
Output-Weighted Sampling for Multi-Armed Bandits with Extreme Payoffs

Yibo Yang¹

Antoine Blanchard²

Themistoklis Sapsis²

Paris Perdikaris¹

¹Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, Philadelphia, PA 19104

²Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139

Abstract

We present a new type of acquisition functions for online decision making in multi-armed and contextual bandit problems with extreme payoffs. Specifically, we model the payoff function as a Gaussian process and formulate a novel type of upper confidence bound (UCB) acquisition function that guides exploration towards the bandits that are deemed most relevant according to the variability of the observed rewards. This is achieved by computing a tractable likelihood ratio that quantifies the importance of the output relative to the inputs and essentially acts as an *attention mechanism* that promotes exploration of extreme rewards. We demonstrate the benefits of the proposed methodology across several synthetic benchmarks, as well as a realistic example involving noisy sensor network data. Finally, we provide a JAX library for efficient bandit optimization using Gaussian processes.

1 INTRODUCTION

Online decision making defines an important branch of modern machine learning in which uncertainty quantification plays a prominent role. In most stochastic optimization settings, evaluating the unknown function is expensive, hence new information needs to be acquired judiciously. Classical applications include recommendation systems for articles and products, where the goal is to maximize the total revenue of the product maker given limited user feedback [Li et al., 2010, Kawale et al., 2015]; control and reinforcement learning, where the reward is obtained after a sequence of experiments or actions and the objective is not only to obtain optimal rewards but also avoid the potentially negative effects of uncertainty [Dearden et al., 1998, Osband et al., 2016, Azizzadenesheli et al., 2018, Li et al., 2019]; environment monitoring, where sensor data is used to identify areas

of interest as in traffic flow estimation [Srinivas et al., 2009] and room temperature monitoring [Krause and Ong, 2011]; and optimal design of expensive experiments [Sacks et al., 1989, Saha et al., 2008].

More recently, new applications have appeared beyond machine learning, including optimal sampling in cardiac electrophysiology and bio-engineering [Sahli Costabal et al., 2019, 2020], multi-fidelity design of experiments [Forrester et al., 2007, Sarkar et al., 2019], hyper-parameter tuning in high-dimensional design spaces [Shan and Wang, 2010, Perdikaris et al., 2016, Bouhlel et al., 2016], and prediction of extreme events in complex dynamical systems [Wan et al., 2018, Mohamad and Sapsis, 2018].

Many of these applications can be formulated as multi-armed bandit problems, for which effective sampling algorithms exist [Auer, 2002, Srinivas et al., 2009, Chu et al., 2011, Krause and Ong, 2011, Schaul et al., 2015, Osband et al., 2016]. These algorithms are generally characterized by two key ingredients. First, they involve building a model for the latent payoff function given scarce and possibly noisy observations of past rewards. To enable effective sampling and exploration of the decision space, uncertainty in the model predictions needs to be accounted for in the predictive posterior distribution of the latent payoffs, which can be obtained via either a frequentist or a Bayesian approach. The second critical ingredient pertains to designing a data acquisition policy that can leverage the model predictive uncertainty to effectively balance the trade-off between exploration and exploitation while ensuring a consistent asymptotic behavior for the cumulative regret.

1.1 PREVIOUS WORK

Multi-armed bandit problems provide a general setting for developing online decision-making algorithms and rigorously studying their performance. Early research in this setting includes the celebrated ϵ -greedy algorithm [Schaul et al., 2015], where random exploration is introduced with a

small probability ϵ to prevent the algorithm from focusing on local sub-optimal solutions. Despite its widespread applicability, ϵ -greedy algorithms employ a heuristic treatment of uncertainty, and often require careful tuning in order to prevent sub-optimal exploration.

To this end, the upper confidence bound (UCB) policy [Agrawal, 1995, Auer, 2002] was proposed to provide a natural way to estimate sub-optimal choices using a model’s predictive posterior uncertainty. However, the original UCB formulation does not take into account correlations between different bandits in a multi-armed setting and, therefore, typically requires a large number of datapoints to be collected before convergence can be observed. Variants of the UCB algorithm have been adapted to the contextual bandit setting with linear payoffs, where the payoff function is modeled via Bayesian linear regression [Chu et al., 2011]. Gaussian process models have also been employed to account for correlated payoffs, and the corresponding GP-UCB criteria have shown great promise in data-scarce and “cold start” scenarios [Dani et al., 2008, Srinivas et al., 2009, Krause and Ong, 2011].

Thompson sampling [Thompson, 1933, Russo et al., 2017] provides an alternative approach to balancing the exploration–exploitation trade-off that only requires access to posterior samples of a parametrized payoff function. Although the algorithm was largely ignored at the time of its inception by Thompson [1933], the results of Chapelle and Li [2011] have initiated a wave of resurgence, leading to significant advances in applications (e.g., recommendation systems [Kawale et al., 2015], hyper-parameter optimization [Kandasamy et al., 2018], reinforcement learning [Dearden et al., 1998, Osband et al., 2016, Azizzadenesheli et al., 2018]), as well as theoretical analyses (e.g., optimal regret bounds [Kaufmann et al., 2012, Leike et al., 2016, Russo and Van Roy, 2016]). More recently, Bayesian deep learning models have been considered [Graves, 2011] for modeling more complex and high-dimensional payoff functions. However, their effectiveness, interpretability, and convergence behavior are still under investigation [Riquelme et al., 2018].

1.2 OUR CONTRIBUTIONS

Primary contribution. All aforementioned approaches have enjoyed success across various applications, however they lack a mechanism for distinguishing and promoting the input/context variables that have the greatest influence on the observed payoffs. Short of such mechanism, regions in the decision space that may have negligible effect on the payoffs will still be sampled as long as they are uncertain. As we will demonstrate, this undesirable behavior can have a deteriorating impact on convergence, and this effect is exacerbated in the presence of extreme payoffs (i.e., situations in which a small number of bandits yield rewards significantly greater than the rest of the bandit population).

