

Bayesian Optimized Monte Carlo Planning

John Mern,¹ Anil Yildiz,¹ Zachary Sunberg,² Tapan Mukerji,³ and Mykel J. Kochenderfer¹

¹Stanford University, Department of Aeronautics and Astronautics, 496 Lomita Mall, Stanford, CA 94305

²University of Colorado Boulder, Department of Aerospace Engineering Sciences, 3775 Discovery Drive, Boulder, CO 80303

³Stanford University, Department of Energy Resources Engineering, 367 Panama Street, Stanford, CA 94305

{jmern91, yildiz, mukerji, mykel}@stanford.edu, zachary.sunberg@colorado.edu

Abstract

Online solvers for partially observable Markov decision processes have difficulty scaling to problems with large action spaces. Monte Carlo tree search with progressive widening attempts to improve scaling by sampling from the action space to construct a policy search tree. The performance of progressive widening search is dependent upon the action sampling policy, often requiring problem-specific samplers. In this work, we present a general method for efficient action sampling based on Bayesian optimization. The proposed method uses a Gaussian process to model a belief over the action-value function and selects the action that will maximize the expected improvement in the optimal action value. We implement the proposed approach in a new online tree search algorithm called Bayesian Optimized Monte Carlo Planning (BOMCP). Several experiments show that BOMCP is better able to scale to large action space POMDPs than existing state-of-the-art tree search solvers.

Introduction

The partially observable Markov decision process (POMDP) is a mathematical model of sequential decision making problems under state uncertainty (Littman, Cassandra, and Kaelbling 1995). In a POMDP, an agent takes actions while receiving noisy observations of the world state in order to accumulate reward. A policy that solves a POMDP maps histories of observations, represented as a belief over the state, to actions that maximize the expected sum of discounted rewards.

Solving POMDPs exactly is generally intractable and has been shown to be PSPACE-*complete* over finite horizons (Papadimitriou and Tsitsiklis 1987). A policy for a POMDP with n distinct states requires a function over an $(n - 1)$ -dimensional continuous belief space (Silver and Veness 2010). Computing exact solutions over this space requires evaluating a number of potential trajectories that is exponential in the time horizon (Pineau, Gordon, and Thrun 2006). These phenomena are referred to as the *curse of dimensionality* and the *curse of history*, respectively.

Because of these challenges, sample-based methods are often used to solve POMDPs approximately (Smith and Simmons 2004; Kurniawati, Hsu, and Lee 2008). Methods

based on Monte Carlo tree search (MCTS) have been particularly effective in recent years as online-solvers. MCTS methods work by incrementally building a search tree over trajectories reachable from the agent’s current belief and using the resulting tree to estimate the value of the available actions. While tree-search methods typically handle POMDPs with large state spaces well, they often perform poorly on problems with large action spaces (Browne et al. 2012b).

Large action and observation spaces tend to lead to excessive branching in approximate tree-search solvers such as POMCP (Silver and Veness 2010). High branching results in wide, shallow trees with relatively fewer visits per-branch, resulting in poor value estimates.

Progressive widening limits tree-width by sampling a subset of the available actions for addition to the search tree (Couëtoux et al. 2011). Random action sampling, however, may often fail to select good actions for evaluation, requiring expert-biased sampling for acceptable performance (Browne et al. 2012a).

This work frames progressive action selection as an optimization problem to be solved within MCTS. The proposed method models the problem at each branching step using a Bayesian optimization framework (Mockus 2002). Under this framework, the optimization objective is to select the action with the highest expected improvement in the action-value returned by the subsequent MCTS search. To calculate expected improvement, a distribution over the unknown action-value function is modeled as a Gaussian process (Rasmussen and Williams 2006). The procedure uses knowledge gained during tree search to guide action selection by conditioning the Gaussian process on the values of nodes already in the tree.

We implemented the proposed method in a new MCTS planner called Bayesian Optimized Monte Carlo Planning (BOMCP). The BOMCP algorithm builds on the state-of-the-art solver, POMCPOW (Sunberg and Kochenderfer 2018), to better scale to large action spaces. To efficiently scale the Bayesian optimization procedure, we do Gaussian process inference using a k -nearest neighbor approximation to model the value function distribution in the Bayesian optimizing procedure. Experiments show that this new algorithm outperforms POMCPOW, on several problems with large action-spaces.

Related Work

There has been significant work in online planning for POMDPs with Monte Carlo tree search (Browne et al. 2012b). One of the first successful planners, POMCP (Silver and Veness 2010) adapted UCT exploration (Kocsis and Szepesvári 2006) by sampling states from an unweighted particle set and searching over action-observation trajectories. POMCP cannot scale well to problems with large action or observation spaces due to its branching strategy that leads to high tree width.

DESPOT (Somani et al. 2013) builds on POMCP by using a deterministic generative model to sample scenarios, resulting in a sparse search tree. Due to the reduced observation branching, the number of nodes in the tree is significantly decreased and value estimates are improved. However, DESPOT can also struggle with large action-space problems without an expert action branching strategy.

Double Progressive Widening (DPW) (Couëtoux et al. 2011) was introduced to scale MDP planners to large discrete and continuous spaces. Progressive widening dynamically limits the number of child nodes that may be added to a parent node based on the number of times the parent has been visited in search. DPW, however, has been shown to be sensitive to the order nodes are selected for addition (Browne et al. 2012a) and has limited effect on scaling to *very* large action spaces on its own.

