
Efficient Rollout Strategies for Bayesian Optimization

Eric Hans Lee

Cornell University, CS

erichanslee@cs.cornell.edu

David Eriksson

Uber AI

eriksson@uber.com

Bolong Cheng

SigOpt

harvey@sigopt.com

Michael McCourt

SigOpt

mike@sigopt.com

David Bindel

Cornell University, CS

bindel@cs.cornell.edu

Abstract

Bayesian optimization (BO) is a class of sample-efficient global optimization methods, where a probabilistic model conditioned on previous observations is used to determine future evaluations via the optimization of an acquisition function. Most acquisition functions are myopic, meaning that they only consider the impact of the next function evaluation. Non-myopic acquisition functions consider the impact of the next h function evaluations and are typically computed through rollout, in which h steps of BO are simulated. These rollout acquisition functions are defined as h -dimensional integrals, and are expensive to compute and optimize. We show that a combination of quasi-Monte Carlo, common random numbers, and control variates significantly reduce the computational burden of rollout. We then formulate a policy-search based approach that removes the need to optimize the rollout acquisition function. Finally, we discuss the qualitative behavior of rollout policies in the setting of multimodal objectives and model error.

1 INTRODUCTION

Bayesian optimization (BO) is a class of methods for global optimization used to minimize expensive black-box functions. BO builds a probabilistic surrogate model of the objective and then determines future evaluations via an acquisition function. Applications of BO include robotic gait control, sensor set selection, and neural network hyperparameter tuning [21, 22, 3]. BO is favored in these tasks because of its sample-efficient nature. Achieving this sample-efficiency demands that BO balance exploration and exploitation. However, standard acquisition functions such as expected improvement (EI) are too

greedy and perform little exploration. As a result, they perform poorly on multimodal problems [7] and have provably sub-optimal performance in certain settings, e.g., bandit problems [23]. A key research goal in BO is developing less greedy acquisition functions [21]. Examples include predictive entropy search (PES) [7] or knowledge gradient (KG) [4]. Lam et al. [15] frame the exploration-exploitation trade-off as a balance between immediate and future rewards in a continuous state and action space Markov decision process (MDP). In this framework, *non-myopic* acquisition functions are optimal MDP policies, and promise better performance by considering the impact of future evaluations up to a given BO budget (also referred to as the *horizon*).

Unfortunately, maximizing the MDP reward is an intractable problem as it involves solving an infinite-dimensional dynamic program [18]. Rollout is a popular class of approximate solutions in which future BO realizations and their corresponding values are simulated using the surrogate and then averaged. This average defines a rollout acquisition function, which is considered the state-of-the-art in non-myopic BO. While more practical than the original MDP problem, rollout acquisition functions are still computationally expensive to the extent that suggesting the next evaluation can take several hours [24]. This paper aims to make non-myopic BO more practical by reducing the time to suggest the next point from hours to seconds. In particular, our main contributions are:

- We compute rollout acquisition functions via Monte Carlo integration, and use variance reduction techniques—quasi-Monte Carlo, common random numbers, and control variates—to decrease the estimation error by several orders of magnitude.
- We introduce a coarser approximation to rollout acquisition functions through policy search. In this setting, we use rollout to select the best acquisition function from a given set at each BO iteration. This removes the need to optimize the rollout acquisition

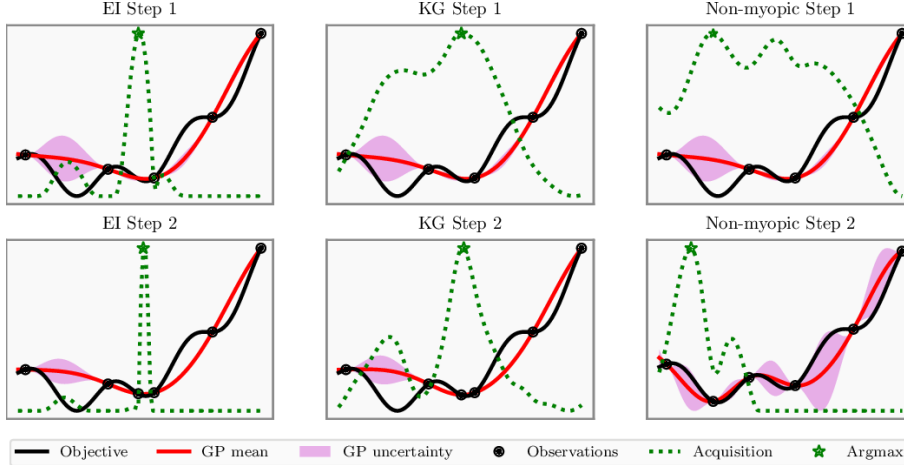


Figure 1: Comparing EI (*left*), KG (*middle*), and a non-myopic acquisition function (*right*) on a carefully chosen objective. We observe five values of $f(x) = \sin(20x) + 20(x - 0.3)^2$. The GP has one promising minimum, but also a region of high uncertainty. For each acquisition function, we perform two steps of BO. EI will ignore the left region and instead greedily evaluate twice in a sub-optimal location. KG is less greedy, but will nonetheless evaluate similarly to EI. The non-myopic acquisition will properly evaluate the region of high uncertainty and identify the global minimum.

function, and makes rollout much more practical.

- We provide experimental results for both rollout acquisitions and policy search. The former suggests that rollout acquisition functions perform better on multimodal problems. The latter shows that policy search performs as well as, if not better than, the best-performing one-step acquisition function.
- We examine the impact of model mis-specification on performance and show that increasing the rollout horizon also increases sensitivity to model error.

2 BACKGROUND

Many papers have been published in the sub-field of non-myopic BO over the last few years [5, 6, 14–16, 24, 25]. Most of this recent research concerns rollout, in which future realizations of BO are simulated over horizon h using the surrogate and averaged to determine the acquisition function. Rollout acquisition functions represent the state-of-the-art in BO and are integrals over h dimensions, where the integrand itself is evaluated through inner optimizations, resulting in an expensive integral. The rollout acquisition function is then maximized to determine the next BO evaluation, further increasing the cost. This large computational overhead is evidenced by Osborne et al. [16], who are only able to compute the rollout acquisition for horizon 2, dimension 1; and later by Lam et al. [15], who use Gauss-Hermite quadrature in horizons up to five and see runtimes on the order of hours for small, synthetic functions [24].

