

Data-Efficient Reinforcement Learning for Malaria Control

Lixin Zou

Baidu Inc.

zoulixin15@gmail.com

Abstract

Sequential decision-making under cost-sensitive tasks is prohibitively daunting, especially for the problem that has a significant impact on people’s daily lives, such as malaria control, treatment recommendation. The main challenge faced by policymakers is to learn a policy from scratch by interacting with a complex environment in a few trials. This work introduces a practical, data-efficient policy learning method, named Variance-Bonus Monte Carlo Tree Search (VB-MCTS), which can copy with very little data and facilitate learning from scratch in only a few trials. Specifically, the solution is a model-based reinforcement learning method. To avoid model bias, we apply Gaussian Process (GP) regression to estimate the transitions explicitly. With the GP world model, we propose a variance-bonus reward to measure the uncertainty about the world. Adding the reward to the planning with MCTS can result in more efficient and effective exploration. Furthermore, the derived polynomial sample complexity indicates that VB-MCTS is sample efficient. Finally, outstanding performance on a competitive world-level RL competition and extensive experimental results verify its advantage over the state-of-the-art on the challenging malaria control task.

1 Introduction

Malaria is a mosquito-borne disease that continues to pose a heavy burden on South Sahara Africa (SSA) [Moran, 2007]. Recently, there has been significant progress in improving treatment efficiency and reducing the mortality rate of malaria. Unfortunately, due to financial constraints, the policymakers face the challenge of ensuring continued success in disease control with insufficient resources. To make intelligent decisions, learning control policies over the years have been formulated as Reinforcement Learning (RL) problems [Bent *et al.*, 2018]. Nevertheless, applying RL to malaria control seems to be a tricky issue since RL usually requires numerous trial-and-error searches to learn from scratch. Unlike simulation-based games, *e.g.*, Atari games [Mnih *et al.*, 2013] and game GO [Silver *et al.*, 2017],

the endless intervention trial is unacceptable to regions over the years since the actual cost of life and money is enormous. Hence, as in many human-in-loop systems [Zou *et al.*, 2019; Zou *et al.*, 2020b; Zou *et al.*, 2020a], it is too expensive to apply RL to learn malaria intervention policy from scratch directly.

Therefore, to reduce the heavy burden of malaria in SSA, it is urgent to improve the data efficiency of policy learning. In [Bent *et al.*, 2018], novel exploration techniques, such as Genetic Algorithm [Holland, 1992], Batch Policy Gradient [Sutton *et al.*, 2000] and Upper/Lower Confidence Bound [Auer and Ortner, 2010], have been firstly applied to learn malaria control policies from scratch. However, these solutions are introduced under the Stochastic Multi-Armed Bandit (SMAB) setting, which myopically ignores the delayed impact of the interventions in the future and might result in serious problems. For example, the large-scale use of spraying may lead to mosquito resistance and bring about the uncontrolled spread of malaria in the coming years. Hence, it requires us to optimize disease control policies in the long run, which is far more challenging than the exploration in SMAB.

Considering the long-term effects, the finite horizon continuous-space Markov Decision Process is employed to model the disease control in this work. Under this setting, we propose a framework named Variance-Bonus Monte Carlo Tree Search (VB-MCTS) for data-efficient policy searching, illustrated in Figure 1. Particularly, it is a model-based training framework, which iterates between updating the world model and collecting data. In model training, Gaussian Process (GP) is used to approximate the state transition function with collected rollouts. As a non-parametric probabilistic model, GP can avoid the model bias and explicitly model the uncertainty about the transitions, *i.e.*, the variance of the state. In data collection, we propose to employ MCTS for generating the policy with the mean MDP plus variance-bonus reward. The variance-bonus reward can decrease the uncertainty at the state-action pairs with high potential reward by explicitly motivating the agent to sample the state-actions with the highest upper-bounded reward. Furthermore, the sample complexity of the proposed method indicates that it is a PAC optimal exploration solution for malaria control. Finally, to verify the effectiveness of our policy search solution, extensive experiments are conducted on the malaria control

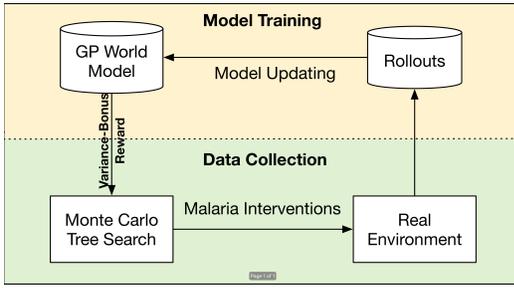


Figure 1: Overview of the proposed learning system on Malaria Control. The system alternates model training and data collection.

simulators¹ [Bent *et al.*, 2018], which are gym-like² environments for KDD Cup 2019 [Zhou *et al.*, 2020]. The outstanding performance on the competition and extensive experimental results demonstrated that our approach could achieve unprecedented data efficiency on malaria control compared to the state-of-the-art methods.

Our main contributions are: (1) We propose a highly data-efficient learning framework for malaria control. Under the framework, the policymakers can successfully learn control policies from scratch within 20 rollouts. (2) We derive the sample complexity of the proposed method and verify that VB-MCTS is an efficient PAC-MDP algorithm. (3) Extensive experiments conducted on malaria control demonstrate that our solution can outperform the state-of-the-art methods.