Motivated by the recent findings of Sapsis [2020] and Blanchard and Sapsis [2020b, 2021], we introduce a novel UCB-type objective for online decision making in multi-armed and contextual bandit problems that can overcome the aforementioned pathologies. This is achieved by introducing an importance weight to effectively promote the exploration of “heavy-tailed” (i.e., rare and extreme) payoffs. We show how such importance weight can be derived from a likelihood ratio that quantifies the relative importance between inputs/contexts and observed rewards, introducing an effective *attention mechanism* that favors exploration of bandits with unusually large rewards over bandits associated with frequent, average payoffs. This output-weighted approach has been shown to outperform classical acquisition functions in active learning [Blanchard and Sapsis, 2020a] and Bayesian optimization [Blanchard and Sapsis, 2021] tasks, and here we set sail for the first time into investigating its effectiveness in online decision making tasks, with a specific focus on multi-armed and contextual bandit problems subject to extreme payoffs.

Comparison to previous work. We demonstrate the effectiveness of the proposed methodology across a collection of synthetic benchmarks, as well as a realistic example involving noisy sensor network data. In all cases, we provide comprehensive quantitative comparisons between the proposed output-weighted sampling criterion and the most widely-used criteria in current practice, including the UCB [Auer, 2002], GP-UCB [Srinivas et al., 2009], Thompson sampling [Thompson, 1933, Chapelle and Li, 2011], and expected improvement [Vazquez and Bect, 2010] methods.

Secondary contributions. We have developed an open-source Python package for bandit optimization using Gaussian processes¹. Our implementation leverages the high-performance package JAX [Bradbury et al., 2018] and thus enables (a) gradient-based optimization of the proposed output-weighted sampling criteria for general Gaussian process priors, (b) the use of GPU acceleration, and (c) scalability and parallelization across multiple computing nodes. This package can be readily used to reproduce all data and results presented in this paper.

2 METHODS

2.1 MULTI-ARMED BANDITS

The multi-armed bandit problem is a prototypical paradigm for sequential decision making. The decision set consists of a discrete collection of M arms where the i th arm may be associated with some contextual information $\mathbf{x}_i \in \mathbb{R}^d$. Pulling arm i produces a reward $y \in \mathbb{R}$ which is determined

¹<https://github.com/PredictiveIntelligenceLab/jax-bandits>.

by some unknown latent function

$$y_t = f(\mathbf{x}_t) + \varepsilon_t, \quad (1)$$

where $\varepsilon_t \sim \mathcal{N}(0, \sigma_n^2)$ accounts for observation noise.

At each round t , we select an arm i and obtain a reward y_t . The goal of sequential decision making is to find a strategy for bandit selection that maximizes the total reward $\sum_{t=1}^T y_t$ for a given budget T . In other words, the goal is to first identify the bandits that provide the best rewards,

$$\mathbf{x}^* = \arg \max_{\mathbf{x}} f(\mathbf{x}), \quad (2)$$

using as few arm pulls as possible, and then to keep on exploiting these optimal bandits to maximize the total reward.

As an alternative metric of success, it is useful to consider the simple regret $r_t = f(\mathbf{x}^*) - f(\mathbf{x})$, as maximizing the total reward is essentially equivalent to minimizing the *cumulative* regret

$$R_T = \sum_{t=1}^T r_t. \quad (3)$$

The holy grail of online decision making is to design an effective no-regret policy satisfying

$$\lim_{T \rightarrow \infty} \frac{R_T}{T} = 0. \quad (4)$$

2.2 GAUSSIAN PROCESSES

Gaussian process (GP) regression provides a flexible probabilistic framework for modeling nonlinear black-box functions [Rasmussen and Williams, 2006]. Given a dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ of input–output pairs (i.e, context–reward pairs), and an observation model of the form $y = f(\mathbf{x}) + \varepsilon$, the goal is to infer the latent function f as well as the unknown noise variance σ_n^2 corrupting the observations.

In GP regression, no assumption is made on the form of the latent function f to be learned; rather, a prior probability measure is assigned to every function in the function space. Starting from a zero-mean Gaussian prior assumption on f ,

$$f(\mathbf{x}) \sim \text{GP}(\mathbf{0}, k(\mathbf{x}, \mathbf{x}'; \boldsymbol{\theta})), \quad (5)$$

the goal is to identify an optimal set of hyper-parameters $\Theta = \{\boldsymbol{\theta}, \sigma_n^2\}$, and then use the optimized model to predict the rewards of unseen bandits. The covariance function $k(\mathbf{x}, \mathbf{x}'; \boldsymbol{\theta})$ plays a key role in this procedure as it encodes prior belief or domain expertise one may have about the underlying function f . In the absence of any domain-specific knowledge, it is common to assume that f is a smooth continuous function and employ the squared exponential covariance kernel with automatic relevance determination (ARD) which accounts for anisotropy with respect to each input variable [Rasmussen and Williams, 2006].

Unlike previous works [Srinivas et al., 2009, Krause and Ong, 2011], here we do not assume that the payoff function f actually comes from a GP prior or that it has low RKHS norm. Instead, we compute an optimal set of hyper-parameters at each round t by minimizing the negative log-marginal likelihood of the GP model [Rasmussen and Williams, 2006]. In our setup, the likelihood is Gaussian and can be computed analytically as

$$\begin{aligned} \mathcal{L}(\Theta) = & \frac{1}{2} \log |\mathbf{K} + \sigma_n^2 \mathbf{I}| \\ & + \frac{1}{2} \mathbf{y}^\top (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y} + \frac{N}{2} \log(2\pi), \end{aligned} \quad (6)$$

where \mathbf{K} is an $N \times N$ covariance matrix constructed by evaluating the kernel function on the input training data \mathbf{X} . The minimization problem is solved with an L-BFGS optimizer with random restarts [Liu and Nocedal, 1989].

Once the GP model has been trained, the predictive distribution at any given bandit \mathbf{x} can be computed by conditioning on the observed data:

$$p(\mathbf{y} | \mathbf{x}, \mathcal{D}) \sim \mathcal{N}(\boldsymbol{\mu}(\mathbf{x}), \boldsymbol{\sigma}^2(\mathbf{x})), \quad (7)$$

where

$$\boldsymbol{\mu}(\mathbf{x}) = k(\mathbf{x}, \mathbf{X})(\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y}, \quad (8a)$$

$$\boldsymbol{\sigma}^2(\mathbf{x}) = k(\mathbf{x}, \mathbf{x}) - k(\mathbf{x}, \mathbf{X})(\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} k(\mathbf{X}, \mathbf{x}). \quad (8b)$$

Here, $\boldsymbol{\mu}(\mathbf{x})$ can be used to make predictions and $\boldsymbol{\sigma}^2(\mathbf{x})$ to quantify the associated uncertainty.