Recent work focuses on making rollout more practical. Wu and Frazier [24] consider horizon two and use a combination of Gauss-Hermite quadrature and Monte Carlo (MC) integration to quickly calculate the acquisition function and its gradient. Non-myopic active learning also uses rollout and recent work develops a fast implementation by truncating the horizon and selecting a batch of points to collect future rewards [10, 11].

Gaussian process regression and BO: Suppose we seek a global minimum of a continuous objective $f(\mathbf{x})$ over a compact set $\Omega \subseteq \mathbb{R}^d$. If $f(\mathbf{x})$ is expensive to evaluate, finding a minimum should ideally be sample-efficient. BO often uses a Gaussian process (GP) to model $f(\mathbf{x})$ from the data $\mathcal{D}_k = \{(\mathbf{x}_i, y_i)\}_{i=1}^k$. The next evaluation location \mathbf{x}_{k+1} is determined by maximizing an acquisition function $\Lambda(\mathbf{x} | \mathcal{D}_k)$: $\mathbf{x}_{k+1} = \arg \max_{\Omega} \Lambda(\mathbf{x} | \mathcal{D}_k)$.

We place a GP prior on $f(\mathbf{x})$, denoted by $f \sim \mathcal{GP}(\mu, K)$, where $\mu : \Omega \rightarrow \mathbb{R}$ and $K : \Omega \times \Omega \rightarrow \mathbb{R}$ are the mean function and covariance kernel, respectively (see the supplement for examples). The kernel $K(\mathbf{x}, \mathbf{x}')$ correlates neighboring points, and may contain *hyperparameters*, such as lengthscales that are learned to improve the quality of approximation [20]. For a given \mathcal{D}_k , we define:

$$\mathbf{y} = \begin{pmatrix} y_1 \\ \vdots \\ y_k \end{pmatrix}, \quad \mathbf{k}(\mathbf{x}) = \begin{pmatrix} K(\mathbf{x}, \mathbf{x}_1) \\ \vdots \\ K(\mathbf{x}, \mathbf{x}_k) \end{pmatrix}, \quad \mathbf{K} = \begin{pmatrix} K(\mathbf{x}_1, \mathbf{x}_1) \\ \vdots \\ K(\mathbf{x}_k, \mathbf{x}_k) \end{pmatrix}.$$

We assume y_i is observed with Gaussian white noise: $y_i = f(\mathbf{x}_i) + \epsilon_i$, where $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$. Given a GP prior and data \mathcal{D}_k , the resulting posterior distribution for




Figure 2: We calculate a standard non-myopic acquisition function in 1D with different values of h using $50h$ samples. (Top) The results of a standard MC estimator. (Bottom) The results of our estimator look far less noisy.

function values at a location \mathbf{x} is the Normal distribution $\mathcal{N}(\mu^{(k)}(\mathbf{x}; \mathcal{D}_k), K^{(k)}(\mathbf{x}, \mathbf{x}; \mathcal{D}_k))$:

$$\begin{aligned} \mu^{(k)}(\mathbf{x}; \mathcal{D}_k) &= \mu(\mathbf{x}) + \mathbf{k}(\mathbf{x})^\top (K + \sigma^2 \mathbf{I}_k)^{-1} (\mathbf{y} - \mu(\mathbf{x})), \\ K^{(k)}(\mathbf{x}, \mathbf{x}; \mathcal{D}_k) &= K(\mathbf{x}, \mathbf{x}) - \mathbf{k}(\mathbf{x})^\top (K + \sigma^2 \mathbf{I}_k)^{-1} \mathbf{k}(\mathbf{x}), \end{aligned}$$

where \mathbf{I}_k is the $k \times k$ identity matrix.

Non-myopic Bayesian optimization: Non-myopic BO frames the exploration-exploitation trade-off as a balance of immediate and future rewards. The strength of this approach is demonstrated in Figure 1, in which we compare EI and KG to a non-myopic acquisition function on a carefully chosen objective. Its GP has a region of uncertainty on the left containing the global minimum and a local minimum at the origin. We run two steps of BO, updating the posterior each step. EI and KG behave greedily by sampling twice near the origin, while the non-myopic approach uses one evaluation to explore the uncertain region, subsequently identifying the global minimum and converging faster than either EI or KG.

Lam et al. [15] formulate non-myopic BO as a finite horizon dynamic program. We present the equivalent Markov decision process (MDP) formulation. We use standard notation from Puterman [19]: an MDP is the collection $\langle T, \mathbb{S}, \mathbb{A}, P, R \rangle$. $T = \{0, 1, \dots, h-1\}$, $h < \infty$ is the set of decision epochs, assumed finite for our problem. The state space, \mathbb{S} , encapsulates all the information needed to model the system from time $t \in T$. \mathbb{A} is the action space. Given a state $s \in \mathbb{S}$ and an action $a \in \mathbb{A}$, $P(s'|s, a)$ is the transition probability of the next state being s' . $R(s, a, s')$ is the reward received for choosing action a from state s , and ending in state s' .

A decision rule, $\pi_t : \mathbb{S} \rightarrow \mathbb{A}$, maps states to actions at time t . A policy π is a series of decision rules $\pi =$

$(\pi_0, \pi_1, \dots, \pi_{h-1})$, one at each decision epoch. Given a policy π , a starting state s_0 , and horizon h , we can define the expected total reward $V_h^\pi(s_0)$ as:

$$V_h^\pi(s_0) = \mathbb{E} \left[\sum_{t=0}^{h-1} R(s_t, \pi_t(s_t), s_{t+1}) \right].$$

In phrasing a sequence of decisions as an MDP, our goal is to find the optimal policy π^* that maximizes the expected total reward, i.e., $\sup_{\pi \in \Pi} V_h^\pi(s_0)$, where Π is the space of all admissible policies.