2 Related Work

As a highly pathogenic disease, malaria has been widely studied from the perspective of predicting the disease spread, diagnosis, and personalized care planning. However, rarely work focuses on applying RL to learn the cost-effectiveness intervention strategies, which plays a crucial role in controlling the spread of malaria [Moran, 2007]. In [Bent *et al.*, 2018], malaria control has firstly been formulated as a stochastic multi-armed bandit (SMAB) problem and solved with novel exploration techniques. Nevertheless, SMAB based solutions only myopically maximize instant rewards, and the ignorance of delayed influences might result in disease outbreaks in the future. Therefore, in this work, a comprehensive solution has been proposed to facilitate policy learning in a few trials under the setting of finite-horizon MDP.

Another topic is data-efficient RL. To increase the data efficiency, we are required to extract more information from available trials [Deisenroth and Rasmussen, 2011], which involves utilizing the samples in the most efficient way (*e.g.*, exploitation) and choosing the samples with more information (*e.g.*, exploration). Generally, for exploitation, model-based methods [Ha and Schmidhuber, 2018; Kamthe and Deisenroth, 2017] are more sample efficient but require more computation time for the planning. Model-free [Szita and Lőrincz, 2006; Krause *et al.*, 2016; Van Seijen *et al.*, 2009]

¹<https://github.com/IBM/ushiriki-policy-engine-library>

²<https://gym.openai.com>

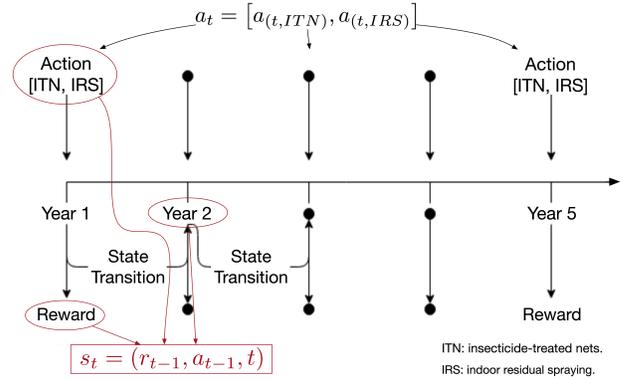


Figure 2: Malaria intervention as MDP.

methods are generally computationally light and can be applied without a planner, but need (sometimes exponentially) more samples, and are usually not efficient PAC-MDP algorithm [Strehl *et al.*, 2009]. For exploration, there are two options: (1) Bayesian approaches, considering a distribution over possible models and acting to maximize expected reward; unluckily, it is intractable for all but very restricted cases, such as the linear policy assumption in PILCO [Deisenroth and Rasmussen, 2011]. (2) intrinsically motivated exploration, implicitly negotiating the exploration/exploitation dilemma by always exploiting a modified reward for directly accomplishing exploration. However, on the one hand, the vast majority of papers only address the discrete state case, providing incremental improvements on the complexity bounds, such as MMDP-RB [Sorg *et al.*, 2012], metric-E³ [Kakade *et al.*, 2003], and BED [Kolter and Ng, 2009]. On the other hand, for more realistic continuous state space MDP, over-exploration has been introduced for achieving polynomial sample complexity in many work, such as KWIK [Li *et al.*, 2011], and GP-Rmax [Grande *et al.*, 2014]. These methods will explore all regions equally until the reward function is highly accurate everywhere. By drawing on the strength of existing methods, our solution is a model-based RL framework, which efficiently plans with MCTS and trade-off exploitation and exploration by exploiting a variance-bonus reward.

3 Proposed Method: VB-MCTS

3.1 Malaria Control as MDP

Finding an optimal malaria control policy can be posed as a reinforcement learning task by illustrating it as a Markov Decision Process (MDP). Specifically, we formulate the task as a finite-horizon MDP, defined by the tuple $\langle \mathcal{S}, \mathcal{A}, P, R, \gamma \rangle$ with \mathcal{S} as the potential infinite state space, \mathcal{A} as the finite set of actions, $P : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ as the deterministic transition function, $R : \mathcal{S} \rightarrow \mathbb{R}$ as the reward function, and $\gamma \in (0, 1]$ as the discount factor. In this case, we face the challenge of developing an efficient policy for a population over a 5 year intervention time frame. As shown in Figure 2, the corresponding components in malaria control are defined as,

Action. The actions are the available means of interventions, including the mass-distribution of long-lasting

insecticide-treated nets (ITNs) and indoor residual spraying (IRS) with pyrethroids in SSA [Stuckey *et al.*, 2014]. In this work, the action space \mathcal{A} is constructed through $a_i \in \mathcal{A} = \{(a_{ITN}, a_{IRS})\}$ with $a_{ITN}, a_{IRS} \in (0, 1]$, which represent the population coverage for ITNs and IRS of a specific area. Without significantly affecting performance, we discretize the action space with an accuracy of 0.1 for simplicity.

Reward. The reward is a scalar $r_t = R(s_{t+1}) \in [R_{\min}, R_{\max}]$, associated with next state s_{t+1} . In malaria control, it is determined through an economic cost-effectiveness analysis. In [Bent *et al.*, 2018], an overview of the reward calculation is specified. Without loss of generality, the reward function $R(\cdot)$ is assumed known to us since an MDP with unknown rewards and unknown transitions can be represented as an MDP with known rewards and unknown transitions by adding additional states to the system.

State. The state contains important observations for decision making in every time step. In malaria control, it includes the number of life with disability, life expectancy, life lost, and treatment expenses [Bent *et al.*, 2018]. We set the state in the form $s_t = (r_{t-1}, a_{t-1}, t)$, including current reward r_{t-1} , previous action a_{t-1} and the current intervention timestamp, which covers the most crucial and useful observations for malaria control, as shown in Figure 2. For the start state s_1 , the reward and action are initialized with 0.