The POMCPOW and PFT-DPW algorithms (Sunberg and Kochenderfer 2018) introduce double progressive widening for POMDPs with large action and observation spaces. The algorithms maintain weighted particle beliefs at internal nodes to prevent the degeneration of multiple beliefs to a single particle. However, due to the random sampling used during action branching, POMCPOW can still struggle on large scale problems. In particular, Sunberg and Kochenderfer only apply POMCPOW to an action space with a single continuous dimension.

Work has been done to develop intelligent action selection methods for progressive widening for a variety of specialized cases. One method suggests using MCTS for a global search of coarsely discretized actions, while applying a finer local search to more promising actions, using value gradients (Lee et al. 2020). The value gradient calculation was found to be computationally expensive for many problems. GPS-ABT (Seiler, Kurniawati, and Singh 2015) uses generalized pattern search to find local optima in POMDPs with discrete observation spaces.

Another method uses Voronoi partitioning to dynamically cluster large action spaces and search by sampling from the Voronoi cells (Kim et al. 2020). However, cells are created at random locations and diameters and do not use knowledge gained during search. The Kernel Regression UCT algorithm (KR-UCT) (Yee, Lisý, and Bowling 2016) introduces information sharing between action nodes of an MCTS tree through smooth kernel regression on value estimates. KR-UCT, however, is restricted to POMDPs with uncertainty in action execution.

The Continuous Belief Tree Search algorithm (CBTS) (Morere, Marchant, and Ramos 2016) selects actions using a Bayesian optimization process. CBTS does

not select actions that maximize the expected Q-value at each belief node, as is proposed in this work, but instead selects actions that maximize the upper confidence bound of immediate reward. CBTS also proposes using a separate Gaussian process for each belief node, which limits how much the action selection process can learn from experiences at other belief nodes. Additionally, the CBTS algorithm does not incorporate progressive widening and instead uses fixed heuristics to control tree growth.

Gaussian processes have also been used to solve MDPs (Imani, Ghoreishi, and Braga-Neto 2018) with uncertain model dynamics or computationally expensive environment models. This work proposes Gaussian Process Temporal Difference (GPTD) learning with Markov Chain Monte Carlo (MCMC) sampling as a method to approximate the action value function.

Background

This section reviews several topics that are foundational to the remainder of the paper. The main discussion assumes familiarity with POMDPs, Monte Carlo tree search, and Gaussian processes, which are all briefly reviewed here.

POMDPs

Partially Observable Markov Decision Processes (POMDPs) are compact representations of sequential decision problems with uncertainty (Kochenderfer 2015). The problem state evolves according to probabilistic dynamics that are conditioned only on the current state and action. In a POMDP, the agent is assumed to only have access to noisy observations of the state. Using these observations, the agent’s task is to plan a sequence of actions that maximizes the total accumulated reward over the problem horizon.

A POMDP is defined by a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{O}, Z, T, R, \gamma)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, and \mathcal{O} is the observation space. The transition model $T(s' | s, a)$ gives the probability of transitioning from state s to state s' after taking action a . The reward function $R(s, a)$ specifies the immediate reward obtained after taking action a at state s . The observation model $Z(o | s, a, s')$ specifies the probability of receiving observation o in state s' given that action a had been taken at the previous state s . The $\gamma \in [0, 1]$ is a time discount factor.

When solving a POMDP, it is common to maintain a distribution over the state, called the belief b . The belief is updated each time the agent takes an action a and receives an observation o , typically with a Bayesian update filter.

For a POMDP, there exists an optimal policy $\pi^*(b)$ that specifies the best action to take in any belief state b . For a given policy, the action-value function $Q(b, a)$ gives the expected utility of taking action a at belief state b and then continuing to act according to the policy. An optimal policy selects an action that maximizes $Q(b, a)$, from belief b .

Monte Carlo Tree Search

Most online POMDP solvers employ sample based tree search methods typically referred to as Monte Carlo tree search. Partially observable MCTS methods construct

search trees as alternating layers of actions taken and observations received, with a root at the current belief. The tree is then used to estimate the value of actions the agent can take from its current state.

A typical MCTS process proceeds by sampling a state from the agent’s current belief. The state is used in a generative model to simulate potential trajectories. The simulation proceeds down the tree, selecting the best action from the existing branches until a leaf node is reached. It then adds a child node to the leaf and performs a rollout using a baseline policy to estimate the value. The value is propagated back through the parent nodes to the root to update their value estimates. Once a specified number of simulations have been run, the search returns the root action with the highest estimated value. Deeper trees tend to result in better estimates of the true optimal value.

Many modern MCTS approaches employ double progressive widening (DPW) to dynamically limit tree width for better value estimates. In DPW, when a node x is visited, a child node is added if

$$|\text{Ch}(x)| \leq kN(x)^\alpha \quad (1)$$

where $|\text{Ch}(x)|$ is the number of children of node x , $N(x)$ is the number of visits to node x , and k and α are hyperparameters. Separate k and α pairs are defined for observation and action progressive widening.