If we can sample from the transition probability P , we can estimate the expected total reward of any *base policy* $\tilde{\pi}$ with MC integration:

$$V_h^{\tilde{\pi}}(s_0) \approx \frac{1}{N} \sum_{i=1}^N \left[\sum_{t=0}^{h-1} R(s_t^i, \tilde{\pi}_t(s_t^i), s_{t+1}^i) \right].$$

Given a GP prior over data \mathcal{D}_t with mean $\mu^{(t)}$ and kernel $K^{(t)}$, we model h steps of BO as an MDP. This MDP's state space is all possible data sets reachable from starting state \mathcal{D}_t with h steps of BO. Its action space is Ω ; actions correspond to sampling a point in Ω . Its transition probability and reward function are defined as follows. Given an action \mathbf{x}_{t+1} , the transition probability from \mathcal{D}_t to \mathcal{D}_{t+1} , where $\mathcal{D}_{t+1} = \mathcal{D}_t \cup \{(\mathbf{x}_{t+1}, y_{t+1})\}$ is:

$$\begin{aligned} P(\mathcal{D}_t, \mathbf{x}_{t+1}, \mathcal{D}_{t+1}) &\sim \\ \mathcal{N}(\mu^{(t)}(\mathbf{x}_{t+1}; \mathcal{D}_t), K^{(t)}(\mathbf{x}_{t+1}, \mathbf{x}_{t+1}; \mathcal{D}_t)). \end{aligned}$$

In other words, the transition probability from \mathcal{D}_t to \mathcal{D}_{t+1} is the probability of sampling y_{t+1} from the posterior $\mathcal{GP}(\mu^{(t)}, K^{(t)})$ at \mathbf{x}_{t+1} . We define a reward according to EI [12]. Let y_t^* be the minimum observed value in the

observed set \mathcal{D}_t , i.e., $y_t^* = \min\{y_0, \dots, y_t\}$. Then our reward is expressed as:

$$R(\mathcal{D}_t, \mathbf{x}_{t+1}, \mathcal{D}_{t+1}) = (y_t^* - y_{t+1})^+ \equiv \max(y_t^* - y_{t+1}, 0).$$

EI can be defined as the optimal policy for horizon one, obtained by maximizing the immediate reward:

$$\begin{aligned} \pi_{EI} &= \arg \max_{\pi} V_1^{\pi}(\mathcal{D}_k) \\ &= \arg \max_{\mathbf{x} \in \Omega} \mathbb{E} \left[(y_k^* - y_{\mathbf{x}})^+ \right] \equiv \arg \max_{\mathbf{x} \in \Omega} \text{EI}(\mathbf{x} \mid \mathcal{D}_k), \end{aligned}$$

where the starting state is \mathcal{D}_k . In contrast, we define the non-myopic policy as the optimal solution to an h -horizon MDP. The expected total reward of this MDP can be expressed as:

$$\begin{aligned} V_h^{\pi}(\mathcal{D}_k) &= \mathbb{E} \left[\sum_{t=k}^{k+h-1} R(\mathcal{D}_t, \pi_t(\mathcal{D}_t), \mathcal{D}_{t+1}) \right] \\ &= \mathbb{E} \left[\sum_{t=k}^{k+h-1} (y_t^* - y_{t+1})^+ \right]. \end{aligned}$$

For $h > 2$, the optimal policy is difficult to compute.

Rollout acquisition functions: In the context of BO, rollout policies [1], which are sub-optimal but yield promising results, are a tractable alternative to optimal policies [24]. For a given current state \mathcal{D}_k , we denote our base policy $\tilde{\pi} = (\tilde{\pi}_0, \tilde{\pi}_1, \dots, \tilde{\pi}_{h-1})$. We introduce the notation $\mathcal{D}_{k,0} \equiv \mathcal{D}_k$ to define the initial state of our MDP and $\mathcal{D}_{k,t}$ for $1 \leq t \leq h$ to denote the random variable that is the state at each decision epoch. Each individual decision rule $\tilde{\pi}_t$ consists of maximizing the base acquisition function Λ given the current state $s_t = \mathcal{D}_{k,t}$,

$$\tilde{\pi}_t = \arg \max_{\mathbf{x} \in \Omega} \Lambda(\mathbf{x} \mid \mathcal{D}_{k,t}).$$

Using this policy, we define the non-myopic acquisition function $\Lambda_h(\mathbf{x})$ as the rollout of $\tilde{\pi}$ to horizon h i.e., the expected reward of $\tilde{\pi}$ starting with the action $\tilde{\pi}_0 = \mathbf{x}$:

$$\Lambda_h(\mathbf{x}_{k+1}) := \mathbb{E} \left[V_h^{\tilde{\pi}}(\mathcal{D}_k \cup \{(\mathbf{x}_{k+1}, y_{k+1})\}) \right],$$

where y_{k+1} is the noisy observed value of f at \mathbf{x}_{k+1} . Λ_h is better than Λ in expectation for a correctly specified GP prior and for any acquisition function. This follows from standard results in the MDP literature:

Definition 2.1 Bertsekas [1]: A policy π is sequentially consistent if, for every trajectory generated from any s_0 :

$$(s_0, a_0), (s_1, a_1), \dots, (s_{h-1}, a_{h-1}),$$

π generates the following trajectory starting at s_1 :

$$(s_1, a_1), (s_2, a_2) \dots, (s_{h-1}, a_{h-1}).$$

Sequential consistency requires the trajectory generated from applying π at s_i for horizon $h - i$ to be a sub-trajectory of the trajectory generated from applying π at s_0 for horizon h . Acquisition functions are sequentially consistent so long as we consistently break ties if they have multiple maxima —though this will not occur generically. Sequential consistency guarantees that rollout does at least as well as its base policy in expectation:

Theorem 2.1 Bertsekas [1]: A rollout policy π_{roll} does as least as well as its base policy $\tilde{\pi}$ in expectation if $\tilde{\pi}$ is sequentially consistent i.e.,

$$V_h^{\pi_{roll}}(s_0) \geq V_h^{\tilde{\pi}}(s_0).$$

Thus, rolling out any acquisition function will do at least as well in expectation as the acquisition function itself.

