Movement Penalized Bayesian Optimization with Application to Wind Energy Systems

Anonymous Author(s)
Affiliation
Address
email

Abstract

Contextual Bayesian optimization (CBO) is a powerful framework for sequential decision-making given side information, with important applications, e.g., in wind energy systems. In this setting, the learner receives context (e.g., weather conditions) at each round, and has to choose an action (e.g., turbine parameters). Standard algorithms assume no cost for switching their decisions at every round. However, in many practical applications, there is a cost associated with such changes, which should be minimized. We introduce the episodic CBO with movement costs problem and, based on the online learning approach for metrical task systems of Coester and Lee [17], propose a novel randomized mirror descent algorithm that makes use of Gaussian Process confidence bounds. We compare its performance with the offline optimal sequence for each episode and provide rigorous regret guarantees. We further demonstrate our approach on the important real-world application of altitude optimization for Airborne Wind Energy Systems. In the presence of substantial movement costs, our algorithm consistently outperforms standard CBO algorithms.

1 Introduction

Bayesian optimization (BO) is a well-established framework for sequential black-box function optimization that relies on Gaussian Process (GP) models [39] to sequentially learn and optimize the unknown objective. In many practical scenarios, however, one wants to additionally use available *contextual* information when making decisions. In this setting, at each round, the learner receives a context from the environment and has to choose an action based upon it. Previous works have developed contextual BO algorithms [29, 15, 28, 37], and applied them to various important applications, e.g., vaccine design, nuclear fusion, database tuning, crop recommender systems, etc.

A potential practical issue with these standard algorithms is that they assume no explicit costs for *switching* between their actions at every round. Frequent action changes can be extremely costly in many real-world applications. This work is motivated by the problem of real-time control of the altitude of an airborne wind energy (AWE) system.[1] In AWE systems, the wind speed is often only measurable at the system's altitude, and determining the optimal operating altitude of an AWE system as the wind speed varies represents a challenging problem. Another fundamental challenge is that additional energy is required for adjusting the altitude, which makes the frequent altitude changes costly. Consequently, this work is motivated by the following question: *How can we efficiently learn to optimize the AWE system's operating altitude despite varying wind conditions while minimizing the energy cost associated with turbine altitude changes?*

[1] AWE system is a wind turbine with a rotor supported in the air without a tower that can benefit from the persistence of wind at different high altitudes [20].

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

In this work, we formalize the *movement penalized* contextual BO problem. When the switching cost is a *metric* (distance function), we propose a novel algorithm that effectively combines ideas from BO with the online learning strategies proposed in [17] for solving the so-called *metrical task system* (MTS) problem [11]. Furthermore, our algorithm relies solely on noisy point evaluations (i.e., bandit feedback), allows for arbitrary context sequences, and besides the standard exploration-exploitation trade-off, it also balances the movement costs. As a result, it outperforms the standard movement-cost-agnostic contextual BO algorithms as well as movement-conservative baselines (see Fig. 1).

Related Work. Bayesian optimization (BO) refers to a sequential approach for optimizing an unknown objective (cost function) from noisy point evaluations. A great number of BO methods have been developed over the years (e.g., [35, 47, 16]). While the focus of standard BO approaches is mainly on optimizing the unknown objective cost function, in this work, we additionally focus on penalizing frequent action changes. This problem makes the most sense in the *contextual* BO setting [29], where the main objective changes with the observed contextual information and the learner also seeks to minimize the cost associated with frequent changes of its actions. Several works have tried to incorporate such and similar cost functions in the BO setting. [32] explicitly consider switching costs based on how deep into the pipeline the change in variable occurs. Other works include cost-aware and multi-objective sampling strategies in various settings such as batch [26, 24], multi-fidelity [10, 25, 46], multi-objective optimization [2] and dynamic programming [31, 30]. Similarly, we consider movement/switching costs, but unlike previous work, we specifically focus on minimizing the movement costs in the contextual setting.