Let $\pi : \mathcal{S} \rightarrow \mathcal{A}$ denote a deterministic mapping from states to actions, and let $V_\pi(s_t) = \mathbb{E}_{a_t \sim \pi(s_t)} \left[\sum_{i=t}^T \gamma^{i-t} r_i \right]$ denote the expected discounted reward by following policy π in state s_t . The objective is to find a deterministic policy π^* that maximizes the expected return at state s_1 as

$$\pi^* = \arg \max_{\pi \in \Pi} V_\pi(s_1).$$

3.2 Model-based Indirect Policy Search

In the following, we detail the key components of the proposed framework VB-MCTS, including the world model, the planner, and variance-bonus reward with its sample complexity.

World Model Learning. The probabilistic world model is implemented as a GP, where we use a predefined feature mapping $x_t = \phi(s_t, a_t) \in \mathbb{R}^m$ as training input and the target state s_{t+1} as the training target. The GP yields one-step predictions

$$\begin{aligned} p(s_{t+1}|s_t, a_t) &= \mathcal{N}(s_{t+1} | \mu_t, \sigma_t^2) \\ \mu_t &= \mathbf{k}_*^\top (K + \omega_n^2 I)^{-1} \mathbf{y}, \\ \sigma_t^2 &= k(x_t, x_t) - \mathbf{k}_*^\top (K + \omega_n^2 I)^{-1} \mathbf{k}_*, \end{aligned}$$

where k is the kernel function, $\mathbf{k}_* \equiv k(\mathbf{X}, x_t)$ denotes the vector of covariances between the test point and all training points with $\mathbf{X} = [x_1, \dots, x_n]$, and $\mathbf{y} = [s_1, \dots, s_n]^\top$ is the corresponding training targets. ω_n is the noise variance. K is the Gram matrix with entries $K_{ij} = k(x_i, x_j)$.

Throughout this paper, we consider a prior mean function $m \equiv 0$ and a squared exponential (SE) kernel with automatic relevance determination. The SE covariance function is defined as

$$k(x, x') = \alpha^2 \exp\left(-\frac{1}{2}(x - x')^\top \mathbf{\Lambda}^{-1}(x - x')\right),$$

where α^2 is the variance of state transition and $\mathbf{\Lambda} \equiv \text{diag}([l_1^2, \dots, l_m^2])$. The characteristic length-scale l_i controls the importance of i -th feature. Given n training inputs \mathbf{X} and the corresponding targets \mathbf{y} , the posterior hyper-parameters of GP (length-scales l_i and signal variance α^2) are determined through evidence maximization technique [Williams and Rasmussen, 2006].

Exploration with Variance-Bonus Reward. The algorithm we propose is itself very straightforward and similar to many previously proposed exploration heuristics [Kolter and Ng, 2009; Srinivas *et al.*, 2009; Sorg *et al.*, 2012; Grande, 2014]. We call the algorithm Variance-Bonus Reward, since it chooses action according to the current mean estimation of the reward plus an additional variance-based reward bonus for state-actions that have not been well explored as

$$\begin{aligned} \tilde{R}(s_{t+1}|s_t, a_t) &= R(s_{t+1})|_{s_{t+1} = \mathbb{E}_{s_{t+1} \sim GP(s_t, a_t)}[s_{t+1}]} + \\ &\beta_1 \text{Var}_{s_{t+1} \sim GP(s_t, a_t)}[R(s_{t+1})] + \beta_2 \text{Var}_{s_{t+1} \sim GP(s_t, a_t)}[s_{t+1}], \end{aligned}$$

where β_1 and β_2 are the parameters that trade-off the balance of exploitation and exploration. $\text{Var}_{s_{t+1} \sim GP(s_t, a_t)}[R(s_{t+1})]$ and $\text{Var}_{s_{t+1} \sim GP(s_t, a_t)}[s_{t+1}]$ are the predicted variances for state and reward. The variance of reward can be exactly computed following the law of iterated variances.

Planning with Mean MDP + Reward Bonus. MCTS is a strikingly successful planning algorithm [Silver *et al.*, 2017], which can find out the optimal solution with enough computation resources. Since disease control is not a real-time task, we propose to apply MCTS (Figure 1) as the planner for generating the policy to maximize the variance-bonus reward. In the executing process, the MCTS planner incrementally builds an asymmetric search tree guided to the most promising direction by a tree policy. This process usually consists of four phases — **selection**, **expansion**, **evaluation**, and **backup** (as shown in Figure 3).

Specifically, each edge (s, a) of the search tree stores an average action value $\tilde{Q}(s, a)$ and visit count $N(s, a)$. In the **selection** phase, starting from the root state, the tree is traversed by simulation (that is, descending the tree with the mean prediction of states without backup). At each time step t of each simulation, an action a_t is selected from state s_t

$$a_t = \arg \max_{a \in \mathcal{A}} \left(\tilde{Q}(s_t, a) + \frac{c_{puct} \sqrt{\sum_{b \in \mathcal{A}} N(s_t, b)}}{|\mathcal{A}| (1 + N(s_t, a))} \right),$$

so as maximize action value plus a bonus that decays with repeated visits to encourage exploration for tree search. Here, c_{puct} is the constant determining the level of exploration. When the traversal reaches a leaf node s_L at step L , the leaf node may be **expanded** with each edge initialized as $\tilde{Q}(s_L, a) = 0$, $N(s_L, a) = 0$, and the corresponding leaf nodes are initialized with the mean GP prediction $s_{L+1} = \mathbb{E}[s_{L+1}|s_L, a]$. Then, the leaf node is **evaluated** by the average outcome $v(s_L)$ of rollouts, which played out until terminal step T using the fast rollout policies, such as random policy and greedy policy. At the end of the simulation, the action values and visit counts of all traversed edges are updated, *i.e.*, **backup**. Each edge accumulates the visit counts