Gaussian Processes

A Gaussian process (GP) is a stochastic process which generates observations y given inputs x such that any subset of observations \mathbf{y} is distributed according to a multivariate Gaussian distribution. Representing a set of observations from a GP as the combination of two subsets \mathbf{y} and \mathbf{y}^* , the GP assumes

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{y}^* \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} m(\mathbf{x}) \\ m(\mathbf{x}^*) \end{bmatrix}, \begin{bmatrix} \Sigma_{xx} & \Sigma_{xx^*} \\ \Sigma_{x^*x} & \Sigma_{x^*x^*} \end{bmatrix} + \mathbb{I}\sigma^2\right) \quad (2)$$

where m is a function of the input x representing the prior mean. In practice, a constant mean is often used. The Σ matrices are the marginal and cross covariance matrices that compose the total covariance matrix. The covariance matrices are generated using a kernel function K such that $\Sigma_{ij} = K(x_i, x_j)$. Observation noise is added as the σ^2 term.

A typical use of a Gaussian process is to model the conditional distribution of unobserved data given some previously observed data. Taking \mathbf{y} and \mathbf{x} to be our previously observed data and \mathbf{y}^* and \mathbf{x}^* to be the new query points, the conditional distribution takes the form

$$P(\mathbf{y}^* | \mathbf{y}, \mathbf{x}, \mathbf{x}^*) = \mathcal{N}(\mathbf{y}^* | \mu^*, \Sigma^*), \quad (3)$$

$$\mu^* = m(\mathbf{x}^*) + \Sigma_{x^*x}(\Sigma_{xx} + \mathbb{I}\sigma^2)^{-1}(\mathbf{y} - m(\mathbf{x})) \quad (4)$$

$$\Sigma^* = \Sigma_{x^*x^*} - \Sigma_{x^*x}(\Sigma_{xx} + \mathbb{I}\sigma^2)^{-1}\Sigma'_{x^*x} \quad (5)$$

The above distributions may be used for calculating likelihoods or sampling new data at previously unobserved points.

BOMCP

Action progressive widening approaches in MCTS are known to require effective action sampling approaches in

order to achieve acceptable performance on tasks with large action spaces. In this work, we introduce a general method for action sampling that uses the information learned during earlier steps of the tree-search process to select subsequent actions for tree expansion. We present it in a new algorithm called Bayesian Optimized Monte Carlo Planning (BOMCP).

The main innovation of BOMCP is the use of Bayesian optimization to select actions for progressive widening during Monte Carlo tree search. Bayesian optimization is a black-box optimization approach that uses a probabilistic model to optimize over an unknown target function. In this case, the unknown optimization target is the action-value function $Q(b, a)$. Each time a new action is selected during tree-expansion, BOMCP selects the optimal action for addition based on the previously branched actions and their current action-value estimates.

Optimal Action Branching

MCTS with progressive widening operates by incrementally adding actions to the search tree when the branching threshold is met. For a belief node b with action child nodes $\text{Ch}(b)$, the maximum child node value estimate can be defined as

$$\hat{Q}(b, a^*) = \max_{a \in \text{Ch}(b)} \hat{Q}(b, a) \quad (6)$$

where $\hat{Q}(b, a)$ is the action-value estimate maintained in the search tree. When the search budget is exhausted, the action with the highest estimated action-value is returned.

The optimal action to add to the tree is the one expected to most improve the maximum value of the child node set. In Bayesian optimization, this is commonly referred to as the *Expected Improvement* acquisition function. We can define the expected improvement of the action-value to be

$$\text{EI}(a | b) := \mathbb{E}_{Q \sim GP} \left[|\hat{Q}(b, a) - \hat{Q}(b, a^*)|^+ \right] \quad (7)$$

where $|x|^+ = \max(x, 0)$. Only non-negative differences are used in calculating the expected improvement. If the expected value of a newly added action is lower than the current maximum, the action will not be selected as the MCTS return value.

To calculate EI, the expectation is taken with respect to a distribution over the unknown $Q(b, a)$ function. To model this distribution, we use a Gaussian process (GP) conditioned on the values of each previously visited node in the search tree $\{(ba_i, q_i)\}_{i=1, \dots, n}$, where q_i is the action-value estimate at belief-action node i , ba_i is a vector representing the corresponding belief and action values, and n is the number of action nodes in the tree.

The distribution of any unobserved belief action-value $q^*(b, a)$ is a Gaussian distribution defined by the GP posterior $P(q^* | ba^*, \{(ba, q)\}; \theta)$. Using this marginal distribution, an analytical expression for the expected improvement at any point can be defined (Jones, Schonlau, and Welch 1998) as

$$\begin{aligned} \text{EI}(a | b) = & |\Delta(ba)|^+ + \sigma(ba)\phi\left(\frac{\Delta(ba)}{\sigma(ba)}\right) \\ & - |\Delta(ba)|\Phi\left(\frac{\Delta(ba)}{\sigma(ba)}\right) \end{aligned} \quad (8)$$

where $\Delta(ba) = \mu(ba) - \hat{Q}(b, a^*)$, $\mu(ba)$ and $\sigma(ba)$ are the GP posterior mean and standard deviation, and ϕ and Φ are the standard normal probability density and cumulative density functions, respectively. The expected improvement can be easily calculated for any candidate action using this closed form and marginal statistics calculated as in eq. (3).