2.3 ONLINE DECISION MAKING

A critical ingredient in online decision making is the choice of the *acquisition function*, which effectively determines which bandits the algorithm should try out and which ones to ignore [Srinivas et al., 2009, Krause and Ong, 2011]. A popular choice of acquisition function is the “vanilla” upper confidence bound (V-UCB),

$$a_{\text{V-UCB}}(\mathbf{x}) = \boldsymbol{\mu}(\mathbf{x}) + \kappa \boldsymbol{\sigma}(\mathbf{x}), \quad (9)$$

and the closely-related GP-UCB criterion [Srinivas et al., 2009],

$$a_{\text{GP-UCB}}(\mathbf{x}) = \boldsymbol{\mu}(\mathbf{x}) + \beta_t^{1/2} \boldsymbol{\sigma}(\mathbf{x}), \quad (10)$$

where κ and $\beta_t = 2 \log(|D|t^2\pi^2/(6\delta))$ are parameters that aim to balance exploration and exploitation. ($|D|$ is the number of bandits in the absence of context, and the dimension of the context otherwise.) In V-UCB, κ is typically considered constant, while in GP-UCB, β_t depends on the round t and comes with convergence guarantees when the payoff function is not too complex [Srinivas et al., 2009].

In this work we also consider the expected improvement,

$$a_{\text{EI}}(\mathbf{x}) = \boldsymbol{\sigma}(\mathbf{x})[\lambda(\mathbf{x})\Phi(\lambda(\mathbf{x})) + \phi(\lambda(\mathbf{x}))], \quad (11)$$

whose convergence properties have been well studied [Vazquez and Bect, 2010], as well as Thomson sampling,

$$a_{\text{TS}}(\mathbf{x}) = \tilde{y}(\mathbf{x}), \quad (12)$$

also known to deliver competitive results in practice [Chapelle and Li, 2011, Agrawal and Goyal, 2012, Riquelme et al., 2018]. In (11), we have defined $\lambda(\mathbf{x}) = (\mu(\mathbf{x}) - y^* - \xi)/\sigma(\mathbf{x})$, with y^* the best reward recorded so far and ξ a user-specified parameter controlling the exploration–exploitation trade-off. The quantity $\tilde{y}(\mathbf{x})$ in (12) denotes a random sample drawn from the posterior distribution of the GP model, that is, $\tilde{y}(\mathbf{x}) \sim \mathcal{N}(\mu(\mathbf{x}), \sigma^2(\mathbf{x}))$.

The goal in bandit optimization is to determine the best bandit to try next by maximizing the acquisition function:

$$\mathbf{x}_{t+1} = \arg \max_{\mathbf{x}} a(\mathbf{x}; \mathcal{D}), \quad (13)$$

where a can be any of (9), (10), (11), or (12), and \mathcal{D} contains all the observed context–reward pairs up to round t .

2.4 OUTPUT-WEIGHTED SAMPLING

Blanchard and Sapsis [2021] recently introduced an efficient and minimally intrusive approach for accelerating the stochastic optimization process in cases where certain regions of the input space have a considerably larger impact on the output of the latent function than others (i.e., extreme payoffs in the bandit problem) by incorporating a sampling weight into several of the acquisition functions commonly used in practice. The sampling weight, referred to as the “likelihood ratio”, was derived from a heavy-tail argument whereby the best next input point to visit is selected so as to most reduce the uncertainty in the *tails* of the output statistics where the extreme payoffs “live” (Figure 1).

The likelihood ratio is defined as

$$w(\mathbf{x}) = \frac{p_{\mathbf{x}}(\mathbf{x})}{p_{\mu}(\mu(\mathbf{x}))} \quad (14)$$

and was derived in Blanchard and Sapsis [2021]. Here, $p_{\mathbf{x}}(\mathbf{x})$ is a prior distribution that can be used to distill prior beliefs about the importance of each bandit or environmental conditions. In this work we assume that no such prior information is available and treat every bandit equally by specifying a uniform prior, $p_{\mathbf{x}}(\mathbf{x}) = 1$ for all \mathbf{x} . The term $p_{\mu}(\mu(\mathbf{x}))$ denotes the output density of the payoff function and plays an important role to determine the best arms to pull.

The intuition behind the likelihood ratio is as follows. Assuming enough data has been collected, the GP posterior mean $\mu(\mathbf{x})$ provides a good estimation about the distribution of rewards for the bandits. Bandits with unusually large rewards are associated with small values of p_{μ} , while bandits with frequent, average rewards are associated with large values of p_{μ} . Because the output density p_{μ} appears in

the denominator of (14), the likelihood ratio assigns more weight to bandits with extreme payoffs. As such, the likelihood serves as an *attention mechanism* which encourages the algorithm to explore bandits whose rewards are thought to be abnormally large, while penalizing the other mediocre bandits by assigning them small weights.

To obtain a well-behaved (i.e., smooth and bounded) analytical approximation of the likelihood ratio, we use a Gaussian mixture model,

$$w(\mathbf{x}) \approx \sum_{k=1}^{n_{\text{GMM}}} \alpha_k \mathcal{N}(\mathbf{x}; \gamma_k, \Sigma_k), \quad (15)$$

where $\mathcal{N}(\mathbf{x}; \gamma_k, \Sigma_k)$ denotes the k th component of the mixture model with mean γ_k and covariance Σ_k . The resulting output-weighted acquisition function for the bandit optimization problem is given by

$$a_{\text{LW-UCB}}(\mathbf{x}) = \mu(\mathbf{x}) + \kappa w(\mathbf{x}) \sigma(\mathbf{x}), \quad (16)$$

where the subscript “LW-UCB” stands for “likelihood-weighted UCB”. Equation (16) is subject to the same bandit-selection policy as the acquisition functions in Section 2.3:

$$\mathbf{x}_{t+1} = \arg \max_{\mathbf{x}} a_{\text{LW-UCB}}(\mathbf{x}; \mathcal{D}). \quad (17)$$

In general, the minimization problem can be efficiently solved with an L-BFGS optimizer with random restarts [Liu and Nocedal, 1989], where the gradient of the acquisition function with respect to the inputs \mathbf{x} can be computed analytically for the squared exponential covariance kernel [Blanchard and Sapsis, 2021], or using automatic differentiation [Baydin et al., 2015] for more general kernel choices. The workflow for output-weighted sampling with LW-UCB is summarized in Algorithm 1.

