations. This process is repeated until the finite budget (e.g., path length or energy [5]) is exhausted. Here, we verify that this process of committing to a fixed number of steps and rollouts at each invocation of the planner can be modified to reduce the total number of rollouts needed over the entire episode. We find that, in information gathering problems, there is a period when the information gathered is sufficient to predict the field accurately enough to make more useful plans. The intuition behind our result is that this period should be allocated more

* Equal contribution. {salhotra, cdennist}@usc.edu
All authors are with the University of Southern California

rollouts than the period when less information is known, or when gathering more samples does not result in as much reward. Additionally, more environment steps can be enacted from a single POMCP planning tree because the reward for future actions can be accurately predicted.

We cast the adaptive sampling problem as a Partially Observable Markov Decision Process (POMDP) which is solved using a local solver that updates the expected rewards for action sequences by sampling some reward distribution in an iterative fashion using rollouts. Specifically, we investigate minimizing the number of rollouts performed to achieve comparable accumulated reward by altering three parameters: the number of rollouts to perform, the choice of which actions to take in the planning tree during a planning iteration, and how many steps of the planned trajectory to follow.

II. BACKGROUND

Gaussian Processes are widely used modeling tools for adaptive sampling because of their non-parametric and continuous representation of the sensed quantity with uncertainty quantification [5], [9], [6]. Gaussian processes approximate an unknown function from its known outputs by computing the similarity between points from a kernel function [10]. Gaussian Processes are specifically useful for modeling the belief distribution of the underlying function from observations in the POMDP formulation of Bayesian optimization [11].

Online Adaptive Sampling consists of constructing an optimal path by alternating between planning and action. A plan is developed which attempts to maximize some objective function F by taking the actions describe by partial trajectory p . The partial trajectory is executed and samples are added to the model of the environment. These partial trajectories are concatenated to form full trajectory P . The plan and act iterations are iteratively interleaved until the cost $c(P)$ exceeds some budget B . Formally this is described by Eq. (1).

$$P^* = \operatorname{argmax}_{P \in \Phi} F(P) | c(P) \leq B \quad (1)$$

where Φ is the space of full trajectories, and P^* is the optimal trajectory [5].

Typically, F is some sub-modular function describing how informative the model of the environment sampled by P is. In this work, the objective function is $\mu_x + c\sigma_x$, where μ_x is the Gaussian process estimate of the underlying function value at x and σ_x^2 is the variance of the Gaussian process estimate at x . This objective function is commonly used in Bayesian adaptive sampling with a c parameter to trade off between exploration and exploitation. Formally a function, F , is submodular if $\forall A \subset B \subset V$ and $\forall s \in V \setminus B$, we have that $F(\{s\} \cup A) - F(A) \geq F(\{s\} \cup B) - F(B)$ [3]. This naturally describes diminishing returns exhibited in many adaptive sampling problems, that is, taking more samples provides more information the fewer samples you have already taken. The exploration term of the objective function, $\sigma(x)$, exhibits these submodularity characteristics [12].

Partially Observable Markov Decision Processes (POMDPs) are a framework for solving estimation problems when observations do not fully describe the state.

It has been shown that formulating observations as samples from the underlying world state and representing the robot’s model of the underlying function as a belief state can be formulated as a Bayesian Search Game [11], a framework for solving Bayesian optimization problems. In this game, an agent has to select points in the domain X that maximize the value of an unknown $f(x)$. Samples from $f(x)$ constitute observations which are partially observable components of the overall state f . If this state is augmented with the state of the robot, x , and constrained to locally feasible robot actions, this formulation can easily be extended to adaptive sampling [6]. To represent the underlying belief b at each state, a Gaussian process may be used. The problem of determining the appropriate planning horizon for related tasks has been investigated [13].