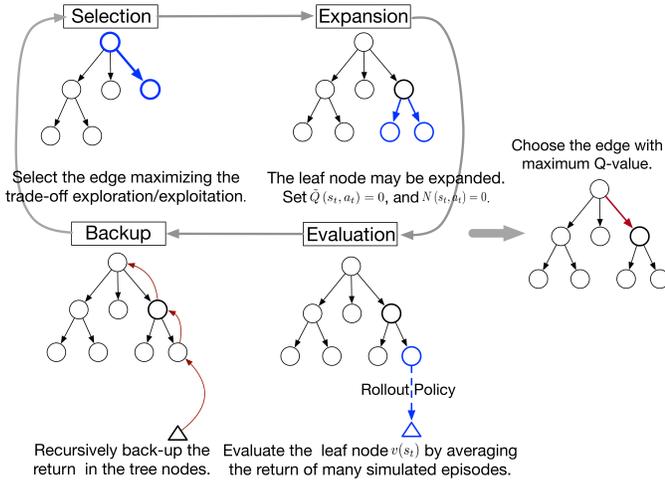


Figure 3: An illustration of MCTS for policy generation.

and means the evaluation of all simulations passing through that edge as

$$\begin{aligned}\tilde{Q}(s_j, a_j) &\leftarrow \frac{N(s_j, a_j) \times \tilde{Q}(s_j, a_j) + v(s_j)}{N(s_j, a_j) + 1}, \\ v(s_j) &\leftarrow v(s_{j+1}) + \tilde{R}(s_{j+1}|s_j, a_j), \\ N(s_j, a_j) &\leftarrow N(s_j, a_j) + 1,\end{aligned}$$

where (s_j, a_j) with $j < L$ is the edge in the forward trace.

This cycle of selection, evaluation, expansion and backup is repeated until the maximum iteration number has been reached. At this point, the best action has been chosen by selecting the action that leads to the highest reward (max child) as follows

$$\pi(s_t) = \arg \max_{a \in \mathcal{A}} \tilde{Q}(s_t, a), \quad (1)$$

where $\pi(s_t)$ is the policy generated by MCTS. To be noticed, the variance-bonus reward \tilde{R} can be replaced with the R for generating the best-known policy since it maximizes the expected reward based on the posterior so far.

Finally, we implement an iterative training procedure, as shown in Algorithm 1, where we specify the order in which each component occurs within the iteration.

3.3 Sample Complexity

In the following, we derive sample complexity (e.g., the required samples to learn near-optimal performances) for our proposed solution. In the worst case, when the reward function is equal everywhere, and all state-action pairs will be equally explored, VB-MCTS will have the same complexity bounds as over-exploration methods, such as GP-Rmax [Grande *et al.*, 2014]. In practice, VB-MCTS will learn the optimal policy in fewer steps than the over-exploration methods because the high uncertainty region with low reward is not worth exploration. Following the general PAC-MDP theorem, Theorem 10 in [Strehl *et al.*, 2009], we derive the polynomial sample complexity of VB-MCTS.

Algorithm 1: Variance-Bonus Monte Carlo Tree Search

Input: Hyper-parameters β , maximum trials number N .

Output: The malaria control policy π .

Initialization.

Generate random policy $\pi = [a_1, \dots, a_T]$.

Apply policy π in real world and collect training samples

$\mathcal{D} = \{(s_i, a_i, s_{i+1})\}_{i=1}^T$.

Update GP world model with \mathcal{D} .

Set number of trials $k = 1$.

for $k \leq N$ **do**

Initialize time step $t = 1$, and state s_t .

for $t \leq T$ **do**

Data Collection.

Generate action $a_t = \arg \max_{a \in \mathcal{A}} Q(s_t, a)$ by

MCTS with **variance-bonus reward** \tilde{R} .

Apply a_t in real world and collect s_{t+1} .

Update training samples $\mathcal{D} = \mathcal{D} \cup \{(s_t, a_t, s_{t+1})\}$.

Model Training.

Update GP world model with \mathcal{D} .

Update $s_t = s_{t+1}$, and $t = t + 1$.

end

Increase the training trials $k = k + 1$.

end

Generate final decision.

Initialize time step $t = 1$, state s_t , and policy $\pi = []$.

for $t \leq T$ **do**

Generate action $a_t = \arg \max_{a \in \mathcal{A}} Q(s_t, a)$ by MCTS with **mean reward** R .

Predict $s_{t+1} = \mathbb{E}[s_{t+1}|s_t, a_t]$ with Gaussian Process.

Update $s_t = s_{t+1}$, and $t = t + 1$.

Append action to final decision $\pi = \pi \oplus a_t$.

end

Return the malaria control policy π .