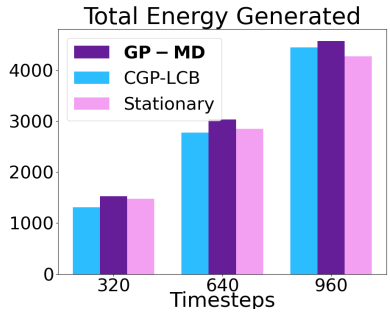

Figure 1: Total energy generated by the AWE system when operated with GP-MD (proposed in this work), CGP-LCB [29] and a stationary baseline. The stationary baseline employs no-learning and does not incur movement costs. GP-MD and CGP-LCB learn an operating strategy, but GP-MD outperforms CGP-LCB since it also considers movement costs.

The *Metrical Task Systems* (MTS) problem [11] is a sequential decision-making problem widely studied in the *online learning* literature. At each round, the learner observes a service cost function, chooses an action, and incurs the corresponding cost together with a movement cost penalizing the distance (according to some metric) between the current action and the one chosen at the previous round. The MTS problem is directly related to our problem setting (see Section 2.2). After a long series of works, in [14] and [17], an $\mathcal{O}\big((\log n)^2\big)$-competitive algorithm for MTS on any finite metric space was established. The approach of Coester and Lee [17] relies on a tree representation of the decision space and an action randomization scheme via a mirror descent procedure. In contrast to our setting, these works assume that the service costs are *known* (i.e., they assume the *full-information* feedback). Moreover, in the online optimization literature, other related works study the online *convex* optimization with switching costs [23, 45, 22, 33] and convex body chasing problems [12, 5, 40, 13]. We make no use of convexity and take a model-based (GP) approach to learn about the unknown service costs. [34] consider non-convex objective functions, however, they impose certain conditions on the objective that are not easily modeled via GP as done in our work.

The use of GP confidence bounds in online learning settings has been previously explored, e.g., in repeated multi-agent [41, 44] and sequential games [43], and to discover randomized max-min strategies [42]. However, none of these works has considered movement costs in the objective. This makes our problem significantly different from the aforementioned ones, and requires a suitable action randomization scheme that can trade-off exploration, exploitation, and movement costs simultaneously.

Various works have applied Bayesian Optimization in the context of *wind energy systems* before. [38, 19, 4] use BO techniques to maximize the total energy yield in a wind farm in a cooperative or a closed-loop framework. [36] use it to tune the parameters of the wind turbine to learn the effective wind speed. [6] use BO for plant design (i.e., physical system design) of airborne wind energy systems. [48] apply BO and other regression techniques to predict short-term wind energy production to make informed production offers. Finally, [3] survey several different methods for efficient wind-power prediction using machine-learning methods including BO. Our problem formulation differs from the above mentioned works, since we explicitly consider movement energy loss caused by the altitude change of the system.

Contributions. We formally introduce Bayesian optimization with movement costs and propose a novel GP-MD algorithm (in Algorithm 1). GP-MD combines the online mirror descent (MD) algorithm with shrinking Gaussian Process (GP) confidence bound to decide on which point to evaluate next. In our theoretical analysis, we establish rigorous sublinear regret guarantees for our algorithm by combining techniques from Bayesian optimization and metrical task systems approaches [17]. Finally, we demonstrate that GP-MD is able to successfully outperform previous contextual Bayesian optimization approaches on both synthetic and real-world data in the presence of movement costs. In particular, we consider the application to airborne wind energy systems and demonstrate that GP-MD can effectively operate such a system by considering movement costs and varying environmental conditions.

2 Problem Statement

Let $f : \mathcal{X} \times \mathcal{E} \to \mathbb{R}_+$ be an *unknown* cost function defined over $\mathcal{X} \times \mathcal{E} \subset \mathbb{R}^p$, where $\mathcal{X}$ is a finite set of actions, i.e., $|\mathcal{X}| = n$, and $\mathcal{E}$ represents convex and compact space of contexts. We denote the *known* metric (i.e., distance function) of $\mathcal{X}$ as $d(\cdot, \cdot)$, and similarly to other works in Bayesian optimization (e.g., [47, 16]) assume that the target cost function f belongs to a reproducing kernel Hilbert space (RKHS) $\mathcal{H}_k$ of functions (defined on $\mathcal{X} \times \mathcal{E}$), that corresponds to a known kernel $k : (\mathcal{X} \times \mathcal{E}) \times (\mathcal{X} \times \mathcal{E}) \to \mathbb{R}_+$ with $k((x, e), (x', e')) \leq 1$ for any action-context pair. In particular, we assume that for some known $B > 0$, the target cost f has a bounded RKHS norm, i.e., $f \in \mathcal{F}_k = \{f \in \mathcal{H}_k : \|f\|_k \leq B\}$. Also, we assume that the diameter of $\mathcal{X}$ ($\max_{x,x' \in \mathcal{X}} d(x, x')$) is bounded and denote it by ψ.