Partially Observable Monte Carlo Planning (POMCP): POMDPs have been used for adaptive sampling and informative path planners in many situations [14], [6], [15]. Many of these use a traditional dynamic programming solution to solve the underlying POMDP. This is infeasible for large state spaces or complex belief representations that are typically present in many common adaptive sampling problems. Recently, attention has focused to solvers which are locally accurate using probabilistic rollout updates [6]. A state of the art algorithm for solving large POMDPs online is the Partially Observable Monte-Carlo Planning solver (POMCP) [8]. POMCP uses a Monte Carlo tree search (MCTS) which propagates reward information from simulated rollouts in the environment. At every iteration in the planner, the agent performs a search through the generated tree \mathcal{T} to choose actions, using Upper-Confidence Tree (UCT) [16] exploration for partially observable environments, until it reaches a leaf node \mathcal{L} of \mathcal{T} . From node \mathcal{L} , the planner performs a rollout using a pre-defined policy (usually a random policy) until the agent reaches the planning horizon. The reward the agent collects while simulating this trajectory is then backpropagated up through the visited states. Finally, the first node the agent visited from the rollout is added as a child of \mathcal{L} , and the tree is expanded. Once the specified number of iterations (rollouts) are completed, the tree is considered adequate for planning. The action from the root node that gives the highest expected reward is then chosen, and the agent executes the action in the environment. At each observation node of \mathcal{T} , the observation is estimated with μ_x , where x is the agent’s position at that node. To update the belief b , the state-observation pair is integrated into the Gaussian Process.

Multi-Armed Bandits (MAB) are a family of problems in which K actions are available and each action has an unknown associated reward distribution, D_i . At each time step, the agent chooses an action $k \in K$ and receives a reward drawn from the distribution D_k . The goal of the agent is to maximize the cumulative reward over time or minimize risk, which is defined as the difference between the agent’s cumulative reward and some optimal policy.

There is a natural exploration-exploitation trade-off in the MAB problem because at each step the agent receives more information about the distribution of one of the actions by sampling the reward [17]. This framework provides the mechanism for action selection in a variety of rollout-based algorithms, including POMCP [8], and is used when each rollout can be viewed as a draw from the reward distribution conditioned on the currently selected actions.

In contrast to optimal action selection algorithms, there is a family of algorithms which seek to identify the best action in terms of mean reward. These algorithms work in two settings: fixed-budget and fixed-confidence. In the fixed-budget setting, the agent finds the best action from a fixed number of samples. In the fixed-confidence setting, the agent finds the best arm having $P[risk > \epsilon] < \delta$, in the fewest number of samples [18].

III. FORMULATION AND APPROACH

Most online informative planning pathfinders use the same planning duration at all points in the planning sequence. We propose to modify the POMCP [8] planner, an online, rollout-based POMDP-planner, given the knowledge about the underlying problem. Our method selects how many rollout iterations to use at each environment interaction and which actions should be tried in the planner. Using this tree, it also adaptively selects how much of the tree can be incorporated into the executed plan, based on the rewards received during these rollouts.

We show that we can produce similar models of the environment in fewer overall iterations of the (expensive) rollout sequence. An overview of how our improvements fit into the planning and action pipeline with POMCP is in Fig. 1b.

A. Rollout Allocation

The first improvement we propose is to alter how the overall rollout budget is handled. Typically, the total rollout budget is divided evenly and a fixed number of rollouts are allocated each time the planner is called to compute a new partial trajectory. This results in a linear increase in the number of rollouts used as the number of environment steps increases. We propose that this rollout allocation method should take advantage of three key ideas: cold-starting, submodularity, and starvation. The idea of cold-starting is well studied in adaptive sampling [9] and captures the notion that the planner can not make useful decisions with little information to plan on. Planning with little information is futile since far away points will generally return to the global mean with high variance. Typically, this is handled by having the robot perform a pre-determined pilot survey to gather initial information about the workspace [9]. This strategy wastes sampling time if too many pre-programmed samples are taken. The lack of information manifests itself in another problem when planning for adaptive sampling: submodularity of the objective function which necessitates effective sampling early on in the episode, since early samples provide the largest benefit. Additionally, because this

information is used in planning later on, the importance of good early information gathering is compounded. There is, of course, a trade-off to allocating rollouts early on: plans which allocate too many rollouts early on could suffer from the problem of starvation of rollouts at later stages. This can cause the planner to make poor decisions when there is rich information to plan on and the areas of interest are well-known and understood. This rollout allocation trade-off in planning is an instance of the exploration-exploitation trade-off which is common in information gathering tasks.