Algorithm 1: The LW-UCB algorithm.

- 1 **Input:** Small initial dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$;
 - 2 **while** $t < T$ **do**
 - 3 Fit GP model to dataset \mathcal{D} using (6) and obtain posterior mean (8a) and variance (8b);
 - 4 Compute likelihood ratio (14) and fit Gaussian mixture model (15) to it ;
 - 5 Select best next bandit \mathbf{x}_{t+1} by maximizing (16);
 - 6 Collect new reward $y_{t+1} = f(\mathbf{x}_{t+1}) + \varepsilon_{t+1}$ and append $(\mathbf{x}_{t+1}, y_{t+1})$ to dataset \mathcal{D} ;
 - 7 **end**
-

3 RESULTS

In all numerical studies considered in this work, we initialize the algorithm with $n = 3$ random input–output pairs and compare the performance of EI, TS, V-UCB, GP-UCB, and




Figure 1: Sketch of the acquisition scheme from which the likelihood ratio is derived. The best next bandit \mathbf{x}^* maximizes the reduction of the uncertainty in the tails of the payoff distribution (quantified by the log-difference between p_{μ^+} and p_{μ^-}).

LW-UCB. Our metric of success is the log-cumulative regret over time. Unless otherwise indicated, we conduct a series of 100 random experiments, each with a different choice of initial data, and report the median of the metric of interest. Variability across experiments is quantified using the median absolute deviation.

3.1 SYNTHETIC BENCHMARKS

We demonstrate the performance of LW-UCB for three synthetic test functions. We consider 2500 bandits arranged on a uniform 50×50 grid with rewards being given by the value of the test function at that point in the domain. The rewards collected during optimization are corrupted by small Gaussian noise with $\sigma_n = 10^{-4}$.

We begin with the Cosine function of Azimi et al. [2010],

$$f(\mathbf{x}) = 1 - [u^2 + v^2 - 0.3 \cos(3\pi u) - 0.3 \cos(3\pi v)] \quad (18)$$

where $u = 1.6x_1 - 0.5$, $v = 1.6x_2 - 0.5$, and $\mathbf{x} \in [0, 1]^2$. For $n_{GMM} = 2$, Figure 2a shows that LW-UCB performs better than the other methods as it leads to faster identification of the best bandit. Moreover, Figure 2a demonstrates how the likelihood ratio highlights the importance of the bandits and favors exploration of those with the highest rewards. We also note the subpar performance of EI, consistent with the discussion in Qin et al. [2017].

Next, we consider the Michalewicz function [Azimi et al., 2010],

$$f(\mathbf{x}) = \sin(\pi x_1) \sin^{20}(\pi x_1^2) + \sin(\pi x_2) \sin^{20}(2\pi x_2^2), \quad (19)$$

with $\mathbf{x} \in [0, 1]^2$. This function is more challenging than the Cosine function as it exhibits large areas of “flatland” (i.e., many mediocre bandits) and a very deep and narrow well located slightly off center (i.e., rare bandits with extreme payoffs). For $n_{GMM} = 4$, Figure 2b shows that LW-UCB outperforms the competition by a substantial margin. Figure 2b also makes it visually clear that the likelihood ratio assigns more weight to the best bandits. Interestingly, we have found that the likelihood ratio sometimes discovers a broader area where other sub-optimal solutions are also captured.

For an even more challenging test case, we introduce a modified version of the Michalewicz function which features multiple small “islands” associated with extreme payoffs. Specifically, the function

$$f(\mathbf{x}) = \sin(\pi x_1) \sin^{20}(2\pi x_1^2) + \sin(\pi x_2) \sin^{20}(3\pi x_2^2) \quad (20)$$

has six extreme local minima and a number of steep valleys in the domain $\mathbf{x} \in [0, 1]^2$, making it quite difficult for the algorithms to identify the best bandits. Figure 2c shows that despite the added difficulty, LW-UCB again exhibits outstanding convergence behavior, with the other acquisition functions struggling to identify the best bandits and therefore yielding poor performance. We also note that the likelihood ratio not only emphasizes the best area for rewards but is also able to identify sub-optimal solutions of somewhat lesser importance, demonstrating the ability of our approach to provide a good balance between exploration and exploitation.

To investigate the effect of the likelihood ratio on runtime, we record the time required to perform one iteration of the Bayesian algorithm. (This includes training the GP model, computing the likelihood ratio and the GMM approximation for LW-UCB, and optimizing the acquisition function.) Consistent with Blanchard and Sapsis [2021], Table S1 in the Supplementary Material shows that the runtimes for LW-UCB are on the same order of magnitude as the other criteria. The additional cost is attributable to the computation and sampling of the likelihood ratio, and presumably can be alleviated using recent advances in sampling methods for GP posteriors [Wilson et al., 2020].

We have also investigated the sensitivity of the LW-UCB criterion to the size of the Gaussian mixture model used in the approximation of the likelihood ratio. For the three synthetic functions (18)–(20), we repeated the experiments with two additional values of n_{GMM} . Figure S1 in the Supplementary Material shows that the performance of LW-UCB is essentially independent of the number of Gaussian components used in (15) when the latent function is relatively simple, and that larger values of n_{GMM} are preferable when the complexity of the landscape grows and the number of optimal regions increases.