Theorem 1. Assume that the feature space $\mathcal{X} \subset \mathbb{R}^m$ of state-action pairs is a compact domain and the target value is bounded $y \in [0, V_m]$. The reward and action value are Lipschitz continuous w.r.t the state-action pairs and let L_q and L_r be the Lipschitz constants for reward and action value respectively. If VB-MCTS is executed with $\beta_1 = \frac{L_r}{2\sigma_n^2}$ and $\beta_2 = \frac{L_p}{2\sigma_n^2}$ for any MDP M , then with probability $1 - \delta$, VB-MCTS will follow a 4ϵ -optimal policy from its current state on all but

$$O\left(\frac{V_m \zeta(\epsilon, \delta)}{\epsilon(1-\gamma)} \log \frac{1}{\delta} \log \frac{1}{\epsilon(1-\gamma)}\right) \quad (2)$$

timesteps, with probability at least $1 - 2\delta$, where $\zeta(\epsilon, \delta) = \left(\frac{4V_m^2}{\epsilon^2(1-\gamma)^2} \log \left(\frac{2|\mathcal{N}_{\mathcal{X}}(\tau(\sigma_{tol}^2))|}{\delta}\right)\right) \mathcal{N}_{\mathcal{X}}(\tau(\sigma_{tol}^2))$, $\sigma_{tol}^2 = \frac{2\omega_n^2 \epsilon_1^2}{V_m^2 \log(\frac{2}{\delta_1})}$. $\mathcal{N}_{\mathcal{X}}(\tau)$ is the covering number of domain \mathcal{X} , which defines the cardinality of the minimal set $C = \{c_1, \dots, c_{N_c}\}$ s.t. $\forall x \in \mathcal{X}, \exists c_j \in C$ s.t. $d(x, c_j) \leq \tau$, where $d(\cdot, \cdot)$ is a distance measure.

Sketch of Proof. The proof is organized by showing the required key properties (optimism, accuracy, and learning complexity) for general PAC-MDP [Strehl *et al.*, 2009] are satisfied.

Optimism. Assume that MCTS can find the optimal decision under every state s with enough computation resources. Then, $\forall (s_t, a_t)$, the optimal action value $Q^*(s_t, a_t)$ satisfies

$$\begin{aligned} Q^*(s_t, a_t) &= \arg \max_{a_{t+1} \in \mathcal{A}} R(s_{t+1}) + \gamma Q^*(s_{t+1}, a_{t+1}) \\ &\leq R(s'_{t+1}) + \frac{L_r}{\omega_n^2} \text{Var}_{s'_{t+1} \sim GP(s_t, a_t)} [R(s'_{t+1})] \\ &\quad + \gamma \tilde{Q}^*(s'_{t+1}, a_{t+1}) + \gamma \frac{L_q}{\omega_n^2} \text{Var}_{s'_{t+1} \sim GP(s_t, a_t)} [s'_{t+1}] \\ &\leq \arg \max_{a_{t+1} \in \mathcal{A}} R(s'_{t+1}) + \frac{L_r}{\omega_n^2} \text{Var}_{s'_{t+1} \sim GP(s_t, a_t)} [R(s'_{t+1})] \\ &\quad + \gamma \tilde{Q}^*(s'_{t+1}, a'_{t+1}) + \gamma \frac{L_q}{\omega_n^2} \text{Var}_{s'_{t+1} \sim GP(s_t, a_t)} [s'_{t+1}] \\ &= \tilde{Q}(s_t, a_t), \end{aligned}$$

where $s_{t+1} = P(s_t, a_t)$, $s'_{t+1} = \mathbb{E}_{s_{t+1} \sim GP(s_t, a_t)} [s_{t+1} | s_t, a_t]$. The maximum error of propagating $Q(s_{t+1}, a_{t+1})$ instead of $Q(s'_{t+1}, a_{t+1})$ is given by $\|s'_{t+1} - s_{t+1}\| L_q$. The norm $\|s'_{t+1} - s_{t+1}\|$ is upper bounded by the regularization error $\frac{L_q}{\omega_n^2} \text{Var}_{s'_{t+1} \sim GP(s_t, a_t)} [s'_{t+1}]$ with probability of $(1 - \delta)$ (Section A.2 in [Grande, 2014]). In the same way, we have $R(s'_{t+1}) + \frac{L_r}{\omega_n^2} \text{Var}_{s'_{t+1} \sim GP(s_t, a_t)} [R(s'_{t+1})] - R(s_{t+1}) \geq 0$ with probability $1 - \delta$. It follows that Algorithm 1 remains optimistic with probability at least $1 - 2\delta$.

Accuracy. Following the Lemma 1 from [Grande *et al.*, 2014], the prediction error at $x_i \in \mathcal{X}$ is bounded in probability $Pr \{ |\hat{\mu}(x_i) - f(x_i)| \geq \epsilon_1 \} \leq \delta_1$ if the predictive variance of the GP at $\sigma_n^2(x_i) \leq \sigma_{\text{tol}}^2 = \frac{2\omega_n^2 \epsilon_1^2}{V_m \log(\frac{2}{\delta_1})}$.

Learning complexity. (a) For $\forall x \in \mathcal{X}$ with $\sigma_0^2 = 1$, let $\rho = \min_{j=1, \dots, n} k(x, x_j)$, $d_{\max} = \max_{j=1, \dots, n} d(x, x_j)$, then d_{\max} can be rewritten as $d_{\max} = \sqrt{2\theta^2 \log \frac{1}{\rho}}$ with $\rho = e^{-\frac{d_{\max}^2}{2\theta^2}}$.