During action expansion at belief node b , the optimal action node

$$\tilde{a} \leftarrow \arg \max_{a \in \mathcal{A} \setminus \text{Ch}(b)} \text{EI}(a | b) \quad (9)$$

is added to the tree, and the resulting value estimate $\hat{Q}(b, \tilde{a})$ is added to the GP. Any appropriate optimizer may be used to solve eq. (9). In the proposed approach, L-BFGS (Liu and Nocedal 1989) is used for continuous action spaces and exhaustive search is used for discrete spaces.

Algorithm

We now define Bayesian Optimized Monte Carlo Planning (BOMCP). The tree search procedures in BOMCP are adapted from the POMCPOW algorithm. In addition to BOMCP, we also developed Bayesian Optimized Monte Carlo tree search (BOMCTS) for fully observable MDPs. BOMCTS is not presented in detail in this work, and more information may be found in the Appendix.

Overview The entry point to BOMCP is the PLAN procedure, which is shown in algorithm 1. The algorithm constructs a search tree by repeatedly calling SIMULATE from the root belief node of the tree for a specified number of times. The algorithm then returns the action at the root node with the highest estimated action-value.

The BOMCP search tree is constructed from alternating layers of belief nodes and action nodes. This is a departure from POMCP and POMCPOW, which use trees of alternating action and observation nodes, with observation nodes containing state particle collections. Using beliefs directly allows BOMCP to maintain less biased sampling of new states during progressive widening steps. Because BOMCP uses intermediate action-value estimates for the Bayesian optimization process, it was important to minimize the bias on these estimates. If desired, BOMCP may be configured to use particle collections to represent belief and recover the behavior of POMCPOW.

The majority of the computation occurs in the SIMULATE procedure. To begin each query, an initial state is sampled from the root belief and passed to SIMULATE. The state and belief are updated during each simulation using the provided generative model, GEN, and update filter, UPDATEBELIEF. Both action and belief progressive widening are used to determine expansion at each internal node. The UCT algorithm is used to select existing action nodes when not expanding.

SIMULATE is called recursively on the updated states and beliefs until a specified maximum depth is reached or the search reaches a leaf node. A rollout simulator is used to estimate the leaf node values. The leaf-node value estimates are then backed-up the tree along the current history to update the value estimates of the parent nodes.

Bayesian optimization is introduced during action progressive widening. When a new action is added to the search

Algorithm 1 Plan

```

1: procedure PLAN( $b, B$ )
2:    $T \leftarrow \text{Node}(b)$ 
3:   for  $i \in 1 : n$ 
4:      $s \sim b$ 
5:     SIMULATE( $s, b, d_{max}, T, B$ )
6:    $B \leftarrow \text{UPDATEBUFFER}(B, T)$ 
7:   return  $\arg \max_a Q(b, a), B$ 

```

tree, instead of selecting the action randomly from the accessible action space, BOMCP uses the procedure defined in algorithm 3. This routine constructs a Gaussian process using all previously visited belief-action node pairs (b, a) and corresponding $Q(b, a)$ estimates currently in the search tree T and buffer B . It then returns the action that maximizes the expected improvement.

Gaussian Process The Gaussian process is defined by the kernel, prior mean, and additive noise. In BOMCP, these are treated as hyper-parameters and must be specified. By default, BOMCP uses a squared-exponential kernel of the form

$$k_{SE}(x_1, x_2) = \sigma^2 \exp\left(-\frac{(x_1 - x_2)^2}{2\ell^2}\right) \quad (10)$$

where σ^2 is the marginal variance, and ℓ is the characteristic length. Larger characteristic lengths imply stronger correlation between belief-action values.

Using a constant prior mean μ_0 is common in Bayesian optimization and is the default setting for BOMCP. The μ_0 value can be used to tune the exploration behavior. Setting μ_0 to a lower bound on the action-value leads to conservative branching, with actions tending to be selected closer to the previously observed optimum, while setting μ_0 to an upper bound tends to encourage more optimistic exploration.

In order to use the Gaussian process, the inputs must be represented as vectors. A function VECTORIZE must be defined to transform belief-action pairs (b, a) to vector representations, ba . For high-dimensional belief spaces, it may be beneficial for VECTORIZE to represent the belief as a reduced dimensional collection of sufficient statistics.

In general, Gaussian process inference is not tractable for large data-sets, as it is $O(n^3)$ in the number of prior observations (Rasmussen and Williams 2006). This cubic dependence is caused by the need to invert the n -by- n covariance matrix. Several approximate methods have been proposed to reduce the size of this inversion (Silverman 1985).

BOMCP uses a k -nearest neighbor approach inspired by Sequential Gaussian Simulation (Gómez-Hernández and Journel 1993), which only considers the effects of the nearest observed points for each posterior calculation. This replaces the n -by- n inversion with a series of k -by- k inversions, where $k \ll n$. The resulting approximation is similar to existing nearest-neighbor Gaussian process regression techniques (Datta et al. 2016). Additional details on this process may be found in the Appendix.