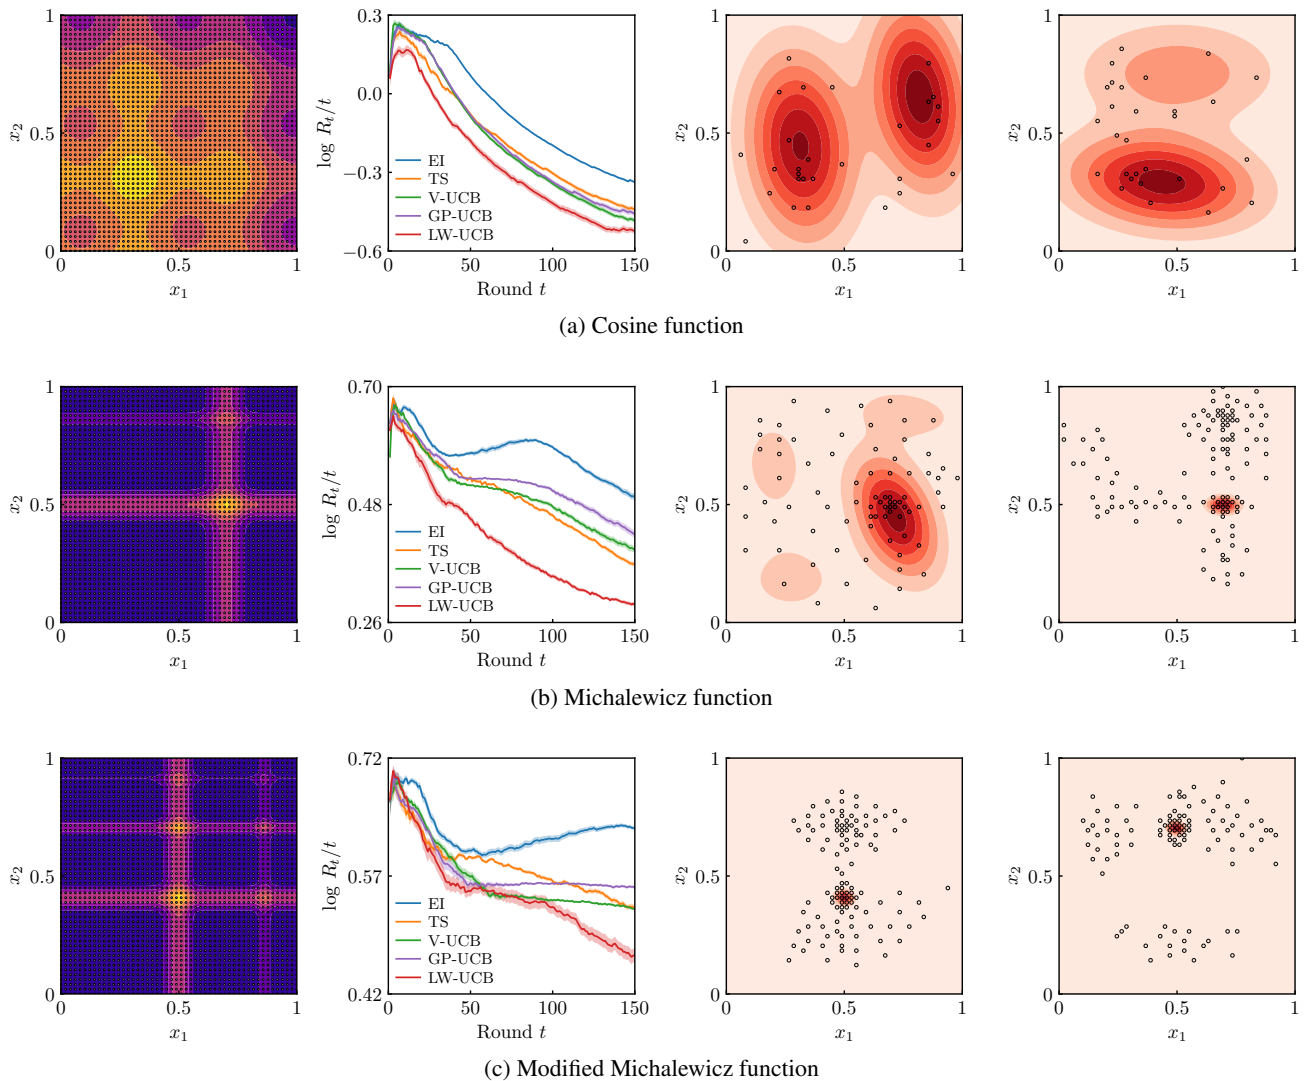


Figure 2: *Synthetic benchmarks*. From left to right: locations of the bandits (white circles) and associated rewards (background color); cumulative regret for various acquisition functions; for two representative trials of LW-UCB, distribution of the likelihood ratio (background color) learned by the GP model from the visited bandits (open circles) after $t = 150$ rounds.

3.2 A SYSTEMATIC STUDY: WHEEL BANDITS

In this section we consider a variant of the contextual wheel bandit problem discussed in Riquelme et al. [2018]. The feasible domain is the unit disk ($0 \leq r \leq 1$) which is divided into five disjoint sectors. The inner disk ($0 \leq r \leq \rho$) is sub-optimal with reward 0.2. The upper left, lower right, and lower left quadrants of the outer ring ($\rho \leq r \leq 1$) are also sub-optimal, with rewards 0.05, 0.1, and 0, respectively (Figure 3). The optimal bandits are located in the upper right quadrant of the outer ring and return a reward of 1, significantly higher than the other quadrants. The parameter ρ determines the difficulty of the problem. For small ρ , the optimal region accounts for a large fraction of the domain, while for large ρ the difficulty significantly increases. We generate the bandits on a 70×70 uniform grid and retain

those lying inside the unit disk. Each bandit produces noisy rewards with $\sigma_n = 10^{-3}$.

For $n_{GMM} = 4$, Figure 3 shows that the proposed LW-UCB criterion leads to significant gains in performance compared to conventional acquisition functions, especially as the value of ρ increases and the optimal bandits become scarcer. Figure 3 also shows that the attention mechanism embedded in the likelihood ratio encourages exploration of the extreme-reward region. It is also interesting to note that in all cases investigated, the expected improvement, Thompson sampling, V-UCB, and GP-UCB deliver nearly identical performance, even in the asymptotic regime, unlike LW-UCB which provides consistently faster convergence.