Unfortunately, while rollout is tractable and conceptually straightforward, it is computationally demanding. To rollout $\tilde{\pi}$ once, we must run h steps of BO with Λ . Many such rollouts must then be averaged to reasonably estimate Λ_h , which is an h -dimensional integral. Estimation can be done either through explicit quadrature or MC integration, and is the primary computational bottleneck of rollout. Our paper significantly lowers the computational burden of rollout through MC variance reduction and a fast policy search method that avoids optimizing Λ_h . We detail these methods in the next section.

3 METHODS

In the context of rollout, MC estimates the expected reduction over h steps of BO using base policy $\tilde{\pi}$:

$$\begin{aligned} \Lambda_h(\mathbf{x}_{k+1}) &= \mathbb{E} \left[V_h^{\tilde{\pi}}(\mathcal{D}_k \cup \{(\mathbf{x}_{k+1}, y_{k+1})\}) \right] \\ &\approx \frac{1}{N} \sum_{i=1}^N \sum_{t=k}^{k+h-1} (y_t^* - y_{t+1})^+. \end{aligned}$$

The distribution for y_{t+1} is a normal distribution whose mean and variance are determined by rolling out $\tilde{\pi}$ to horizon t and examining the posterior GP:

$$\begin{aligned} y_{t+1} &\sim \mathcal{N}(\mu^{(t)}(\mathbf{x}_{t+1}; \mathcal{D}_t), K^{(t)}(\mathbf{x}_{t+1}, \mathbf{x}_{t+1}; \mathcal{D}_t)), \\ \mathbf{x}_{t+1} &= \pi(\mathcal{D}_t) = \arg \max_{\mathbf{x} \in \Omega} \Lambda(\mathbf{x} \mid \mathcal{D}_t). \end{aligned} \quad (1)$$

A sample path in this context may be represented as the sequence $(\mathbf{x}_k, y_k), (\mathbf{x}_{k+1}, y_{k+1}), \dots, (\mathbf{x}_{k+h}, y_{k+h})$ produced by Equation 1. We parameterize the vector of y values, \mathbf{y} , with an h -dimensional vector \mathbf{z} drawn from $\mathcal{N}(0, \mathbf{I}_h)$. y_{t+1} is distributed according to $\mathcal{N}(\mu^{(t)}(\mathbf{x}_{t+1}; \mathcal{D}_t), K^{(t)}(\mathbf{x}_{t+1}, \mathbf{x}_{t+1}; \mathcal{D}_t))$, so we map \mathbf{z}_{t+1} to y_{t+1} by a simple scale-and-shift. This map is done




Figure 3: We estimate a function via MC. (Left) Standard MC is noisy, and the estimate’s argmin is not the function’s. (Right) Using CRN makes a significant difference. We shade point-wise variances above and below. Note that they are unchanged; CRN only affects the estimate’s smoothness, which in this case yields the correct argmin.

sequentially from time step one to time step h . MC integration thus involves sampling N times from $\mathcal{N}(0, I_h)$, mapping each of the samples, and averaging. The mapping step is equivalent to applying our rollout policy $\tilde{\pi}$, and is the dominant cost of integration.

Compared to other quadrature schemes, MC is well-suited to high-dimensional integration. MC converges at a rate of σ/\sqrt{N} , the standard deviation of the MC estimator, where σ is the sample variance and N is the total number of samples. MC’s primary drawback is slow convergence. Increasing precision by an order of magnitude requires two orders of magnitude more samples. If σ is high, many samples may be required to converge. In this section, we focus on two strategies that significantly decrease the overhead of rollout: variance reduction and policy search.

3.1 Variance reduction

Variance reduction is a class of methods that improve convergence by decreasing the variance of the estimator. Effective variance reduction methods can reduce σ by orders of magnitude. We use a combination of quasi-Monte Carlo, common random numbers, and control variates, which significantly reduces the number of MC samples needed, as evidenced by Figure 2.

Quasi-Monte Carlo (QMC): Instead of sampling directly from the probability distribution, QMC instead uses a low-discrepancy sequence as its sample set.

Theorem 3.1 *Caflisch [2]: QMC converges at a rate bounded above by $\log(N)^h/N$, where N is the number of samples and h is the integral’s dimension.*

This bound stems from the well-known Koksma-Hlawka inequality [2], and is roughly linear for large N . In practice, this bound is often loose and convergence proceeds

faster [17]. QMC is inapplicable if a low-discrepancy sequence does not exist for the target distribution.

In our case, the distributions we integrate over are Normal, for which low-discrepancy sequences do exist. We generate low-discrepancy Sobol sequences in the h -dimensional uniform distribution $\mathcal{U}[0, 1]^h$ and map them to the unit multivariate Gaussian via the Box-Muller transform. This yields a low-discrepancy sequence for $\mathcal{N}(0, I_h)$. Recall the parameterization of samples from $\mathcal{N}(0, I_h)$ to sample rollout trajectories in Equation 1. We apply QMC by replacing the unit multivariate Gaussian samples with our low-discrepancy sequence.

Common random numbers (CRN): CRN is used when estimating a quantity to be optimized over parameter \mathbf{x} , and is implemented by using the same random number stream for all values of \mathbf{x} . CRN does not decrease the point-wise variance of an estimate, but rather decreases the covariance between two neighboring estimates, which smooths out the function. Consider estimating $\mathbb{E}_y[f(y, \mathbf{x}_1)]$ and $\mathbb{E}_y[f(y, \mathbf{x}_2)]$ for two points \mathbf{x}_1 and \mathbf{x}_2 using N samples, with variances σ_1/N and σ_2/N respectively. We define the differences δ and $\hat{\delta}$ as:

$$\delta = \mathbb{E}_y[f(y, \mathbf{x}_1)] - \mathbb{E}_y[f(y, \mathbf{x}_2)],$$

$$\hat{\delta} = \mathbb{E}_y[f(y, \mathbf{x}_1) - f(y, \mathbf{x}_2)],$$

$$\text{Var}[\delta] = (\sigma_1^2 + \sigma_2^2)/N,$$

$$\text{Var}[\hat{\delta}] = (\sigma_1^2 + \sigma_2^2 - 2\text{Cov}[f(y, \mathbf{x}_1), f(y, \mathbf{x}_2)])/N.$$

$\hat{\delta}$ uses the same number stream for \mathbf{x}_1 and \mathbf{x}_2 . If \mathbf{x}_1 and \mathbf{x}_2 are close and f is continuous in \mathbf{x} , $\text{Cov}[f(y, \mathbf{x}_1), f(y, \mathbf{x}_2)] > 0$, implying $\text{Var}[\hat{\delta}] < \text{Var}[\delta]$. This increased consistency between neighboring points improves optimization accuracy, as seen in Figure 3.