Experience Buffer The actions added at the root of the tree are the most important to the performance of the search,

Algorithm 2 Simulate

```
1: procedure SIMULATE( $s, b, d, T, B$ )
2:   if  $d = 0$ 
3:     return 0
4:   if  $|\text{Ch}(b)| \leq k_a N(b)^{\alpha_a}$ 
5:      $a \leftarrow \text{BAYESOPT}(b, T, B)$ 
6:      $\text{Ch}(b) \leftarrow \text{Ch}(b) \cup \{a\}$ 
7:      $\text{Ch}(b, a) \leftarrow \emptyset$ 
8:    $a \leftarrow \arg \max_{a \in \text{Ch}(b)} \text{UCB}(b, a)$ 
9:    $s', o, r \leftarrow \text{GEN}(s, a)$ 
10:   $\text{New Node} \leftarrow \text{False}$ 
11:  if  $(b, a, b') \in T$ 
12:     $M(b, a, b') \leftarrow M(b, a, b') + 1$ 
13:  else if  $|\text{Ch}(b, a)| \leq k_b N(b, a)^{\alpha_b}$ 
14:     $b' \leftarrow \text{UPDATEBELIEF}(b, o)$ 
15:     $\text{INITIALIZENODE}(b, a, b', T)$ 
16:     $\text{New Node} \leftarrow \text{True}$ 
17:  else
18:     $b' \sim P(\text{Ch}(b, a) | b, a) \propto M(b', a, b)$ 
19:     $s' \sim b'$ 
20:     $r \leftarrow R(s, a, s')$ 
21:  if  $\text{New Node}$ 
22:     $q \leftarrow r + \gamma \text{ROLLOUT}(s', d - 1)$ 
23:  else
24:     $q \leftarrow r + \gamma \text{SIMULATE}(s', b', d - 1)$ 
25:   $N(b) \leftarrow N(b) + 1$ 
26:   $N(b, a) \leftarrow N(b, a) + 1$ 
27:   $Q(b, a) \leftarrow Q(b, a) + \frac{q - Q(b, a)}{N(b, a)}$ 
28:  return  $q$ 
```

since one of these actions will be returned. We would like to ensure good actions are added to the tree at the root. The Bayesian optimization procedure requires value estimates from existing nodes to be effective. Unfortunately, early in the search process, when many of the root actions are added, there are very few nodes. To offset this, we introduced an experience buffer, B , containing belief-action nodes from prior planning steps. The observations from this buffer are used with the current tree nodes to build the Gaussian process. The buffer is stochastically re-filled at the end of PLAN.

Algorithm 3 Bayesian Optimization

```
1: procedure BAYESOPT( $b, T, B$ )
2:    $X \leftarrow ()$ 
3:    $Y \leftarrow ()$ 
4:   for  $(b, a) \in T \cup B$ 
5:     append  $\text{VECTORIZE}(b, a)$  to  $X$ 
6:     append  $Q(b, a)$  to  $Y$ 
7:    $gp \leftarrow \text{GAUSSIANPROCESS}(X, Y)$ 
8:    $X^* \leftarrow \text{ACTIONS}(b) \setminus X$ 
9:    $\mu^*, \sigma^* \leftarrow \text{POSTERIOR}(gp, X^*)$ 
10:  return  $\arg \max_a \text{EI}(\mu^*, \sigma^*)$ 
```

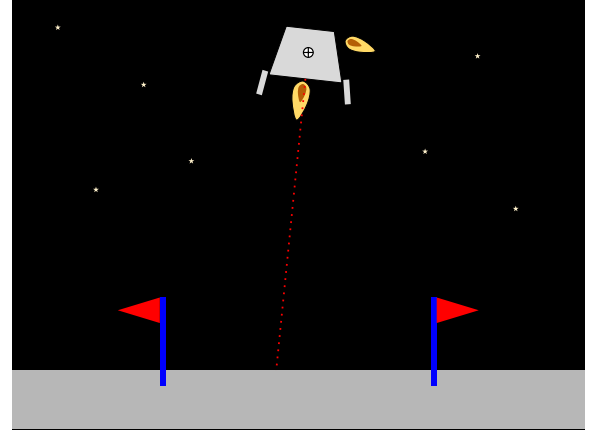


Figure 1: Partially Observable Lunar Lander. Control is applied through vertical and lateral thrust, depicted as the flames. Noisy observations of the altitude give the distance from the vehicle’s center to the floor along the vertical axis of the lander, as depicted by the dashed red line.

Experiments

We implemented BOMCP and BOMCTS in Julia building upon the POMDPs.jl package (Egorov et al. 2017). To evaluate the effectiveness of BOMCP, we conducted a series of experiments on three distinct POMDPs. We evaluated the performance of BOMCP against the performance of POMCPOW and expert policies for each problem. On the Lunar Lander experiment, we also compared BOMCP to an implementation of CBTS.

For each experiment, we recorded the task score as well as the wall clock run time per-search to measure the computation cost. We ran each experiment with varying numbers of queries-per-search. For all tests, the same values were used for hyper-parameters shared by BOMCP and POMCPOW such as K_{action} and α_{action} . Source code for BOMCP is available at <https://github.com/sisl/BOMCP.jl>.

Partially Observable Lunar Lander

The first problem studied was a partially-observable variant of the popular lunar-lander problem. The objective of the task is to guide a vehicle to land in a target area with low impact force. The environment is shown in fig. 1.