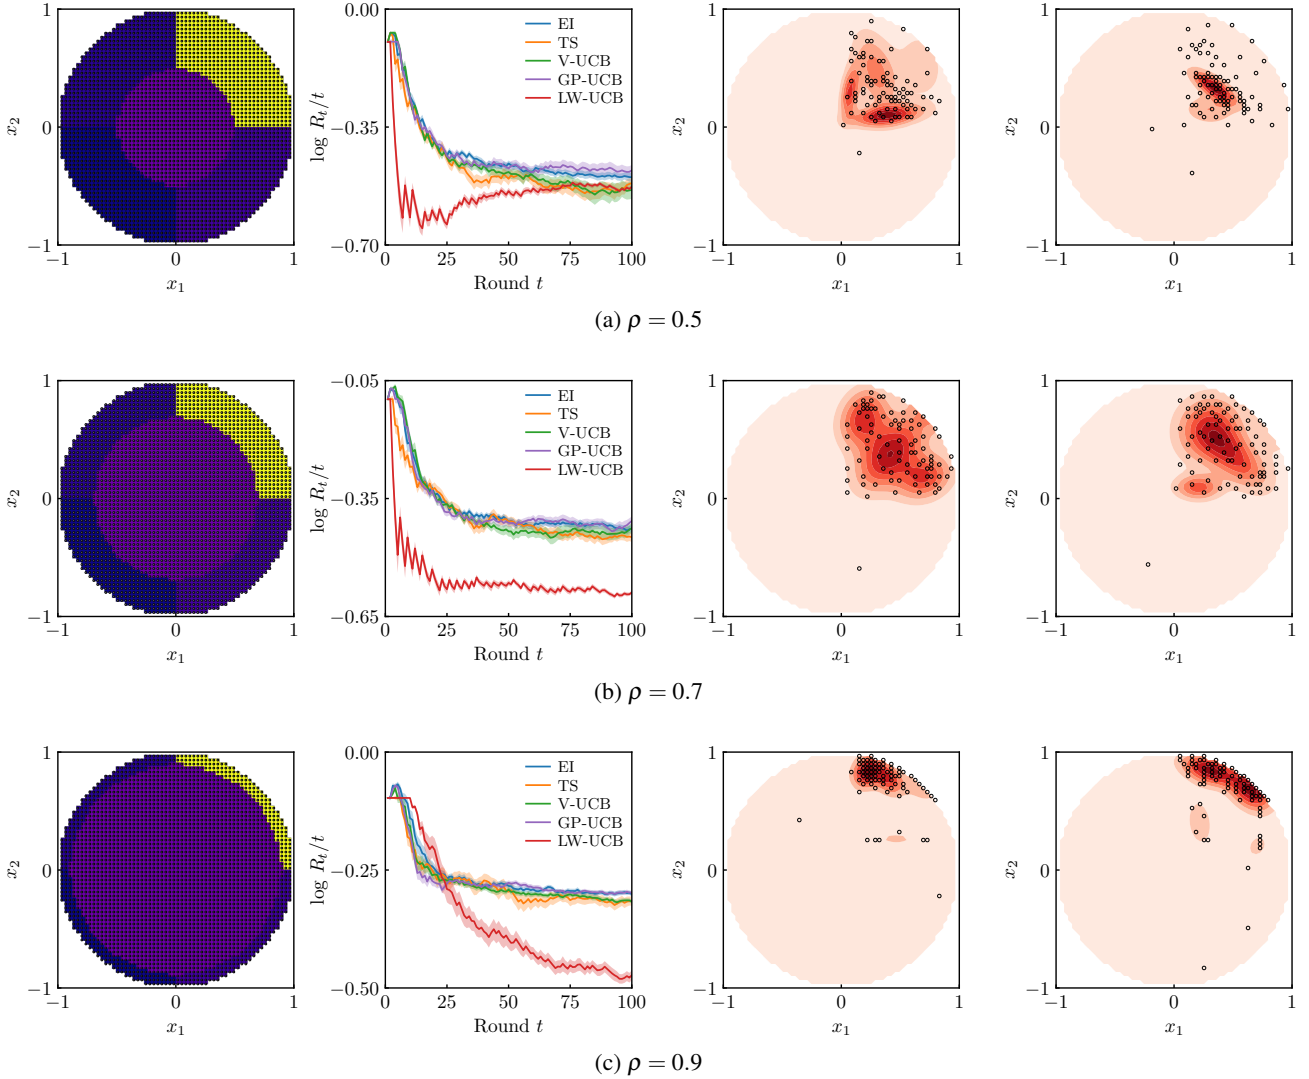


Figure 3: *Wheel bandit problem*. From left to right: locations of the bandits (white circles) and associated rewards (background color); cumulative regret for various acquisition functions; and for two representative trials of LW-UCB, distribution of the likelihood ratio (background color) learned by the GP model from the visited bandits (open circles) after $t = 100$ rounds.


3.3 SPATIO-TEMPORAL ENVIRONMENT MONITORING WITH SENSOR NETWORKS

Finally, we demonstrate the approach using the real-world dataset considered in Srinivas et al. [2009]. The dataset² contains temperature measurements collected by 46 sensors deployed in the Intel Berkeley Research lab (Figure 4a). As in Srinivas et al. [2009], our goal is to find locations of highest temperature by sequentially activating the available sensors while using as few sensor switches as possible in order to save electric power. Our working dataset consists of 500 temperature snapshots collected every ten minutes over a three-day period. For each temperature snapshot, we initialize the algorithm by randomly activating $n = 3$

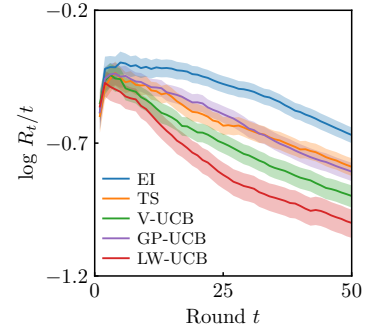
sensors. The sensors (i.e., the bandits) produce rewards that are corrupted by small Gaussian noise with $\sigma_n = 10^{-4}$. We use $n_{GMM} = 2$ for the GMM approximation of the likelihood ratio.