Control variates: The general idea behind control variates is to find a covariate $g(y)$ with the same distribution as $f(y)$ and a known mean. The quantity $c(y) = f(y) - \beta g(y)$, known as a regression control variate (RCV), is estimated, and then de-biased afterwards.

Theorem 3.2 *Consider the estimator $\mathbb{E}[c(y)] = \mathbb{E}[f(y) - \beta g(y)]$. Let $\text{Var}[f(y)], \text{Var}[g(y)] = \sigma_f, \sigma_g$. $g(y)$ is sufficiently correlated with $f(y)$ if:*

$$\beta^2 \sigma_g^2 - 2\beta \text{Cov}[f(y), g(y)] < 0.$$

If $g(y)$ is sufficiently correlated with $f(y)$, then the estimator $\mathbb{E}[c(y)]$ is strictly more accurate than $\mathbb{E}[f(y)]$ i.e., $\text{Var}[c(y)] < \text{Var}[f(y)]$.

The optimal value that minimizes the variance of $c(y)$ is $\beta = \text{Cov}[f(y), g(y)]/\sigma_g^2$. In practice, both $\text{Cov}[f(y), g(y)]$ and σ_g^2 must be estimated. This formula

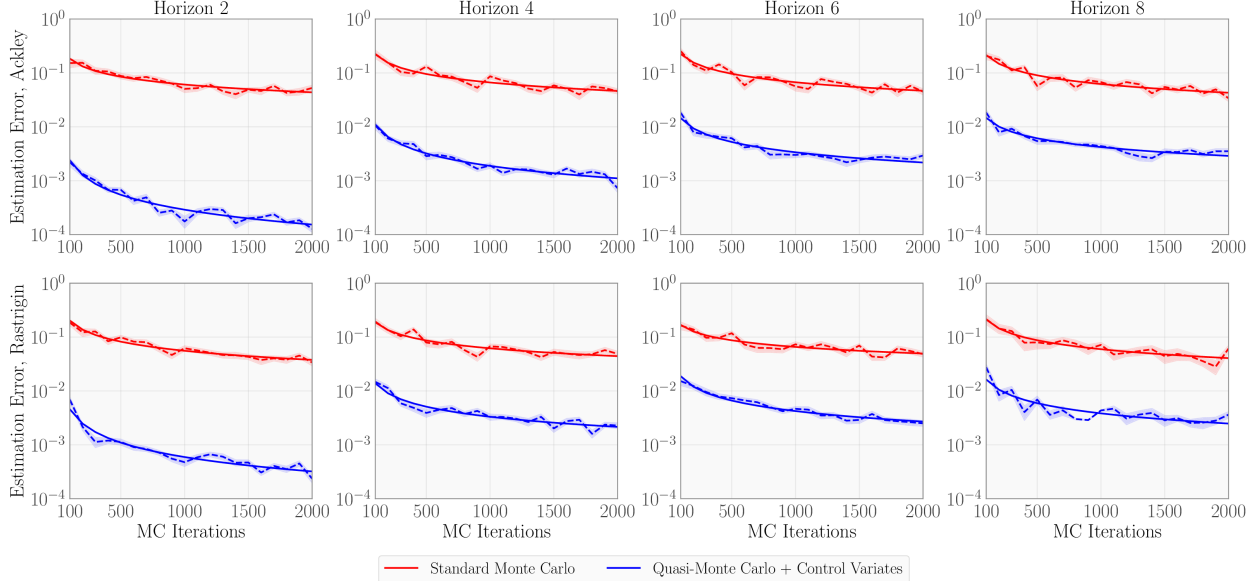


Figure 4: The estimation errors of MC (red) and our reduced-variance estimator (blue).

generalizes to multiple control variates (included in the supplementary). We note that variates whose derivatives are also correlated with the derivatives of $f(y)$ yield superior performance, especially when using QMC [8].

In the context of BO, we opt to use covariates derived from existing acquisition functions with known means. EI and PI are straightforward options. We expect their value to be at least somewhat correlated with the value of the rollout acquisition function; a promising candidate point should ideally score highly among all acquisition functions, and vice-versa. We will demonstrate the effectiveness of these variates in Section 4. EI and PI are defined as the expectation of a random variable of the form $\Lambda(\mathbf{x}) = \mathbb{E}[\gamma(y|\mathbf{x})] = \int_{-\infty}^{\infty} \gamma(y|\mathbf{x})p(y|\mathbf{x})dy$:

- *Probability of improvement (PI)*

$$\gamma(y|\mathbf{x}) = \begin{cases} 1, & y < y^* \\ 0, & y \geq y^*. \end{cases}$$

- *Expected improvement (EI)*

$$\gamma(y|\mathbf{x}) = \begin{cases} y^* - y, & y < y^* \\ 0, & y \geq y^*. \end{cases}$$

3.2 Policy search

While we have dramatically lowered the cost of evaluating the rollout acquisition function, there still remains the problem of its optimization. Wu and Frazier [24] use the reparameterization trick to estimate the gradient of EI for horizon two and use stochastic gradient descent to

maximize it. However, their method does not immediately extend to horizons larger than two.