The vehicle state is represented by a six dimensional tuple $(x, y, \theta, \dot{x}, \dot{y}, \omega)$, where x and y are the horizontal and vertical positions, θ is the orientation angle, \dot{x} and \dot{y} are the horizontal and vertical speeds, and ω is the angular rate.

The vehicle makes noisy observations of its angular rate, horizontal speed, and altitude, which is represented as the dashed red line in fig. 1. The action space is a three-dimensional continuous space defined by the tuple (T, F_x, δ) . T is the main thrust which acts along the vehicle’s vertical axis through its center of mass, and is in the range $[0, 15]$. F_x is the corrective thrust, which acts along a horizontal axis offset from the center of mass by a distance δ . F_x is in the range $[-5, 5]$ and $\delta \in [-1, 1]$.

The initial vehicle state is sampled from a multivariate

Algorithm	Queries	Score	Time (seconds)
POMCPOW	10	-92 ± 9	0.003 ± 0.001
	50	-5 ± 6	0.009 ± 0.002
	100	13 ± 5	0.013 ± 0.007
	500	31 ± 4	0.051 ± 0.013
	1000	32 ± 3	0.130 ± 0.012
CBTS	10	-149 ± 14	0.029 ± 0.001
	50	21 ± 5	0.112 ± 0.010
	100	24 ± 5	0.156 ± 0.009
	500	34 ± 5	0.694 ± 0.019
BOMCP	10	-50 ± 1	0.028 ± 0.001
	50	36 ± 1	0.070 ± 0.011
	100	57 ± 2	0.141 ± 0.010
	500	61 ± 2	0.727 ± 0.017
Expert	-	-320 ± 28	-

Table 1: Lunar Lander Results. The mean and one standard error bound are given for the episode score and per-search wall clock runtime.

Gaussian with mean $\mu = (x = 0, y = 50, \theta = 0, \dot{x} = 0, \dot{y} = -10, \omega = 0)$. The reward function is defined as

$$r(s, a, s') = \begin{cases} -1000, & \text{if } x \geq 15 \vee \theta \geq 0.5 \\ 100 - x - v_y^2, & \text{if } y \leq 1 \\ -1, & \text{otherwise} \end{cases} \quad (11)$$

The first term in the reward function provides a penalty for the vehicle entering an unrecoverable state. The second term provides a positive reward for landing the vehicle, minus a penalty for drift away from center and for impact velocity. The final term is a constant penalty for fuel consumption.

We tested BOMCP, POMCPOW, and CBTS with varying numbers of queries per tree search. An Extended Kalman Filter was used to maintain a multi-variate Gaussian belief over vehicle state. An expert policy was also tested. This expert policy was also used as the rollout policy for leaf node evaluations for BOMCP, POMCPOW, and CBTS. The results are summarized in table 1.

For each number of queries-per-search tested, BOMCP outperforms both POMCPOW and CBTS by a statistically significant margin. Despite the low-dimensional action-space, optimal action selection still significantly improves BOMCP performance. BOMCP outperforming CBTS suggests that expected improvement is a better action selection criteria than expected reward.

BOMCP and CBTS take significantly longer per-search than POMCPOW. Profiling BOMCP PLAN revealed that approximately 30% of the time was spent in action selection. In POMCPOW, negligible time was spent in action selection.

We ran an additional set of tests on POMCPOW for 1000 queries-per-search. With 1000 queries, POMCPOW took as much time as BOMCP with 100 queries. BOMCP at 100 queries still outperformed POMCPOW at 1000, with both taking approximately 0.13 seconds per-search. This suggests

BOMCP outperforms POMCPOW by selecting better actions, rather than by just using additional computation.

Wind Farm Planning

The second task was large-scale wind farm planning. In this task, the agent sequentially selects locations to install sensor towers in a large three-dimensional wind field. The objective is to generate accurate maps of high wind areas which will later be used to plan turbine layouts.

The environment state is represented by a three-dimensional wind map of average annual wind-speed at discrete locations in the wind field. We used data from the Global Wind Atlas at the Altamont Pass wind farm, which covers an area of approximately 392 km². The wind-field at a single altitude is shown in fig. 2.

We selected a sub-region of the field for our experiments. The map of this region was a $20 \times 20 \times 3$ array, covering a 4440 m \times 4440 m area at altitudes 50 m, 100 m, and 150 m. In addition to the wind map, the state also contains the location of all previously placed sensor towers.

At each time step, the agent may choose to place a sensor tower at any unoccupied grid location on the map, at any of the three altitudes. The action space, therefore, contains at most 1200 distinct actions. The agent receives a noiseless observation of the wind value at each sensor tower location, but nowhere else. A sensor tower observes wind values all altitudes at and below its height. It is assumed that some initial knowledge of the wind field is available, represented as a set of sparse prior observations in a Gaussian process.

Reward is generated by first passing the current belief to a turbine-layout optimizer. The optimizer produces a risk-sensitive turbine arrangement by greedily selecting locations with the maximum $1 - \sigma$ lower confidence bound. The resulting layout is then used to estimate total annual power production based on the true wind field state. An additional cost is incurred for each tower, linearly proportional to its height. In this way, the reward encourages sensor arrangements that reduce variance in areas of high wind.

We tested BOMCP and POMCPOW with varying numbers of queries per tree search for 200 trials each. Mean returns with one standard error bounds are shown in table 2.