B. Exploration Algorithm

MCTS-based planners, such as POMCP, treat each action at each layer in the tree as an MAB problem. In the usual MCTS algorithm, the objective is to optimally explore possible trajectories by choosing an action at each layer according to some optimal action selection criterion [8]. We propose using optimal arm identification algorithms instead of optimal exploration algorithms because the final goal of the POMCP is to choose and execute the action with the highest mean reward, not to maximize the cumulative reward during searching.

Shifting from the cumulative-reward setting to the fixed-budget setting allows the exploration algorithm to decide the exploration-exploitation trade-off based on the number of allocated rollouts at each planning step. When many rollouts are allowed, the algorithm can be more explorative and consider each action for longer, while with fewer rollouts the algorithm typically becomes more exploitative. A fixed-budget action selection algorithm can only be used for the first action in the tree as these actions are the only ones for which the total number of combined rollouts is fixed.

In this work, we investigate three exploration algorithms. The first, Upper Confidence Tree (UCT), is an optimal exploration algorithm and is the default choice for most MCTS-based solvers, including POMCP [8]. UCT provides a trade-off between exploration and exploitation by adding an exploration bonus to each action using the number of times the parent and child have been explored. UCT does not take into account a budget, but instead tries to maximize the sum of the reward samples during planning. Because of this, UCT may tend to be highly explorative [19].

The remaining two algorithms are fixed-budget algorithms that explicitly incorporate the amount of rollouts allotted and attempt to maximize the probability that the chosen action is, in fact, the best action. The first algorithm, UGapEb, uses an upper bound on the simple regret of the actions to choose the action which is most likely to switch to become the best one to exploit [18]. This algorithm incorporates the total budget into the confidence interval to make the most efficient use of the provided budget. It contains a (difficult to estimate) parameter, H_ϵ , which requires the gap between actions to be known ahead of time.

Both UCT and UGapEb additionally depend on a (difficult to estimate) parameter, b , which is multiplied by the bound to create the exploration-exploitation trade-off. In this work, we use the difference between the highest and lowest discounted

rewards ever seen, but this may not be the optimal bound and requires the agent to know these values before exploring the workspace.

The final exploration algorithm is Successive Rejects. This algorithm is unique in that it does not attempt to use a confidence bound like UCT and UGapEb. Instead, it chooses each action a number of times in successive rounds and eliminates an action in each round until a single action is found [19]. This algorithm is preferable in some situations because it is parameter-free, while UCT and UGapEb have hard to tune parameters. However it may waste some rollouts when an action is obviously inferior.

C. Plan Commitment

Each call to the POMCP planner produces a planning tree of estimated rewards for action sequences. Typically, the agent executes the first action with the highest expected reward and replans with the part of the tree it went down. Since an adaptive sampling agent is approximating the underlying function using an initial belief, the generative model is different after each call to POMCP planner. This is because every time the agent takes an action in the environment and receives new samples, the initial belief for the underlying POMDP changes. Hence, the tree must be discarded and completely replanned after incorporating the observation from the environment into the initial belief state. We propose to take more than one action from the tree when there is certainty at lower levels about the optimality of the action. In order for the agent’s performance to be unaffected by taking further actions, there are two considerations it must have.

The first consideration is of the quality of the estimate of the reward from an action. If rollouts are spread evenly, the number of rollouts performed for lower actions (deeper in the tree) will be exponentially less than higher actions (closer to the root), causing the estimate of their reward to be poor.