We consider an episodic setting, wherein each episode runs over a finite time horizon H. Let the *initial state* of the system in the first episode correspond to action $x_{0,1} \in \mathcal{X}$. At the end of every episode, the system resets to a new given initial action $x_{0,m} \in \mathcal{X}$ where $m \in \{1, 2, \ldots, N_{ep}\}$ denotes the episode index. In each episode m and at every time step $h \in \{1, 2, \ldots, H\}$, the environment reveals the context $e_{h,m} \in \mathcal{E}$ to the learner. We make no assumptions on the context sequence provided by the environment (i.e., it can be arbitrary and different across episodes). The learner then chooses $x_{h,m} \in \mathcal{X}$ and observes the noisy function value:

$$y_{h,m} = f(x_{h,m}, e_{h,m}) + \xi_{h,m}, \tag{1}$$

where $\xi_{h,m} \sim \mathcal{N}(0, \sigma^2)$ with known σ, and independence over time steps. The goal of the learner is to minimize the cost incurred over the rounds in every episode, but at the same time to minimize the distance between its subsequent decisions as measured by $d(x_{h-1,m}, x_{h,m})$.

Let $D_m = \{x_{1,m}, x_{2,m}, \ldots, x_{H,m}\}$ denote the set of actions chosen by the learner over H rounds in episode m. We recall that each action $x_{h,m} \in D_m$ is chosen after observing the corresponding context $e_{h,m}$. The objective is to minimize the cumulative episodic cost for each episode m,

$$\text{cost}_m(D_m) = \underbrace{\sum_{h=1}^{H} f(x_{h,m}, e_{h,m})}_{S_m(D_m)} + \underbrace{\sum_{h=1}^{H} d(x_{h,m}, x_{h-1,m})}_{M_m(D_m)}, \tag{2}$$

where we refer to the two terms in Eq. (2) as *service cost* S_m and *movement cost* M_m.

When f is known, the problem can be seen as a MTS instance as detailed in Section 2.2. Even in such a case, we cannot hope to solve this problem optimally, and nearly-optimal approximate algorithms were recently proposed (see Coester and Lee [17]). Hence, the learner's performance in episode m is measured via (α, β)-approximate regret:

$$r_m^{\alpha,\beta} = \text{cost}_m(D_m) - \alpha \cdot \text{cost}_m(D_m^*) - \beta, \tag{3}$$

where $D_m^* := \arg\min_{D \subset \mathcal{X},\, |D|=H} \text{cost}_m(D)$ is the offline optimal action sequence obtained assuming the knowledge of the true sequence of contexts $\{e_{h,m}\}_{h=1}^{H}$ *in advance*, and α and β are approximation constants (independent of N_{ep}). In contrast, in our setting, the learner only gets to see the *current context* when making a decision and has no knowledge about the future ones.

After N_e episodes, the total cumulative regret is defined as

$$R_{N_{ep}}^{\alpha,\beta} = \sum_{m=1}^{N_{ep}} r_m^{\alpha,\beta}. \tag{4}$$

We seek an algorithm whose total cumulative regret grows sublinearly in N_{ep}, so that $\lim_{N_{ep}\to\infty} R^{\alpha,\beta}_{H,N_{ep}}/N_{ep} = 0$, for any set of initial states $\{x_{0,m}\}_{m=1}^{N_{ep}} \subset \mathcal{X}$.

2.1 Gaussian Process Model

In standard Bayesian optimization, a surrogate Gaussian Process model is typically used to model the target cost function. A Gaussian Process $GP(\mu(\cdot), k(\cdot,\cdot))$ over the input domain $\mathcal{X}\times\mathcal{E}$, is a collection of random variables $(f(x,e))_{x\in\mathcal{X},e\in\mathcal{E}}$ where every finite number of them $(f(x_i,e_i))_{i=1}^n$, $n\in\mathbb{N}$, is jointly Gaussian with mean $\mathbb{E}[f(x_i)] = \mu(x_i,e_i)$ and covariance $\mathbb{E}[(f(x_i,e_i)-\mu(x_i,e_i))(f(x_j,e_j)-\mu(x_j,e_j))] = k((x_i,e_i),(x_j,e_j))$ for every $1\le i,j\le n$.