If $n \geq \frac{\omega_n^2}{\sigma_{\text{tol}}^2 - 1 + \rho^2}$, the posterior variance of x satisfies $\sigma_n^2(x) \leq 1 - \frac{n\rho^2}{n + \omega_n^2} \leq \frac{n(1 - \rho^2) + \omega_n^2}{n} \leq \sigma_{\text{tol}}^2$. (b) Since \mathcal{X} is a compact domain with the length of i -th ($i \in [1, m]$) dimension as L_i , we have $\mathcal{N}_{\mathcal{X}}(d_{\max}) \leq \mathcal{N}_{\mathcal{X}}(\frac{d_{\max}}{2})$ and \mathcal{X} can be covered by $2^m \prod_{j=1}^m L_j / d_{\max}^m$ balls $\{B(c_j, \frac{d_{\max}}{2})\}_j$, which centers at c_j with a radius of $\frac{d_{\max}}{2}$.

Given (a) and (b), $\forall x \in \mathcal{X}$, $\exists c_j \in \mathcal{X}$ s.t. $B(x, d_{\max}) \supset B(x, d_{\max}/2)$. If there are at least $\frac{\omega_n^2}{\sigma_{\text{tol}}^2 - 1 + \rho^2}$ observations in $B(x, d_{\max}/2)$, then (2) satisfied. Combining the Lemma 8 from [Strehl *et al.*, 2009] and denoting $\tau(\sigma_{\text{tol}}^2) = d_{\max}/2$, the total number of updates occurs will be bounded by $\zeta(\epsilon, \delta) = \left(\frac{4V_m^2}{\epsilon^2(1-\gamma)^2} \log \left(\frac{2|\mathcal{N}_{\mathcal{X}}(\tau(\sigma_{\text{tol}}^2))|}{\delta} \right) \right) \mathcal{N}_{\mathcal{X}}(\tau(\sigma_{\text{tol}}^2))$ with probability of $1 - \delta$.

Now that the key properties of optimism, accuracy, and learning complexity have been established, the general PAC-MDP Theorem 10 of [Strehl *et al.*, 2009] is invoked. \square

Theorem 1 allows us to guarantee that the number of steps in which the performance of VB-MCTS is significantly worse

than that of an optimal policy starting from the current state is at most log-linear in the covering number of the state-action space with probability $1 - \delta$.

4 Experiments

In this section, we conduct extensive experiments on two different OpenMalaria [Smith *et al.*, 2008] based simulators: *SeqDecChallenge* and *ProveChallenge*, which are the testing environments used in KDD Cup 2019. In these simulators, the parameters of the simulator are hidden from the RL agents since the “simulation parameters” for SSA are unknown for policymakers. Additionally, to simulate the real disease control problem, only **20 trials** were allowed before generating the final decision, which is much more challenging than traditional RL tasks. These simulating environments are available at <https://github.com/IBM/ushiriki-policy-engine-library>.

Agents for Comparison. To show the advantage of VB-MCTS, many benchmarking reinforcement learning methods and open source solutions have been deployed to verify its effectiveness: **Random Policy:** The random policy is executed in 20 trials and chooses the generated policy with the maximum reward as the final decision. **SMAB:** This kind of policy treats the problem as a Stochastic Multi-Armed Bandit problem and independently optimizes the policy every year with Thompson sampling [Chapelle and Li, 2011]. **CEM:** Cross-Entropy Method is a simple gradient-free policy searching method [Szita and Lőrincz, 2006]. **CMA-ES:** CMA-ES is a gradient-free evolutionary approach to optimizing non-convex objective functions [Krause *et al.*, 2016]. **Q-learning-GA:** It learns the malaria control policy by combining Q-learning and Genetic Algorithm. **Expected-Sarsa:** It collects 13 random episodes and runs expected value SARSA [Van Seijen *et al.*, 2009] for 7 episodes to improve the best policy using the collected statistics. **GP-Rmax:** It uses GP learners to model T and R and replaces the value of any $Q^*(s, a)$ where $T(s, a)$ is “unknown” with a value of $\frac{R_{\max}}{1-\gamma}$ [Li *et al.*, 2011; Grande *et al.*, 2014]. **GP-MC:** It employs Gaussian Process to regress the world model. The policy is generated by sampling from the posterior and choosing the max rewarded action. **VB-MCTS:** Our proposed method.

Implementation Details. We build a feature map of 14-dimension for this task, which includes the periodic feature and cross term feature. Specifically, the feature map is set as,

$$\begin{aligned} \phi(s_t, a_t) &= [t, t\%2, t\%3, r_{t-1}, a_{(t-1, ITN)}, a_{(t-1, IRS)}, \\ &\quad a_{(t, ITN)}, a_{(t, IRS)}, a_{(t, ITN)} \times a_{(t, IRS)}, \\ &\quad a_{(t-1, ITN)} \times a_{(t-1, IRS)}, a_{(t, ITN)} \times a_{(t-1, ITN)}, \\ &\quad a_{(t, IRS)} \times a_{(t-1, IRS)}, a_{(t, ITN)} \times (1 - a_{(t-1, ITN)}), \\ &\quad a_{(t, IRS)} \times (1 - a_{(t-1, IRS)})], \end{aligned}$$

where $t, t\%2, t\%3$ are the periodic features and $a_{(*,*)} \times a_{(*,*)}$ is the cross term feature. Since the predicted variance of state and reward are the same in our setting, we empirically set the sum of exploration/exploitation parameters as $\beta_1 + \beta_2 = 3.5$ in the experiments. In MCTS, c_{puct} is 5, and only the top 50 rewarded child nodes are expanded. The number of iterations does not exceed 100,000. For Gaussian Process,