The second consideration is the quality of the estimate of the observation at locations further away. The initial belief greatly affects the quality of observation estimates. Trajectories further away from what the agent has seen will generally have worse estimates for what their observations will likely be.

With these two considerations the agent may be able to extract more than one plan step from the tree without significant deterioration in the accumulated reward.

The simplest method is to take a fixed number of actions from the MCTS and execute them. This does not take into account any considerations of the quality of estimates of the reward for actions. With this method, the agent may have an inadequate understanding of the underlying world state at some point, but still take actions from the tree.

If the agent accounts for the statistics of the samples for the reward for actions, more complex methods can be used. We use a method based on UGapEc, a fixed-confidence MAB scheme [18]. UGapEc determines if there is an action that has a mean reward higher than all other actions with probability $1 - \delta$. This algorithm can be used to verify if a fixed-confidence threshold is met by checking if the algorithm




Fig. 2: Results from a grid search over possible rollout curves. Fig. 2a presents the curves that were searched over. Fig. 2b shows the three most optimal curves and a linear allocation curve, colored by their mean accumulated reward at the end of the episode.

would choose an action to explore. If it would not, the confidence criterion is passed and the action is added to the sequence.

Another method which uses a statistical test similar to UGapEc is a two tailed Welch’s t-test [20]. This test assumes the distributions the samples are from are Gaussian but does not assume the standard deviations are known or equal. Since the error in the estimate of standard deviation is quadratic in the number of samples, our estimate of the standard deviation deteriorates much faster than our estimate of the mean. Because of this, a more complex test must be used than a simple Gaussian confidence interval test since it may underestimate the sample standard deviation [21]. The unequal variances two tailed t-test tests the null hypothesis that the reward distributions for two actions have identical expected values. This method returns an estimate of the probability (p-value) that the means are equal. A threshold is set and if the p-value of the null hypothesis is below the threshold, the action is considered safe to take. This method is statistically robust. It causes the action to not be chosen in two cases. The first case is that there are not enough samples of each action and the second is that the means are too close to distinguish with the number of samples gathered. This method uses a Student’s t-distribution [22] and calculates the t statistic and the v value with Eq. (2) which can be used to compute the p-value with a Student’s t-distribution.

$$t = \frac{\bar{\mu}_1 - \bar{\mu}_2}{\sqrt{\frac{\bar{\sigma}_1^2}{N_1} + \frac{\bar{\sigma}_2^2}{N_2}}} \quad v \approx \frac{(\frac{\bar{\sigma}_1^2}{N_1} + \frac{\bar{\sigma}_2^2}{N_2})^2}{\frac{\bar{\sigma}_1^4}{N_1^2(N_1-1)} + \frac{\bar{\sigma}_2^4}{N_2^2(N_2-1)}} \quad (2)$$

where $\bar{\mu}_i$ is the sample mean reward for distribution i , $\bar{\sigma}_i$ is the sample standard deviation for distribution i , and N_i is the sample size for distribution i . We compare the reward distributions of the top two actions, with highest expected reward, to determine the p-value. We ignore other actions because of the asymmetrical nature of an MCTS tree causing the worst actions to have very few rollouts.

IV. EXPERIMENTS

We assess the performance of our improvements on three environments. The first environment is a test function for testing the effectiveness of sequential Bayesian Optimization using POMCP [6]. We use a dynamic (time-varying) two dimensional function as the underlying ground truth for




Fig. 3: Comparison of Exploration Algorithms (a, b, and c) and Plan Commitment Algorithms (d and e). *UCT, Fixed* is the baseline which evenly splits the rollouts at each step and uses the UCT exploration algorithm (the default for MCTS). Other results use a curved rollout allocation. For plan commitment, Fig. 3d shows the reward accumulation and Fig. 3e shows the number of rollouts used in all POMCP calls for the whole episode. A small offset value is added to methods which overlap.

testing our individual improvements. It corresponds to a Gaussian curve circling a fixed point twelve times.

$$f(x, y, t) = e^{-\left(\frac{x-2-1.5\sin(24\pi t)}{0.7}\right)^2} e^{-\left(\frac{y-2-1.5\cos(24\pi t)}{0.7}\right)^2} \quad (3)$$

where $x \in [0, 5], y \in [0, 5], t \in [0, 1]$. In this environment the agent starts out at the bottom-center of the time box and progresses towards the top, choosing actions in the x-y plane.