For this real-world problem, Figure 4b shows that LW-UCB performs better than the other acquisition schemes. Figures 4c–4f show that the likelihood ratio draws the algorithm’s attention to the bandits whose rewards are high by artificially inflating the model uncertainty for these bandits. We note that, in contrast to the examples considered previously, here the bandits are few and far between. For instance, there is no sensor data available in the server room and the stairwell (see Figure 4a). Because of the sparsity of the data, finding the best sensor to activate is more challenging. But this does not seem to negatively affect the LW-UCB acquisition



²<http://db.csail.mit.edu/labdata/labdata.html>





(a)





(b)





(c)



(d)



(e)



(f)

Figure 4: *Spatio-temporal monitoring with sensor networks.* (a) Sensor locations; (b) cumulative regret for various acquisition functions; and (c–f) for four representative trials of LW-UCB, spatial distribution of temperature (left panel) and the likelihood ratio (right panel) learned by the GP model from the activated sensors (circles) after $t = 50$ rounds.

criterion, which is able to identify and explore the relevant areas more intelligently than the other acquisition functions.

4 CONCLUSIONS

We have proposed a novel output-weighted acquisition function (LW-UCB) for sequential decision making. Our approach leverages the information provided by the GP regression model to regularize uncertainty and favor exploration of abnormally large payoff values. The regularizer takes the form of a sampling weight—the likelihood ratio—and can be efficiently approximated by a Gaussian mixture model. The likelihood ratio provides a principled way to balance exploration and exploitation in multi-armed bandit optimization problems where the goal is to maximize the cumulative reward. The benefits of the proposed method have been systematically established via several benchmark examples which demonstrated superiority of our method compared to classical acquisition functions (expected improvement, Thompson sampling, and two variants of UCB).

Though the proposed LW-UCB criterion yields superior performance in bandit problems, several questions remain open. First, a theoretical analysis of the convergence behavior of LW-UCB is needed, in the same way that information gain has helped characterize the convergence of GP-UCB [Srinivas et al., 2009, Krause and Ong, 2011]. The second avenue is to investigate more complex cases with high-dimensional contexts and multi-output GP priors. The latter can be readily accommodated in our JAX implementation which leverages automatic differentiation to allow efficient gradient-based optimization of the LW-UCB criterion for arbitrary GP priors. The third question has to do with extending the proposed approach to other Bayesian inference schemes, e.g., Bayesian linear regression [Chu et al., 2011], Bayesian neural networks [Riquelme et al., 2018], and variational inference [Hoffman et al., 2013]. Finally, there is the question of how to adapt the proposed framework for use in more general Markov decision processes and reinforcement learning problems [Sutton and Barto, 2018] where contextual information is typically high-dimensional and rewards are obtained after multiple trials rather than instantaneously.

5 BACK MATTER

Author Contributions

Y.Y., A.B., T.S and P.P conceived the study, implemented the methods, performed the simulations, and wrote the manuscript.

Acknowledgements

Y.Y. and P.P. received support from the US Department of Energy under the Advanced Scientific Computing Research program (Grant No. DE-SC0019116) and the Air Force Office of Scientific Research (Grant No. FA9550-20-1-0060). A.B. and T.S. would like to thank the support from the AFOSR-MURI Grant No. FA9550- 21-1-0058 and the ARO-MURI Grant No. W911NF-17-1-0306.

References

- Rajeev Agrawal. Sample mean based index policies with $O(\log n)$ regret for the multi-armed bandit problem. *Advances in Applied Probability*, pages 1054–1078, 1995.
- Shipra Agrawal and Navin Goyal. Analysis of Thompson sampling for the multi-armed bandit problem. In *Conference on Learning Theory*, pages 39–1, 2012.
- Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3:397–422, 2002.
- Javad Azimi, Alan Fern, and Xiaoli Fern. Batch Bayesian optimization via simulation matching. In *Advances in Neural Information Processing Systems*, pages 109–117, 2010.
- Kamyar Azizzadenesheli, Emma Brunskill, and Animashree Anandkumar. Efficient exploration through Bayesian deep q-networks. In *2018 Information Theory and Applications Workshop (ITA)*, pages 1–9. IEEE, 2018.
- Atilim Gunes Baydin, Barak A Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. Automatic differentiation in machine learning: a survey. *arXiv preprint arXiv:1502.05767*, 2015.
- Antoine Blanchard and Themistoklis Sapsis. Informative path planning for anomaly detection in environment exploration and monitoring. *arXiv preprint arXiv:2005.10040*, 2020a.
- Antoine Blanchard and Themistoklis Sapsis. Output-weighted importance sampling for Bayesian experimental design and uncertainty quantification. *arXiv preprint arXiv:2006.12394*, 2020b.
- Antoine Blanchard and Themistoklis Sapsis. Bayesian optimization with output-weighted importance sampling. *Journal of Computational Physics*, 425:109901, 2021.
- Mohamed Amine Bouhlef, Nathalie Bartoli, Abdelkader Otsmane, and Joseph Morlier. Improving kriging surrogates of high-dimensional design models by partial least squares dimension reduction. *Structural and Multidisciplinary Optimization*, 53:935–952, 2016.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- Olivier Chapelle and Lihong Li. An empirical evaluation of Thompson sampling. In *Advances in Neural Information Processing Systems*, pages 2249–2257, 2011.
- Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics*, pages 208–214, 2011.
- Varsha Dani, Thomas Hayes, and Sham Kakade. Stochastic linear optimization under bandit feedback. In *The 21st Annual Conference on Learning Theory*, pages 355–366, 2008.
- Richard Dearden, Nir Friedman, and Stuart Russell. Bayesian Q-learning. In *AAAI/IAAI*, pages 761–768, 1998.
- Alexander Forrester, András Sóbester, and Andy Keane. Multi-fidelity optimization via surrogate modelling. *Proceedings of the Royal Society A*, 463:3251–3269, 2007.
- Alex Graves. Practical variational inference for neural networks. In *Advances in Neural Information Processing Systems*, pages 2348–2356, 2011.
- Matthew Hoffman, David Blei, Chong Wang, and John Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 14, 2013.
- Kirthevasan Kandasamy, Akshay Krishnamurthy, Jeff Schneider, and Barnabás Póczos. Parallelised Bayesian optimisation via Thompson sampling. In *International Conference on Artificial Intelligence and Statistics*, pages 133–142, 2018.
- Emilie Kaufmann, Nathaniel Korda, and Rémi Munos. Thompson sampling: An asymptotically optimal finite-time analysis. In *International Conference on Algorithmic Learning Theory*, pages 199–213, 2012.