Figure 2: Wind Map. Figure shows wind map for Altamont Pass, CA at 100m altitude. The colors represent the average annual wind speed in m/s.

Algorithm	Queries	Score	Time (seconds)
POMCPOW	10	15708 ± 229	2.25 ± 0.07
	25	16234 ± 217	4.80 ± 0.07
	50	16374 ± 212	6.27 ± 0.08
	100	16018 ± 262	11.98 ± 0.07
	200	15787 ± 233	20.67 ± 0.09
BOMCP	10	18095 ± 183	2.55 ± 0.08
	25	18154 ± 158	5.21 ± 0.07
	50	18015 ± 163	6.71 ± 0.06
	100	18225 ± 119	13.39 ± 0.07
	200	18113 ± 157	25.14 ± 0.08
Expert	–	8130 ± 51	–

Table 2: Wind Farm Planning Results. The mean and one standard error bound are given for the episode score and per-search wall clock runtime.

For this large, discrete action-space problem, BOMCP significantly outperformed POMCPOW. Both BOMCP and POMCPOW seem to reach maximum performance at approximately 50 queries. POMCPOW’s performance never exceeds the performance of BOMCP with even just 1 query. Somewhat surprisingly, with a single query each, BOMCP still outperforms POMCPOW. This is likely due to the BOMCP experience buffer, which allows BOMCP to choose intelligent actions while POMCPOW samples randomly.

Cyber Security

The final task was a network cyber security task in which the agent scans nodes on a network in order to detect and eliminate a spreading malware infection. This task was selected because a distance metric between Gaussian process observations was not as clearly apparent. The problem tests the robustness of BOMCP to more difficult to represent spaces.

The network is composed of a set of Local Area Networks (LANs), where each LAN is a set of fully connected host nodes and a single server node. Edges between server nodes are generated randomly, though the graph is constrained to be complete. The network also contains one special “vault” server node, which is the malware target.

The state is represented by the infection states of the individual nodes as well as the edge topology. There are 4 LANs and 10 host nodes per LAN. The infection spreads according to a known stochastic adversary policy.

Each step, the agent may scan every node on a single LAN or scan an individual node. When scanning a LAN, the agent detects malware on all LAN nodes with probability 0.3. When scanning single nodes, the agent detects malware with probability 0.5 and cleans it with probability 0.8.

The episode terminates if the vault node is infected or 250 timesteps have passed. The agent receives no reward on non-terminal steps. For terminal steps, the reward function is

$$r(s, a, s') = \begin{cases} -100 - 0.5|S_i| - 0.1|H_i|, & \text{if vault infected} \\ -0.5|S_i| - 0.1|H_i|, & \text{otherwise} \end{cases} \quad (12)$$

Algorithm	Queries	Score	Time (seconds)
POMCPOW	50	-62 ± 4	0.13 ± 0.07
	100	-49 ± 5	0.23 ± 0.07
	500	-31 ± 4	1.01 ± 0.22
BOMCP	50	-27 ± 4	0.71 ± 0.14
	100	-23 ± 4	1.38 ± 0.27
	500	-23 ± 4	6.12 ± 0.49
Expert	–	-55 ± 4	–

Table 3: Cyber Security Results. The mean and one standard error bound are given for the episode score and per-search wall clock runtime.

where S_i is the set of infected server nodes and H_i is the set of infected host nodes.

A Dynamic Bayes Network was used to represent the belief model, with a discrete Bayes filter used for updates. The geodesic distance between nodes was used in the Gaussian process kernel. The results are summarized in table 3.

As with the previous experiments, BOMCP outperforms POMCPOW for a given number of queries, though at higher cost per query. Despite using a non-physical distance in the Gaussian Process kernel, BOMCP still solves the problem.

Conclusions

In this work, we framed MCTS action selection as a Bayesian optimization problem. We solved this optimization problem in a new POMDP planning algorithm called BOMCP. Experiments showed that BOMCP significantly outperforms POMCPOW on a variety of large scale POMDPs, though at much higher computational cost. These results suggest that the current BOMCP implementation may be best suited for problems with computationally expensive rollout simulators.

This work demonstrated significant improvements to MCTS through optimized action branching. The work also showed that direct application of Bayesian optimization can result in significant increase in computational cost, even with Gaussian process approximations. Future work will explore additional methods to reduce computational cost, including use of different distributional models.

Future work will also investigate different formulations of the Bayesian Optimization problem, for example using multi-fidelity Bayesian Optimization. Alternative acquisition functions to Expected Improvement will be considered. In some instances, tuning the Gaussian process hyperparameters required significant trial and error. Methods, to dynamically tune these parameters will also be developed.

Using the Gaussian process distribution over $Q(b, a)$ to generalize a search tree to unobserved trajectories will also be explored to allow it to be used in an offline setting.

Acknowledgements

The research reported in this work was supported by The ExxonMobil Research and Engineering Company through

the Stanford Strategic Energy Alliance.