In the other two environments, called Validation Environment 1 (Fig. 5a) and Validation Environment 2 (Fig. 5d), we use chlorophyll concentration data collected from a YSI Ecomapper robot as input data. These datasets are $186m$ by $210m$, by $15m$ deep. We interpolate this with a Gaussian process to create the underlying function to estimate. In these scenarios the robot travels $3m$ between each sample and can freely travel in any direction at any point. The robot starts at the center of the environment at 0 depth.

For all environments, the agent is allowed to take 200 environment steps. We use the objective function $\mu_x + c\sigma_x$ and use $c = 10$ for the dynamic function and $c = 100$ for the validation environments. All experiments are run for five seeds each.

A. Grid Search for Rollout Allocation

To find the proper form of the rollout allocation and test the assertion that different parts of the POMDP planning process need different number of rollouts we perform a grid search over different curves that describe the rollout allocation. For each curve, if it would allocate less than one rollout per action, we allow the planner to perform a single rollout per action. We parameterize these curves by cumulative beta distributions because of their flexibility in representing many different kinds of curves. These curves are parameterized by an α and β parameter which determine the exact form of the curve. We search over $\alpha = [.75, 1, 2, 3, 4, 5, 6]$ and $\beta = [.75, 1, 2, 3, 4, 5, 6]$. These curves can be seen in Fig. 2a.

The results of this experiment are shown in Fig. 2b, which indicate that an exponential increase in the rollout allocations is desirable and a very flat curve is undesirable.

We find that the best curve for the dynamic function to be $\alpha = 6, \beta = 1$. We empirically find that this curve does worse on the validation environments, and that a curve with $\alpha = 4, \beta = 4$ works best. We use this for tests involving a curved rollout allocation with them. This may be due to the grid search overfitting the curve to the specific environment.

B. Comparison of Exploration Algorithms

We test the effectiveness of alternative exploration algorithms to UCT and the interaction between the rollout allocation method and exploration algorithm. We test three exploration algorithms described in Section III-B: UGapEb, UCT, and Successive-Rejects on three environments. In Fig. 3a all beta curve-based methods outperform the fixed method and all allocators work almost equally well, with UCT having a slight performance boost. In Fig. 3b, UGapEb and Successive Rejects with curved rollout allocation perform approximately equally but out-perform UCT with both a fixed and curved rollout allocation. In Fig. 3c all three curved allocators are out-performed by a fixed rollout allocation curve. This is likely because the rollout curve is poorly chosen for this environment due to not being chosen by grid search. UGapEb outperforms all curved allocators by a significant margin.

C. Comparison of Plan Commitment Algorithms

We test the methods described in Section III-C for determining how many steps to take once the MCTS tree is generated. We test the unequal variances t-test and UGapEc methods with different parameters against a baseline method which takes only the first action, across 5 seeds. Fig. 3d and Fig. 3e shows the comparison of all these combinations against the baseline. UGapEc and the baseline largely overlap because UGapEc cannot confidently predict whether the chosen action is the best action with such few rollouts and a larger epsilon does not make sense for the scale of the rewards. We believe that UGapEc may be of use for complex environments where the agent cannot make strong assumptions about the underlying reward distributions and many more rollouts will be required for the POMCP algorithm. The unequal variances t-test performs the best amongst the options. Within the t-test parameters, the p-value of 0.1 requires slightly fewer rollouts than a p-value of 0.05 for similar reward. However, choosing 0.1 implies riskier behavior which can have a negative effect in complex environments and real-world datasets (such as our validation environments). Hence, we choose the unequal variance t-test with $p = 0.05$ as our best choice for plan commitment. Fig. 3d shows the accumulated reward for a trajectory execution between the baseline and our choice. In Fig. 3e it is clear that each of the algorithms take vastly different amounts of rollouts for obtaining this result. Hence, we see that the plan commitment t-test algorithm helps to significantly reduce the rollouts needed to solve the adaptive sampling problem.