- Jaya Kawale, Hung Bui, Branislav Kveton, Long Tran-Thanh, and Sanjay Chawla. Efficient Thompson sampling for online matrix-factorization recommendation. In *Advances in Neural Information Processing Systems*, pages 1297–1305, 2015.
- Andreas Krause and Cheng Soon Ong. Contextual Gaussian process bandit optimization. In *Advances in Neural Information Processing Systems*, pages 2447–2455, 2011.
- Jan Leike, Tor Lattimore, Laurent Orseau, and Marcus Hutter. Thompson sampling is asymptotically optimal in general environments. *arXiv preprint arXiv:1602.07905*, 2016.
- Chunyuan Li, Ke Bai, Jianqiao Li, Guoyin Wang, Changyou Chen, and Lawrence Carin. Adversarial learning of a sampler based on an unnormalized distribution. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 3302–3311, 2019.
- Lihong Li, Wei Chu, John Langford, and Robert Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web*, pages 661–670, 2010.
- Dong Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45:503–528, 1989.
- Mustafa Mohamad and Themistoklis Sapsis. Sequential sampling strategy for extreme event statistics in nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 115:11138–11143, 2018.
- Ian Osband, Charles Blundell, Alexander Pritzel, and Benjamin Van Roy. Deep exploration via bootstrapped dqn. In *Advances in Neural Information Processing Systems*, pages 4026–4034, 2016.
- Paris Perdikaris, Daniele Venturi, and George Em Karniadakis. Multifidelity information fusion algorithms for high-dimensional systems and massive data sets. *SIAM Journal on Scientific Computing*, 38(4):B521–B538, 2016.
- Chao Qin, Diego Klabjan, and Daniel Russo. Improving the expected improvement algorithm. In *Advances in Neural Information Processing Systems*, volume 30, pages 5381–5391, 2017.
- Carl Edward Rasmussen and Christopher Williams. *Gaussian processes for machine learning*. MIT Press, Cambridge, MA, 2006.
- Carlos Riquelme, George Tucker, and Jasper Snoek. Deep bayesian bandits showdown: An empirical comparison of bayesian deep networks for thompson sampling. *arXiv preprint arXiv:1802.09127*, 2018.
- Daniel Russo and Benjamin Van Roy. An information-theoretic analysis of Thompson sampling. *The Journal of Machine Learning Research*, 17:2442–2471, 2016.
- Daniel Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, and Zheng Wen. A tutorial on Thompson sampling. *arXiv preprint arXiv:1707.02038*, 2017.
- Jerome Sacks, William Welch, Toby Mitchell, and Henry Wynn. Design and analysis of computer experiments. *Statistical Science*, pages 409–423, 1989.
- U.K. Saha, S. Thotla, and D. Maity. Optimum design configuration of Savonius rotor through wind tunnel experiments. *Journal of Wind Engineering and Industrial Aerodynamics*, 96:1359–1375, 2008.
- Francisco Sahli Costabal, Paris Perdikaris, Ellen Kuhl, and Daniel Hurtado. Multi-fidelity classification using gaussian processes: Accelerating the prediction of large-scale computational models. *Computer Methods in Applied Mechanics and Engineering*, 357:112602, 2019.
- Francisco Sahli Costabal, Yibo Yang, Paris Perdikaris, Daniel Hurtado, and Ellen Kuhl. Physics-informed neural networks for cardiac activation mapping. *Frontiers in Physics*, 8:42, 2020.
- Themistoklis Sapsis. Output-weighted optimal sampling for Bayesian regression and rare event statistics using few samples. *Proceedings of the Royal Society A*, 476: 20190834, 2020.
- Soumya Sarkar, Sudeepta Mondal, Michael Joly, Matthew Lynch, Shaunak Bopardikar, Ranadip Acharya, and Paris Perdikaris. Multifidelity and multiscale Bayesian framework for high-dimensional engineering design and calibration. *Journal of Mechanical Design*, 141, 2019.
- Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*, 2015.
- Songqing Shan and G Gary Wang. Survey of modeling and optimization strategies to solve high-dimensional design problems with computationally-expensive black-box functions. *Structural and Multidisciplinary Optimization*, 41:219–241, 2010.
- Niranjana Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.
- Richard Sutton and Andrew Barto. *Reinforcement learning: An introduction*. MIT Press, Cambridge, MA, 2018.
- William Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25:285–294, 1933.

Emmanuel Vazquez and Julien Bect. Convergence properties of the expected improvement algorithm with fixed mean and covariance functions. *Journal of Statistical Planning and Inference*, 140:3088–3095, 2010.

Zhong Yi Wan, Pantelis Vlachas, Petros Koumoutsakos, and Themistoklis Sapsis. Data-assisted reduced-order modeling of extreme events in complex dynamical systems. *PLOS One*, 13:e0197704, 2018.

James Wilson, Viacheslav Borovitskiy, Alexander Terenin, Peter Mostowsky, and Marc Deisenroth. Efficiently sampling functions from gaussian process posteriors. In *International Conference on Machine Learning*, pages 10292–10302. PMLR, 2020.

SUPPLEMENTARY MATERIAL

For the synthetic test functions considered in Section 3.1, we provide results on a) the effect of the likelihood ratio on computational runtime, and b) the sensitivity of the cumulative regret with respect to the size of the Gaussian mixture model used in LW-UCB.

Table S1: Single-iteration runtime (in seconds) averaged over ten experiments.

	Cosine	Michalewicz	Modified Michalewicz
EI	0.49	0.52	0.68
TS	0.55	0.53	0.63
V-UCB	1.36	1.28	1.50
GP-UCB	1.36	1.28	1.50
LW-UCB	4.19	3.94	4.51




Figure S1: For the synthetic functions in Section 3.1, performance of LW-UCB with various values of n_{GMM} compared to the other acquisition functions considered in this work.