Policy search is an alternative method for approximately solving MDPs, in which a best performing policy is chosen out of a (possibly infinite) set of policies $\Pi = \{\pi_1, \pi_2, \dots\}$ [1]. It is performed either by computing the expected reward for each policy in the set, or using a gradient-based method to maximize the expected reward given a parameterization of the policy set. In this paper, we use a finite policy set:

$$\Pi_{ps} = \{\pi(\mathcal{D}_k) \mid \pi(\mathcal{D}_k) := \arg \max \Lambda(\mathbf{x}|\mathcal{D}_k) \in A\},$$

where A is any arbitrary set of acquisition functions. We then select the best-performing policy:

$$\pi_{ps} = \arg \max_{\pi_i \in \Pi_{ps}} V_h^{\pi_i}(\mathcal{D}_k).$$

A less formal explanation follows: at every step of BO, we roll out each acquisition function in A on its argmax, and use the one with the highest h -step reward. A related approach by Hoffman et al. [9] employs a bandit strategy to switch between different acquisition functions. Our policy search method does not maximize the rollout acquisition function, and is thus significantly faster, though it likely reduces performance. Experiments in Section 4 suggest that our policy search method performs at least as well as the best-performing acquisition in Π_{ps} .

4 EXPERIMENTS AND DISCUSSION

Unless otherwise stated, we use a GP with the Matérn 5/2 ARD kernel [22] and learn its hyperparameters via

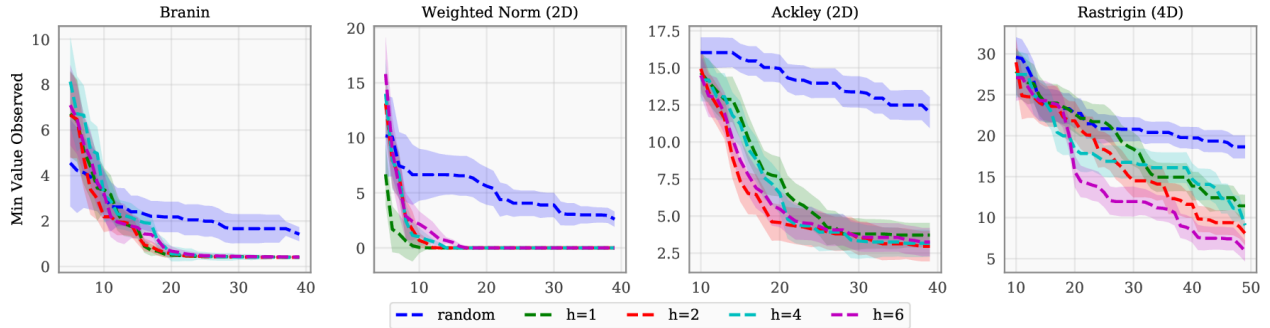


Figure 5: Empirically, looking longer horizons only seems to help on multimodal functions. On unimodal functions (not necessarily convex), there is little to no performance gain.

maximum likelihood estimation [20]. When rolling out acquisition functions, we maximize them with L-BFGS-B using five restarts, selected by evaluating the acquisition on a Latin hypercube of $10d$ points and picking the five best. EI is used as the base rollout policy except for in the policy search experiments. All synthetic functions are found in the supplementary. Code to reproduce our experiments is found at https://github.com/ericlee0803/lookahead_release.

Variance reduction experiments: We compare the estimation error and convergence rate between the standard MC estimator and our estimator. We take $2d$ random points in the domain and evaluate the Ackley and Rastrigin functions in 2D and 4D, respectively. We roll out EI for horizons 2, 4, 6, and 8, and calculate the variance of both estimators for MC sample sizes in $[100, 200, 300, \dots, 2000]$, using 50 trials each. We take the ground truth to be estimation with 10^4 samples. The mean error of the standard MC estimator (*red*) and our reduced-variance estimator (*blue*) are plotted with dotted lines in Figure 4, with standard error shaded above and below. We also compute a best-fit line to each mean error, which is plotted with a solid line.

Table 1: Estimated convergence rate and error reduction $\sigma/\hat{\sigma}$ for the standard MC estimator and our estimator

Objective	Horizon	MC Rate	Our Rate	$\sigma/\hat{\sigma}$
Ackley	2	0.53	0.95	410
Ackley	4	0.52	0.82	63
Ackley	6	0.55	0.64	28
Ackley	8	0.53	0.54	26
Rastrigin	2	0.56	0.90	150
Rastrigin	4	0.48	0.63	31
Rastrigin	6	0.47	0.68	30
Rastrigin	8	0.42	0.64	25

Table 1 summarizes our experimental results, and includes our estimates for the convergence rate of both estimators and the relative reduction in estimation error $\sigma/\hat{\sigma}$. Our estimator has significantly lower estimation error—the

maximum reduction in estimation error we achieve is a factor of 410. Standard MC clearly converges at a $N^{-1/2}$ rate. Our estimator converges like N^{-1} for smaller horizons, but its convergence rate drops as h increases. This is due to QMC’s $\log(N)^h/N$ convergence. N is not large enough for longer horizons to exhibit N^{-1} convergence; increasing it past 2000 should yield N^{-1} convergence.

Another trend is the increase in estimation error as the horizon increases, which is expected given that the dimensionality of the underlying integral increases. Fortunately, the error seems to increase only linearly—and by a small constant—rather than exponentially, suggesting that MC samples proportional to h is sufficient to achieve a high quality of approximation. Finally, the reduction in estimation error levels off to around a factor of 25, suggesting that the correlation between the rollout acquisition function and our control variates decreases when h increases. We include an ablation study in the supplement to quantify the individual contributions of QMC and control variates.

A factor of 25 error reduction is still significant; the standard MC estimator would need 625 times more samples to achieve comparable accuracy.

Full rollout on synthetic functions: We roll out EI for $h = 2, 4,$ and 6 on the Branin, weighted-two-norm (2D), Ackley (2D), and Rastrigin (4D) functions in Figure 5 using $200h$ MC samples, and compare to both standard EI and random search. To optimize the acquisition functions quickly, we employ the following strategy: we evaluate the acquisition function on a Sobel sequence of size $10d$, as well as an additional point which is the argmax of EI. We then use this as an initial design and run BO for $50d$ more iterations. We run 50 iterations for each horizon and provide random search as a baseline. The mean results and the standard error are plotted in Figure 5.