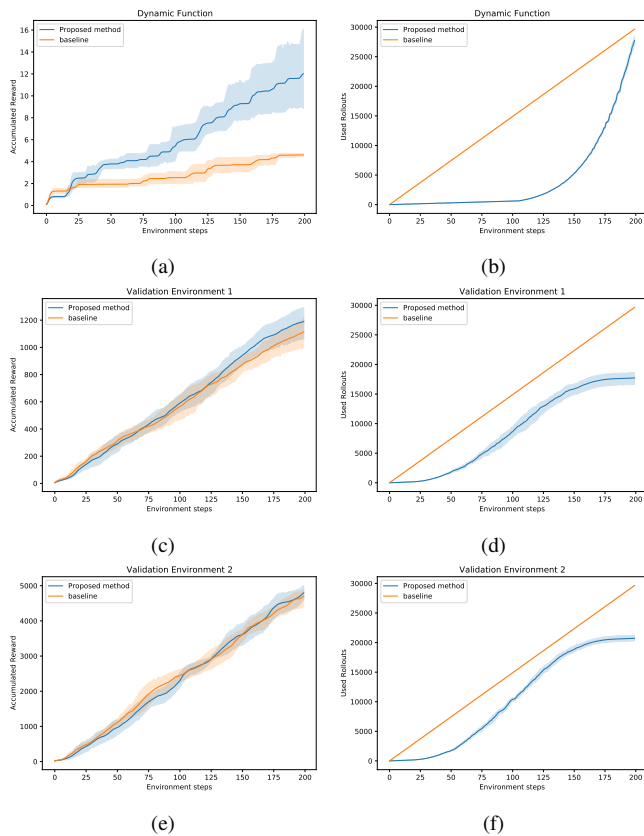


Fig. 4: Comparison of the combined proposed improvements (Proposed Method) against the baseline for all environments. Fig. 4a and Fig. 4b are the reward and number of rollouts used by the agent in the dynamic function environment. Fig. 4c and Fig. 4d are the reward and number of rollouts used by the agent in Validation Environment 1, Fig. 4e and Fig. 4f are the reward and number of rollouts used by the agent in Validation Environment 2.

D. Comparison of Baseline with Combined Improvements

Fig. 4 shows the combined effect of all improvements from the preceding experiments: using a curved rollout allocation, using the UGapEb exploration algorithm and using the t-test plan commitment algorithm. We compare against a baseline which uses an equal number of rollouts at each step, uses the UCT exploration algorithm and takes only one action before replanning. We compare our method and this baseline for each environment. Fig. 4a and Fig. 4b show that the combined features achieve a much higher reward in fewer rollouts on the dynamic environment. Fig. 4c and Fig. 4f show that again the agent receives a higher reward in many fewer rollouts than the baseline method. Fig. 4e and Fig. 4f indicate that our method is comparable to the baseline in terms of reward but achieves this reward in fewer rollouts.

Because our proposed method reduces the number of rollouts, the time to compute a plan is reduced. Additionally, rollouts from later environment steps are cheaper to compute because the remaining budget is lower. This causes our method to finish the episode faster because it allocates more rollouts to later environment steps. These two effects combine to produce a saving in wall-clock time for the entire episode, as shown in Table I. Experiments were run on a server with 2 Intel[®] Xeon Gold processors and 256GB RAM.