		<i>SeqDecChallenge</i>		
Agents		Med. Reward	Max Reward	Min Reward
Random Policy		167.79	193.24	135.06
SMAB		209.05	386.28	-6.44
CEM		179.30	214.87	120.92
CMA-ES		185.34	246.12	108.18
Q-Learning-GA		247.75	332.40	171.33
Expected-Sarsa		462.76	495.03	423.93
GP-Rmax		233.95	292.99	200.35
GP-MC		475.99	499.60	435.51
VB-MCTS		533.38	552.78	519.61
		<i>ProveChallenge</i>		
Agents		Med. Reward	Max Reward	Min Reward
Random Policy		248.25	464.92	55.24
SMAB		18.02	135.37	-56.86
CEM		229.61	373.83	20.09
CMA-ES		289.03	314.57	92.95
Q-Learning-GA		242.97	325.24	88.70
Expected-Sarsa		190.08	296.16	140.86
GP-Rmax		287.45	371.49	153.98
GP-MC		300.37	447.15	263.96
VB-MCTS		352.17	492.23	259.97

Table 1: Performance comparisons between different agents.

to avoid overfitting problems, 5-fold cross-validation is performed during the updates of the GP world model. Particularly, we use 1-fold for training and 4-fold for validation, which ensures the generalizability of the GP world model over different state-action pairs. Our implementation and all baseline codes are available at <https://github.com/zoulixin93/VB-MCTS>.

4.1 Results

Main Results. In Table 1, we report the median reward, the maximum reward, and the minimal reward over 10 independent repeat runs. The results are quite consistent with our intuition. We have the following observations: **(1)** For the finite-horizon decision-making problem, treating it as SMAB does not work, and the delayed influence of actions can not be ignored in malaria control. As presented in Table 1, SMAB’s performances have large variance and are even worse than random policy in *ProveChallenge*. **(2)** Overall, two model-based methods (GP-MC and VB-MCTS) consistently outperform the model-free methods (CEM, CMA-ES, Q-learning-GA, and Expected-Sarsa) in *SeqDecChallenge* and *ProveChallenge*, which indicates that empirically model-based solutions are generally more data-efficient than model-free solutions. From the results in Table 1, performances of model-free methods are defeated by model-based methods with a large margin in *SeqDecChallenge*, and their performances are almost the same as the random policy in *ProveChallenge*. **(3)** The proposed VB-MCTS can outperform all the baselines in *SeqDecChallenge* and *ProveChallenge*. Compared with GP-MC and GP-Rmax, the advantage from the efficient MCTS with variance-bonus reward leads to the success on *SeqDecChallenge* and *ProveChallenge*.

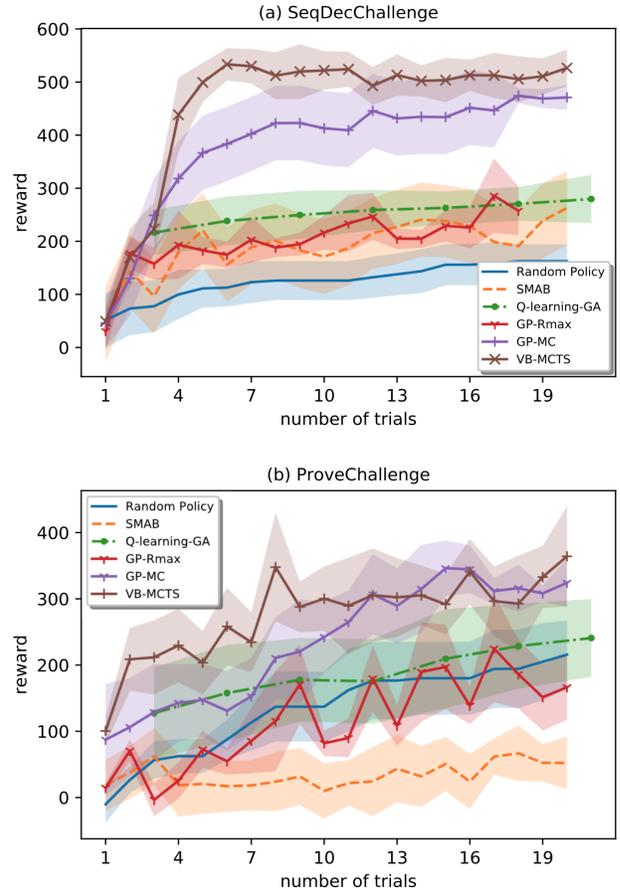


Figure 4: Malaria policy learning curves in five different agents.

Data Efficiency. This paragraph compares data-efficiency (required trials) of VB-MCTS with other RL methods that learn malaria policies from scratch. In Figure 4(a) and 4(b), we report agents’ performances after collecting every trial episode in *SeqDecChallenge* and *ProveChallenge*. The horizontal axis indicates the number of trials. The vertical axis shows the average performance after collecting every episode. Figure 4(a) and 4(b) highlighted that our proposed VB-MCTS approach (brown) requires on average only 8 trials to achieve the best performances in *SeqDecChallenge* and *ProveChallenge*, including the first random trial. The results indicate that VB-MCTS can outperform the state-of-the-art method on both data efficiency and performance.

5 Conclusion

We proposed a model-based approach employing the Gaussian Process to regress the state transition for data-efficient RL in malaria control. By planning with variance-bonus reward, our method can naturally deal with the dilemma of exploration and exploitation by efficiently MCTS planning. Extensive experiments conducted on the challenging malaria control task have demonstrated the advantage of VB-MCTS over state-of-the-arts both on performance and efficiency.