BO algorithms typically use zero-mean GP priors to model uncertainty in f, i.e., $f\sim GP(0,k(\cdot,\cdot))$, and Gaussian likelihood models for the observed data. As more data points are observed, GP (Bayesian) posterior updates are performed in which noise variables are assumed to be drawn independently across t from $\mathcal{N}(0,\lambda)$. Here, λ is a hyperparameter that might be different from the true noise variance σ^2. More precisely, given a sequence of previously queried points and their noisy observations the posterior is again Gaussian, with the posterior mean and variance given by:

$$\mu_t(x,e) = k_t(x,e)^T(K_t+\lambda I_t)^{-1}Y_t \tag{5}$$

$$\sigma_t^2(x,e) = k((x,e),(x,e)) - k_t(x,e)^T(K_t+\lambda I_t)^{-1}k_t(x,e), \tag{6}$$

where $Y_t := [y_1,\dots,y_t]$ denotes a vector of observations, $K_t = [k((x_s,e_s),(x_{s'},e_{s'})]_{s,s'\le t}$ is the corresponding kernel matrix, and $k_t(x,e) = [k((x_1,e_1),(x,e)),\dots,k((x_t,e_t),(x,e))]^T \in \mathbb{R}^{t\times 1}$.

Maximum Information Gain. In the standard Bayesian optimization, the main quantity that characterizes the complexity of optimizing the target cost function is the maximum information gain [47] defined at time t as:

$$\gamma_t = \max_{\{(x_i,e_i)\}_{i=1}^t} I(Y_t;f), \tag{7}$$

where $I(Y_t;f)$ denotes the mutual information between random observations Y_t and GP model f. The mutual information for the GP model is given as:

$$I(Y_t;f) = \frac{1}{2}\log\det(I_t+\lambda^{-1}K_t). \tag{8}$$

This quantity is kernel-specific and for compact and convex domains γ_t is sublinear in t for various classes of kernel functions [47] as well as for kernel compositions (e.g., additive kernels in [29]).

Confidence Bounds. We also use the following result ([47, 1, 16]) that is frequently used in Bayesian optimization to provide confidence bounds around the unknown function.

Lemma 1. *Assume the σ-sub-Gaussian noise model as in Eq. (1), and let f belong to $\mathcal{F}_k$. Then, the following holds with probability at least $1-\delta$ simultaneously over all $t\ge 1$ and $x\in\mathcal{X}$, $e\in\mathcal{E}$:*

$$|\mu_t(x,e) - f(x,e)| \le \beta_t\sigma_t(x,e), \tag{9}$$

where $\beta_t = \frac{\sigma}{\lambda^{1/2}}\sqrt{2\ln(1/\delta)+2\gamma_t} + B$, and μ_t and σ_t are defined in Eqs. (5) and (6) with $\lambda > 0$.

Based on the previous, we also define the lower confidence bound for every $x\in\mathcal{X}, e\in\mathcal{E}$ as:

$$\text{lcb}_t(x,e) := \mu_t(x,e) - \beta_t\sigma_t(x,e). \tag{10}$$

We use $\text{lcb}_m(x,e)$ when it is computed based on data collected before episode m.

2.2 Relation to Metrical Task Systems (MTS)

When f is known, our optimization objective in Eq. (2) can be seen as a particular type of MTS problem, where $f(\cdot,e_{h,m})$ is the MTS service cost that changes for every h and m. Compared to a standard MTS (see Appendix A), our problem formulation is more challenging since the learner can only learn about f from previously observed data. The approach proposed in this paper builds on the algorithm by Coester and Lee [17] for standard MTS problems. However, to cope with the aforementioned challenge, our approach exploits the regularity assumptions regarding f and utilizes the constructed lower confidence bounds Eq. (10) to *hallucinate* information about the unavailable