Method	Dynamic Function	Validation Environment 1	Validation Environment 2
Baseline	2061.84	2687.73	3542.72
Proposed Method	1371.77	2497.36	3120.90

TABLE I: Wall-clock time (seconds) required to complete five episodes.

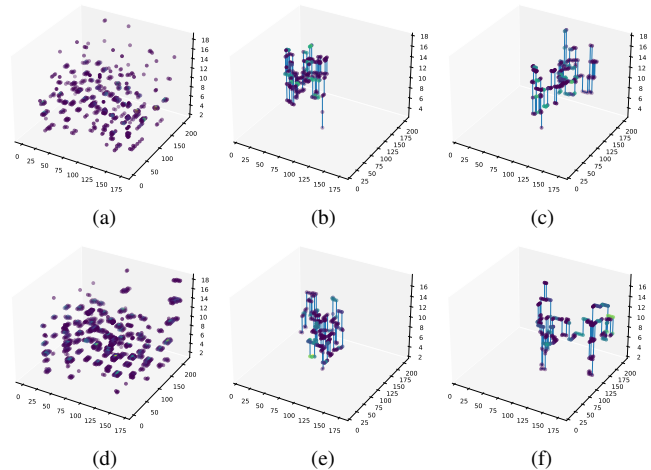


Fig. 5: Fig. 5a shows a dataset collected with an underwater robot, and Fig. 5b and Fig. 5c show example trajectories from a baseline implementation and our proposed implementation respectively. Fig. 5d shows another, more complex, dataset collected in the same location, Fig. 5e and Fig. 5f show example trajectories from a baseline implementation and our proposed implementation respectively.

V. CONCLUSION

We present improvements for online adaptive sampling with a Monte Carlo-based POMDP solver which uses specific knowledge of the adaptive sampling problem structure. This allows the agent to estimate a non-parametric function by taking samples of the underlying phenomenon such as the concentration of chlorophyll in a body of water.

First, we show that by changing the amount of rollouts that are allocated to more heavily favor later stages in planning, a better overall model of the environment can be created. We believe this is due to later stages having more information to plan on and therefore able to develop better and longer plans. We show that searching for an optimal curve can lead to high performance increases and that reasonable curves chosen can lead to increased performance. Second, we show that the agent's total reward can increase by changing the action exploration algorithm to one that explicitly incorporates knowledge of the number of rollouts allocated for each planning step. This can work with the rollout allocation to improve selection when few rollouts are allocated. We also show that by modifying the amount of steps the agent takes from a planning tree, the overall planning can be made more efficient. We show a statistical test can be used to determine if an action can be confidently determined to be the best action. With this test we are able to reduce the number of rollouts needed to reach a comparable accumulated reward. Finally, we show that these proposed improvements are synergistic and when used together can greatly improve the planning over a fixed-step, optimal exploration, fixed-rollout allocation planner.