References

- [Auer and Ortner, 2010] Peter Auer and Ronald Ortner. Ucb revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Periodica Mathematica Hungarica*, 61(1-2):55–65, 2010.
- [Bent *et al.*, 2018] Oliver Bent, Sekou L Remy, Stephen Roberts, and Aisha Walcott-Bryant. Novel exploration techniques (nets) for malaria policy interventions. In *AAAI*, 2018.
- [Chapelle and Li, 2011] Olivier Chapelle and Lihong Li. An empirical evaluation of thompson sampling. In *NIPS*, 2011.
- [Deisenroth and Rasmussen, 2011] Marc Deisenroth and Carl E Rasmussen. Pilco: A model-based and data-efficient approach to policy search. In *ICML*, 2011.
- [Grande *et al.*, 2014] Robert Grande, Thomas Walsh, and Jonathan How. Sample efficient reinforcement learning with gaussian processes. In *ICML*, 2014.
- [Grande, 2014] Robert Conlin Grande. *Computationally efficient gaussian process changepoint detection and regression*. PhD thesis, Massachusetts Institute of Technology, 2014.
- [Ha and Schmidhuber, 2018] David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. In *NeurIPS*, 2018.
- [Holland, 1992] John H Holland. Genetic algorithms. *Scientific american*, 267(1):66–73, 1992.
- [Kakade *et al.*, 2003] Sham Kakade, Michael J Kearns, and John Langford. Exploration in metric state spaces. In *ICML*, 2003.
- [Kamthe and Deisenroth, 2017] Sanket Kamthe and Marc Peter Deisenroth. Data-efficient reinforcement learning with probabilistic model predictive control. *arXiv:1706.06491*, 2017.
- [Kolter and Ng, 2009] J Zico Kolter and Andrew Y Ng. Near-bayesian exploration in polynomial time. In *ICML*, 2009.
- [Krause *et al.*, 2016] Oswin Krause, Dídac Rodríguez Arbonès, and Christian Igel. Cma-es with optimal covariance update and storage complexity. In *NIPS*, 2016.
- [Li *et al.*, 2011] Lihong Li, Michael L Littman, Thomas J Walsh, and Alexander L Strehl. Knows what it knows: a framework for self-aware learning. *Machine learning*, 82(3):399–443, 2011.
- [Mnih *et al.*, 2013] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv:1312.5602*, 2013.
- [Moran, 2007] Mary Moran. *The malaria product pipeline: planning for the future*. Health Policy Division, The George Institute for International Health, 2007.
- [Silver *et al.*, 2017] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354, 2017.
- [Smith *et al.*, 2008] T Smith, N Maire, A Ross, M Penny, N Chitnis, A Schapira, A Studer, B Genton, C Lengeler, Fabrizio Tediosi, et al. Towards a comprehensive simulation model of malaria epidemiology and control. *Parasitology*, 135(13):1507–1516, 2008.
- [Sorg *et al.*, 2012] Jonathan Sorg, Satinder Singh, and Richard L Lewis. Variance-based rewards for approximate bayesian reinforcement learning. *arXiv:1203.3518*, 2012.
- [Srinivas *et al.*, 2009] Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv:0912.3995*, 2009.
- [Strehl *et al.*, 2009] Alexander L Strehl, Lihong Li, and Michael L Littman. Reinforcement learning in finite mdps: Pac analysis. *JMRL*, 2009.
- [Stuckey *et al.*, 2014] Erin M Stuckey, Jennifer Stevenson, Katya Galactionova, Amrsh Y Baidjoe, Teun Bousema, Wycliffe Odongo, Simon Kariuki, Chris Drakeley, Thomas A Smith, Jonathan Cox, et al. Modeling the cost effectiveness of malaria control interventions in the highlands of western kenya. *PLoS one*, 9(10):e107700, 2014.
- [Sutton *et al.*, 2000] Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *NIPS*, 2000.
- [Szita and Lörincz, 2006] István Szita and András Lörincz. Learning tetris using the noisy cross-entropy method. *Neural computation*, 18(12):2936–2941, 2006.
- [Van Seijen *et al.*, 2009] Harm Van Seijen, Hado Van Hasselt, Shimon Whiteson, and Marco Wiering. A theoretical and empirical analysis of expected sarsa. In *DPRL*, 2009.
- [Williams and Rasmussen, 2006] Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.
- [Zhou *et al.*, 2020] Wenjun Zhou, Taposh Dutta Roy, and Iryna Skrypnik. The kdd cup 2019 report. *ACM SIGKDD Explorations Newsletter*, 22(1):8–17, 2020.
- [Zou *et al.*, 2019] Lixin Zou, Long Xia, Zhuoye Ding, Jiaying Song, Weidong Liu, and Dawei Yin. Reinforcement learning to optimize long-term user engagement in recommender systems. In *SIGKDD*, 2019.
- [Zou *et al.*, 2020a] Lixin Zou, Long Xia, Pan Du, Zhuo Zhang, Ting Bai, Weidong Liu, Jian-Yun Nie, and Dawei Yin. Pseudo dyna-q: A reinforcement learning framework for interactive recommendation. In *WSDM*, 2020.
- [Zou *et al.*, 2020b] Lixin Zou, Long Xia, Yulong Gu, Xiangyu Zhao, Weidong Liu, Jimmy Xiangji Huang, and Dawei Yin. Neural interactive collaborative filtering. In *SIGIR*, 2020.