Algorithm 1 GP-MD

Require: Action space $\mathcal{X}$, kernel function $k(\cdot,\cdot)$, metric $d(\cdot,\cdot)$
Run FRT$(\mathcal{X}, d(\cdot,\cdot))$ and obtain τ-HST $\mathcal{T} = (V, E)$ with leaves $\mathcal{L} = \mathcal{X}$
for $m = 1, \dots, N_{ep}$ **do**
  Receive $x_{0,m}$ and initialize $z_{0,m}$ (Eq. (27)), conditional prob. $q_0 = \Delta^{-1}(z_{0,m})$ as in Eq. (15)
  for $h = 1, \dots, H$ **do**
    Observe context $e_{h,m}$ and initialize costs: $\text{lcb}_m(v, e_{h,m}) = 0, \forall v \in V \backslash \mathcal{L}$
    for $u \in \mathcal{OD}(V \backslash \mathcal{L})$ **do**
      Update vertex prob. $q_h^{(u)}$ from $q_{h-1}^{(u)}$ and $\text{lcb}_m(\cdot, e_{h,m})$ via Mirror Descent (Eq. (13))
      Update cost for vertex u:

$$\text{lcb}_m(u, e_{h,m}) = \langle q_h^{(u)}, \text{lcb}_m(\cdot, e_{h,m}) \rangle = \sum_{\nu \in \mathcal{C}(u)} q_{h,\nu} \cdot \text{lcb}_m(\nu, e_{h,m})$$

    end for
    Compute prob. vector $z_{h,m} = \Delta(q_h)$ (Eq. (15)) and leaves' prob. $l(z_{h,m})$ (Eq. (11))
    Estimate optimal coupling $\zeta_{h-1,h,m}$ between $l(z_{h-1,m})$ and $l(z_{h,m})$ as in Eq. (12)
    Sample action $x_{h,m} \sim \zeta_{h-1,h,m}(\cdot | x_{h-1,m})$ and observe $y_{h,m} = f(x_{h,m}, e_{h,m}) + \xi_{h,m}$
  end for
  Update $\mu_{m+1}(\cdot,\cdot)$ and $\sigma_{m+1}(\cdot,\cdot)$ as per Eq. (5) and Eq. (6)
end for

service cost at each round. Before presenting our overall approach, we describe a preliminary step proposed by [17], which consists of representing our metric space $(\mathcal{X}, d)$ by a *Hierarchically Separated Tree* (HST) metric space.

HST metric space. Consider a tree $\mathcal{T} = (V, E)$ with root r, leaves $\mathcal{L} \subset V$ and non-negative weights w_v, for each $v \in V$, which are non-increasing along root-leaf paths. Let $d_{\mathcal{T}}(l, l')$ denote a distance metric between any two leaves $l, l' \in \mathcal{L}$ given as the sum of the encountered weights on the path from l to l' (see Fig. 4). $(\mathcal{L}, d_{\mathcal{T}})$ is a HST metric space, and τ-HST metric space if the weights are exponentially decreasing, i.e., $w_u \leq w_v/\tau$, with v being the parent of u.

Similarly to [17], we use the algorithm from [21] (which we name via the authors' surnames as FRT) to approximate the given metric space $(\mathcal{X}, d)$ by a τ-HST one. In particular, we use FRT in Algorithm 1 as a computationally efficient preprocessing step to create a tree $\mathcal{T}$ with leaves $\mathcal{L}$ corresponding to actions in $\mathcal{X}$, distance metric $d_{\mathcal{T}}$, and root node r. We explain the intrinsic MTS motivation for this preprocessing step in Appendix B.2, and defer additional details to Appendix B.3.

3 The GP-MD Algorithm

In this section, we introduce GP-MD, a novel algorithm for the contextual BO problem with movement costs defined in Section 2. At each episode m and round h, the state of GP-MD can be summarized by a vector of probabilities $z_{h,m} \in K_{\mathcal{T}}$ over the vertices of $\mathcal{T}$, where $K_{\mathcal{T}} := \Big\{ z \in \mathbb{R}_+^{|V|} : z_r = 1, z_u = \sum_{\nu \in \mathcal{C}(u)} z_\nu \quad \forall u \in V \backslash \mathcal{L} \Big\}$, and $\mathcal{C}(u)$ denotes the children of u. Each entry z_ν represents the probability that the selected action $x_{h,m}$ belongs to the leaves of the subtree rooted at ν, i.e., $z_\nu = \mathbb{P}(x_{h,m} \in \mathcal{L}(\nu))$. Below, we specify how $z_{h,m}$ is computed at each round. Moreover, given any $z \in K_{\mathcal{T}}$, the vector