REFERENCES

- [1] Jimin Hwang, Neil Bose, and Shuangshuang Fan. AUV Adaptive Sampling Methods: A Review. *Applied Sciences*, 9(15):3145, January 2019.
- [2] J. Das, K. Rajany, S. Frolovy, F. Pyy, J. Ryany, D. A. Caronz, and G. S. Sukhatme. Towards marine bloom trajectory prediction for AUV mission planning. In *2010 IEEE International Conference on Robotics and Automation*, pages 4784–4790, May 2010. ISSN: 1050-4729.
- [3] Andreas Krause and Carlos Guestrin. Submodularity and its applications in optimized information gathering. *ACM Transactions on Intelligent Systems and Technology*, 2(4):1–20, July 2011.
- [4] J. Binney and G. S. Sukhatme. Branch and bound for informative path planning. In *2012 IEEE International Conference on Robotics and Automation*, pages 2147–2154, May 2012. ISSN: 1050-4729.
- [5] Geoff A Hollinger and Gaurav S Sukhatme. Sampling-based robotic information gathering algorithms. *IJRR*, 33(9):1271–1287, 2014.
- [6] Roman Marchant, Fabio Ramos, and Scott Sanner. Sequential Bayesian optimisation for spatial-temporal monitoring. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence*, UAI’14, pages 553–562, Arlington, Virginia, USA, July 2014. AUAI Press.
- [7] Graeme Best, Oliver M Cliff, Timothy Patten, Ramgopal R Mettu, and Robert Fitch. Dec-MCTS: Decentralized planning for multi-robot active perception. *The International Journal of Robotics Research*, 38(3):316–337, 2019.
- [8] David Silver and Joel Veness. Monte-Carlo Planning in Large POMDPs. In J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 2164–2172. Curran Associates, Inc., 2010.
- [9] Stephanie Kemna, Oliver Kroemer, and Gaurav S. Sukhatme. Pilot Surveys for Adaptive Informative Sampling. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6417–6424, May 2018. ISSN: 2577-087X.
- [10] C E Rasmussen and C K I Williams. *Gaussian Processes for Machine Learning*. MIT press, 2006. ISSN: 0129-0657.
- [11] Marc Toussaint. The Bayesian Search Game. In Yossi Borenstein and Alberto Moraglio, editors, *Theory and Principled Methods for the Design of Metaheuristics*, pages 129–144. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014. Series Title: Natural Computing Series.
- [12] Carlos Guestrin, Andreas Krause, and Ajit Paul Singh. Near-optimal sensor placements in Gaussian processes. In *Proceedings of the 22nd international conference on Machine learning - ICML ’05*, pages 265–272, Bonn, Germany, 2005. ACM Press.
- [13] Xubo Yue and Raed AL Kontar. Why Non-myopic Bayesian Optimization is Promising and How Far Should We Look-ahead? A Study via Rollout. volume 108 of *Proceedings of Machine Learning Research*, pages 2808–2818, Online, August 2020. PMLR.
- [14] Pierre F.J. Lermusiaux, Tapovan Lolla, Patrick J. Haley Jr., Konuralp Yigit, Mattheus P. Ueckermann, Thomas Sondergaard, and Wayne G. Leslie. Science of Autonomy: Time-Optimal Path Planning and Adaptive Sampling for Swarms of Ocean Vehicles. In Manhar R. Dhanak and Nikolaos I. Xiros, editors, *Springer Handbook of Ocean Engineering*, pages 481–498. Springer International Publishing, Cham, 2016.
- [15] Jiu Hai-Feng, Chen Yu, Deng Wei, and Pang Shuo. Underwater Chemical Plume Tracing Based on Partially Observable Markov Decision Process. *International Journal of Advanced Robotic Systems*, 16(2):1729881419831874, March 2019. Publisher: SAGE Publications.
- [16] Levente Kocsis and Csaba Szepesvári. Bandit based monte-carlo planning. In *European conference on machine learning*, pages 282–293. Springer, 2006.
- [17] D. Bouneffouf, I. Rish, and C. Aggarwal. Survey on Applications of Multi-Armed and Contextual Bandits. In *2020 IEEE Congress on Evolutionary Computation (CEC)*, pages 1–8, July 2020.
- [18] Victor Gabillon, Mohammad Ghavamzadeh, and Alessandro Lazaric. Best Arm Identification: A Unified Approach to Fixed Budget and Fixed Confidence. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 3212–3220. Curran Associates, Inc., 2012.
- [19] Jean-Yves Audibert and Sébastien Bubeck. Best Arm Identification in Multi-Armed Bandits. page 13 p., June 2010.
- [20] B. L. Welch. The generalisation of student’s problems when several different population variances are involved. *Biometrika*, 34(1-2):28–35, 1947.
- [21] John Gurland and Ram C. Tripathi. A Simple Approximation for Unbiased Estimation of the Standard Deviation. *The American Statistician*, 25(4):30–32, October 1971. Publisher: Taylor & Francis.
- [22] WS Gosset. The Probable Error of a Mean, 1908. Originally published under the pseudonym.