On the Branin function, all horizons performed comparably and converge in 20 iterations. Rollout performed best on the Ackley and Rastrigin functions, which are multi-

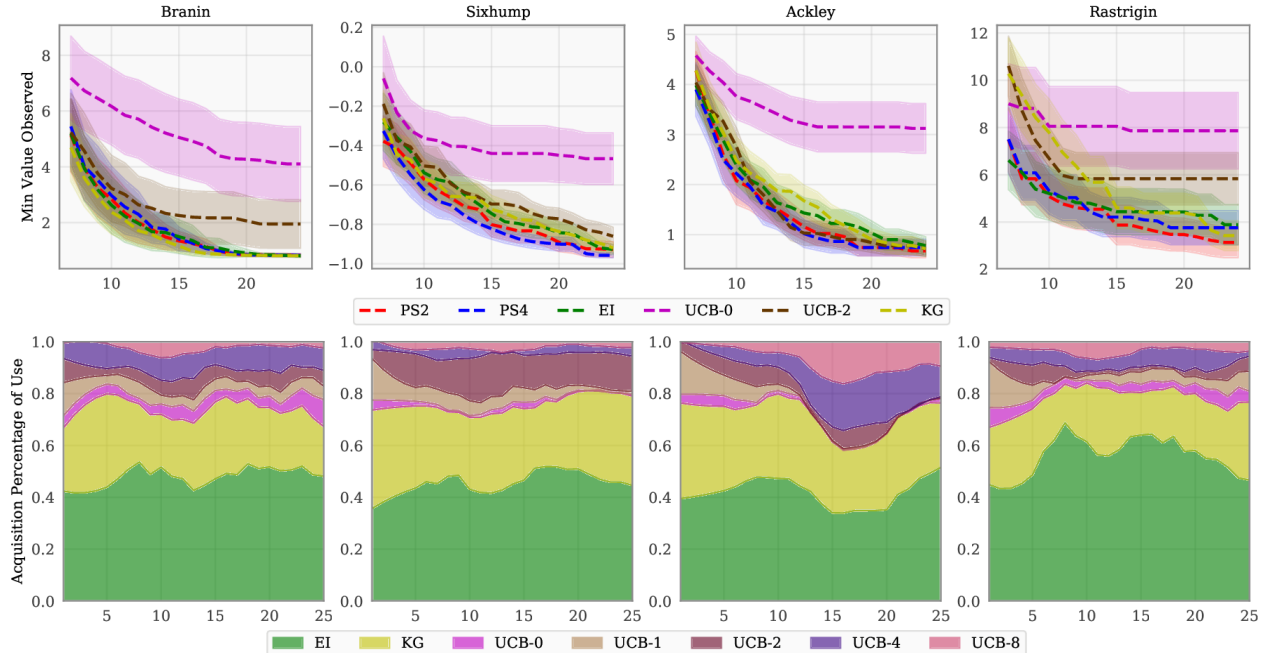


Figure 6: (*Top*) Policy search performs at least as well as the best acquisition function, if not better. (*Bottom*) For each of the corresponding objectives, we plot the percentage of use of each acquisition function per iteration for PS4. EI and KG are chosen more often than any of the UCB acquisition functions. The worst-performing acquisition, UCB-0, is chosen the least, suggesting correlation between an acquisition’s performance and its percentage of use.

modal. On the weighted norm function, which is strongly convex, EI converges within 10 iterations, and looking ahead further yielded poorer results. These results suggest that more exploratory acquisitions are needed for a multi-modal objective, whereas more exploitative acquisitions suffice for reasonably simple objective functions.

Policy search: We consider policy search (PS) with an acquisition set of EI, KG, and Upper Confidence Bound (UCB- κ) for $\kappa \in \{0, 1, 2, 4, 8\}$ [22], which contains acquisitions that tend towards both exploitation and exploration. We run policy search for horizons 2 and 4 on the Branin, Sixhump, Ackley, and Rastrigin synthetic functions, all in 2D, using $200h$ MC samples.

All acquisition functions are maximized via L-BFGS-B with five random restarts, except for KG, which uses grid search of size 900. The mean results and standard error over 50 trials are plotted in Figure 6, in which policy search for horizons 2 and 4, labeled PS2 and PS4 respectively, do better or on par with the best-performing acquisition function. This robustness is a key strength of policy search, as the performance of each acquisition function is often problem-dependent. We also examine the choice of acquisition function as a function of iteration. The percentage of use of each acquisition function is shown for its corresponding objective, and for plotting purposes we smooth the percentage with a box filter of

size five. EI and KG are chosen more often than the other acquisition functions. UCB-0, the worst performing method representing a pure exploitation policy, is chosen significantly less than others, while UCB-2 was chosen the most often out of the UCB family. Of particular interest is the Ackley function (third column, Figure 6): when UCB-2 starts to outperform the other acquisitions functions, a clear spike in its percentage of use is seen in the corresponding histogram.

Neural architecture search (NAS) benchmark: We run policy search on the NAS tabular benchmarks in [13], which are an exhaustive set of evaluated configurations for multi-layer perceptrons trained on different datasets. We optimize the perceptrons’ layer sizes, batch size, and training epochs with 60 iterations of BO for each dataset and plot average regret in Figure 8. For space’s sake, we describe the search space and the regret metric in the supplement. Policy search for $h = 2$ and $h = 4$ performs better than EI, UCB0, UCB2, and KG.

The impact of model mis-specification: We believe that any probabilistic model only supports a limited horizon due to the effects of model error. Errors in the GP model result in errors to the MDP transition probabilities, which grow as they are propagated through time. This likely renders long-horizon rollout ineffectual. We support this hypothesis by comparing the performance of policies in




Figure 7: (Left) The expected performance of EI-based rollout for $h = 1, 2, 3, 4,$ and 5 . (Middle, Right) The observed performance of rollout, given model error in the form of a smoother and less smooth kernel, respectively. When the model has large error, the resulting performance of non-myopic policies can be reversed from the expected performance.




Figure 8: We compare the performance of policy search for horizons 2 and 4 in red and blue, respectively, with that of EI, UCB0, UCB2, and KG. PS2 and PS4 outperform the others after about 20 iterations.

the MDP setting they were designed in with their observed performance on objectives drawn from a different MDP.