References

- Browne, C.; Powley, E. J.; Whitehouse, D.; Lucas, S. M.; Cowling, P. I.; Rohlfshagen, P.; Tavener, S.; Liebana, D. P.; Samothrakis, S.; and Colton, S. 2012a. A Survey of Monte Carlo Tree Search Methods. *IEEE Transactions on Computational Intelligence and AI in Games* 4(1): 1–43.
- Browne, C. B.; Powley, E.; Whitehouse, D.; Lucas, S. M.; Cowling, P. I.; Rohlfshagen, P.; Tavener, S.; Perez, D.; Samothrakis, S.; and Colton, S. 2012b. A survey of Monte Carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in Games* 4(1): 1–43.
- Couëtoux, A.; Hoock, J.; Sokolovska, N.; Teytaud, O.; and Bonnard, N. 2011. Continuous Upper Confidence Trees. In *Learning and Intelligent Optimization (LION)*.
- Datta, A.; Banerjee, S.; Finley, A. O.; and Gelfand, A. E. 2016. On nearest-neighbor Gaussian process models for massive spatial data. *Computational Statistics* 8(5): 162–171.
- Egorov, M.; Sunberg, Z. N.; Balaban, E.; Wheeler, T. A.; Gupta, J. K.; and Kochenderfer, M. J. 2017. POMDPs.jl: A Framework for Sequential Decision Making under Uncertainty. *Journal of Machine Learning Research* 18: 26:1–26:5.
- Gómez-Hernández, J. J.; and Journel, A. G. 1993. Joint sequential simulation of multigaussian fields. In *Geostatistics Troia*, 85–94. Springer.
- Imani, M.; Ghoreishi, S. F.; and Braga-Neto, U. M. 2018. Bayesian Control of Large MDPs with Unknown Dynamics in Data-Poor Environments. In *Advances in Neural Information Processing Systems (NeurIPS)*, 8157–8167.
- Jones, D. R.; Schonlau, M.; and Welch, W. J. 1998. Efficient Global Optimization of Expensive Black-Box Functions. *J. Global Optimization* 13(4): 455–492.
- Kim, B.; Lee, K.; Lim, S.; Kaelbling, L. P.; and Lozano-Pérez, T. 2020. Monte Carlo Tree Search in Continuous Spaces Using Voronoi Optimistic Optimization with Regret Bounds. In *AAAI Conference on Artificial Intelligence (AAAI)*, 9916–9924.
- Kochenderfer, M. 2015. *Decision Making Under Uncertainty: Theory and Application*. MIT Press.
- Kocsis, L.; and Szepesvári, C. 2006. Bandit Based Monte-Carlo Planning. In *European Conference on Machine Learning (ECML)*.
- Kurniawati, H.; Hsu, D.; and Lee, W. S. 2008. SARSOP: Efficient Point-Based POMDP Planning by Approximating Optimally Reachable Belief Spaces. In *Robotics: Science and Systems*.
- Lee, J.; Jeon, W.; Kim, G.-H.; and Kim, K.-E. 2020. Monte-Carlo Tree Search in Continuous Action Spaces with Value Gradients. In *AAAI Conference on Artificial Intelligence (AAAI)*, 4561–4568.
- Littman, M. L.; Cassandra, A. R.; and Kaelbling, L. P. 1995. Learning Policies for Partially Observable Environments: Scaling Up. In *International Conference on Machine Learning (ICML)*.
- Liu, D. C.; and Nocedal, J. 1989. On the limited memory BFGS method for large scale optimization. *Math. Program.* 45(1-3): 503–528.
- Mockus, J. 2002. Bayesian heuristic approach to global optimization and examples. *Journal of Global Optimization* 22(1-4): 191–203.
- Morere, P.; Marchant, R.; and Ramos, F. 2016. Bayesian optimisation for solving continuous state-action-observation POMDPs. *Advances in Neural Information Processing Systems (NIPS)*.
- Papadimitriou, C. H.; and Tsitsiklis, J. N. 1987. The complexity of Markov decision processes. *Mathematics of Operations Research* 12(3): 441–450.
- Pineau, J.; Gordon, G. J.; and Thrun, S. 2006. Anytime Point-Based Approximations for Large POMDPs. *Journal of Artificial Intelligence Research* 27: 335–380.
- Rasmussen, C. E.; and Williams, C. K. I. 2006. *Gaussian Processes for Machine Learning*. MIT Press.
- Seiler, K. M.; Kurniawati, H.; and Singh, S. P. N. 2015. An online and approximate solver for POMDPs with continuous action space. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2290–2297.
- Silver, D.; and Veness, J. 2010. Monte-Carlo Planning in Large POMDPs. In *Advances in Neural Information Processing Systems (NIPS)*.
- Silverman, B. W. 1985. Some aspects of the spline smoothing approach to non-parametric regression curve fitting. *Journal of the Royal Statistical Society: Series B (Methodological)* 47(1): 1–21.
- Smith, T.; and Simmons, R. G. 2004. Heuristic Search Value Iteration for POMDPs. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 520–527.
- Somani, A.; Ye, N.; Hsu, D.; and Lee, W. S. 2013. DESPOT: Online POMDP Planning with Regularization. In *Advances in Neural Information Processing Systems (NIPS)*.
- Sunberg, Z. N.; and Kochenderfer, M. J. 2018. Online Algorithms for POMDPs with Continuous State, Action, and Observation Spaces. In *International Conference on Automated Planning and Scheduling (ICAPS)*.
- Yee, T.; Lisý, V.; and Bowling, M. H. 2016. Monte Carlo Tree Search in Continuous Action Spaces with Execution Uncertainty. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 690–697.