We do this by drawing objective functions from a GP with fixed kernel every step of BO. More concretely, evaluating the objective function at any point \mathbf{x} is performed by sampling from the GP posterior distribution at \mathbf{x} . Because policies are designed to maximize this MDP’s reward for a fixed horizon and because the objective is drawn from the MDP itself, policies looking further ahead perform better by definition. We then draw objectives from a GP using a *different* kernel with a *different* lengthscale, and check if policies looking further ahead still perform better.

We rollout EI in 1D with a budget of seven and we model our objective with a GP using the Matérn 5/2 kernel with $\ell = 0.2$. In Figure 7, the left panel depicts expected performance of rollout for $h = 1, 2, 3, 4,$ and 5 . The middle and right panels depict observed performance of

rollout when we sample objectives from a GP that has a far smoother kernel (Square Exponential with $\ell = 0.8$) and far less smooth kernel (Matérn 3/2 with $\ell = 0.05$), respectively. All plots use 2000 replications to achieve high accuracy.

The result is perhaps unsurprising; the ranking of the observed performance of policies is reversed with that of the expected performance. Myopic BO performed the best; more generally, policies with shorter horizons outperformed those with longer horizons. This demonstrated sensitivity to model error suggests non-myopic BO must carefully strike a balance between model accuracy and horizon, and justifies use of modest horizons over the full BO budget. This confirms experimental results by Yue and Kontar [25], who suggest looking ahead to short horizons is preferable to long horizons in practice.

5 CONCLUSION

We have shown that a combination of quasi-Monte Carlo, control variates, and common random numbers significantly lowers the overhead of rollout in BO. We have introduced a policy search which further decreases computational cost by removing the need to maximize the rollout acquisition function. Finally, we have illustrated the penalties incurred by using inaccurate GP models in the non-myopic setting.

This work raises several interesting research directions. Decreasing the variance of our estimator may be possible with additional variance reduction methods such as stratified or antithetic sampling. Developing a more comprehensive policy search space, such as a parameterized set of all convex combinations of acquisition functions, may further strengthen the policy search performance.

References

- [1] D. P. Bertsekas. *Dynamic Programming and Optimal Control*, volume I. Athena scientific Belmont, MA, 4th edition, 2017.
- [2] R. E. Caflisch. Monte Carlo and quasi-Monte Carlo methods. *Acta numerica*, 7:1–49, 1998.
- [3] P. I. Frazier. Bayesian optimization. In E. Gel and L. Ntaimo, editors, *Recent Advances in Optimization and Modeling of Contemporary Problems*, pages 255–278. INFORMS, 2018. doi: 10.1287/educ.2018.0188.
- [4] P. I. Frazier, W. B. Powell, and S. Dayanik. A knowledge-gradient policy for sequential information collection. *SIAM Journal on Control and Optimization*, 47(5):2410–2439, 2008.
- [5] D. Ginsbourger and R. Le Riche. Towards Gaussian process-based optimization with finite time horizon. In *Proceedings of the 9th International Workshop in Advances in Model-Oriented Design and Analysis*, pages 89–96. Springer, 2010.
- [6] J. González, M. A. Osborne, and N. Lawrence. GLASSES: Relieving the myopia of Bayesian optimisation. In *Artificial Intelligence and Statistics*, pages 790–799, 2016.
- [7] J. M. Hernández-Lobato, M. W. Hoffman, and Z. Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In *Proceedings of the 28th Conference on Neural Information Processing Systems*, pages 918–926, 2014.
- [8] F. J. Hickernell, C. Lemieux, A. B. Owen, et al. Control variates for quasi-monte carlo. *Statistical Science*, 20(1):1–31, 2005.
- [9] M. D. Hoffman, E. Brochu, and N. de Freitas. Portfolio allocation for Bayesian optimization. In *UAI*, pages 327–336. Citeseer, 2011.
- [10] S. Jiang, G. Malkomes, G. Converse, A. Shofner, B. Moseley, and R. Garnett. Efficient nonmyopic active search. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1714–1723, 2017.
- [11] S. Jiang, G. Malkomes, M. Abbott, B. Moseley, and R. Garnett. Efficient nonmyopic batch active search. In *Proceedings of the 32nd Conference on Neural Information Processing Systems*, pages 1107–1117, 2018.
- [12] D. R. Jones, M. Schonlau, and W. J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13(4): 455–492, 1998. ISSN 1573-2916.
- [13] A. Klein and F. Hutter. Tabular benchmarks for joint architecture and hyperparameter optimization. *arXiv preprint arXiv:1905.04970*, 2019.
- [14] R. Lam and K. Willcox. Lookahead Bayesian optimization with inequality constraints. In *Proceedings of the 31st Conference on Neural Information Processing Systems*, pages 1890–1900, 2017.
- [15] R. Lam, K. Willcox, and D. H. Wolpert. Bayesian optimization with a finite budget: An approximate dynamic programming approach. In *Proceedings of the 30th Conference on Neural Information Processing Systems*, pages 883–891, 2016.
- [16] M. A. Osborne, R. Garnett, and S. J. Roberts. Gaussian processes for global optimization. In *Proceedings of the 3rd International Conference on Learning and Intelligent Optimization (LION3)*, pages 1–15, 2009.
- [17] A. Papageorgiou. Sufficient conditions for fast quasi-Monte Carlo convergence. *Journal of Complexity*, 19(3):332–351, 2003.
- [18] W. B. Powell. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley, 2nd edition, 2007.
- [19] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., 2014.
- [20] C. E. Rasmussen and C. K. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [21] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, Jan 2016.
- [22] J. Snoek, H. Larochelle, and R. P. Adams. Practical Bayesian optimization of machine learning algorithms. In *Proceedings of the 26th Conference on Neural Information and Processing Systems*, pages 2951–2959, 2012.
- [23] N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *Proceedings of the 27th International Conference on Machine Learning*, pages 1015–1022, 2010.
- [24] J. Wu and P. Frazier. Practical two-step lookahead Bayesian optimization. In *Advances in Neural Information Processing Systems*, pages 9810–9820, 2019.
- [25] X. Yue and R. A. Kontar. Why non-myopic Bayesian optimization is promising and how far should we look-ahead? A study via rollout. *Artificial Intelligence and Statistics*, 2020.