$$l(z) := [z_l,\ l \in \mathcal{L}] \in [0,1]^n, \tag{11}$$

defines a probability distribution over the leaves $\mathcal{L}$, and hence the actions $\mathcal{X}$. As in [17], given probability vectors $z_{h,m}$ and $z_{h-1,m}$, GP-MD computes the *minimal distance* distribution

$$\zeta_{h-1,h,m} = \operatorname*{arg\,inf}_{\zeta \in \Pi(l(z_{h-1,m}), l(z_{h,m}))} \mathbb{E}_\zeta[d_{\mathcal{T}}(U_{h-1,m}, U_{h,m})], \tag{12}$$

where $U_{h-1,m}$ and U_h are random variables having marginals $l(z_{h-1,m})$ and $l(z_{h,m})$ respectively. Finally, action $x_{h,m}$ is sampled from the conditional minimal distance distribution $x_{h,m} \sim$

$\zeta_{h-1,h,m}(\cdot|x_{h-1,m})$ (Line 13 in Algorithm 1). At the end of each episode m, the newly observed data are then used to update posterior mean and standard deviation about the cost function.

Finally, we describe how probability vectors $z_{h,m}$ are computed at each round, a key step of GP-MD (Lines 8–12 in Algorithm 1). We follow the recursive Mirror Descent (MD) procedure proposed by [17], with the important difference that we are dealing with an *unknown* context-dependent cost function. Hence, we make use of the Gaussian process model and corresponding confidence estimates as defined in Section 2.1.

To obtain probabilities $z_{h,m}$, we consider *conditional* probability vectors $q \in Q_{\mathcal{T}}$, where $Q_{\mathcal{T}}$ is the set of valid conditional probabilities $Q_{\mathcal{T}} := \Big\{ q \in \mathcal{R}_+^{|V \setminus r|} : \sum_{\nu \in \mathcal{C}(u)} q_\nu = 1 \quad \forall u \in V \setminus \mathcal{L} \Big\}$. For each vertex ν with parent u, q_ν represents the conditional probability $\mathbb{P}(x_{h,m} \in \mathcal{L}(\nu) | x_{h,m} \in \mathcal{L}(u))$. Moreover, given $q_h \in Q_{\mathcal{T}}$ we define the vector $q_h^{(u)} := [q_{h,\nu},\ \nu \in \mathcal{C}(u)]$ as the conditional distribution over children of u, and let $Q_{\mathcal{T}}^{(u)}$ be the set of all valid distributions $q_h^{(u)}$.

In each episode m, conditional probability vector q_h for round h is obtained recursively, from leaves to root, as a function of q_{h-1}, the observed context $e_{h,m}$, and the current estimate about the cost associated to each particular vertex. More precisely, let $\mathcal{OD}(V \setminus \mathcal{L})$ be a topological ordering of the internal vertices $V \setminus \mathcal{L}$ so that every child in $\mathcal{T}$ occurs before its parent. Then, for each $u \in \mathcal{OD}(V \setminus \mathcal{L})$ conditional probabilities $q_h^{(u)}$ are obtained via the Mirror Descent update:

$$q_h^{(u)} = \arg\min_{p \in Q_{\mathcal{T}}^{(u)}} \Big\{ D^{(u)}(p \| q_{h-1}^{(u)}) + \langle p, \mathrm{lcb}_m^{(u)}(\cdot, e_{h,m}) \rangle \Big\}. \tag{13}$$

Function $D^{(u)}$ is the Bregman divergence with respect to a suitable potential function (see Appendix D.1), while $\mathrm{lcb}_m^{(u)}(\cdot, e_{h,m}) := [\mathrm{lcb}_m(\nu, e_{h,m}),\ \forall \nu \in \mathcal{C}(u)]$ is a lower confidence bound estimate of the costs corresponding to children of vertex u. For $v \in \mathcal{L}$, $\mathrm{lcb}_m(\nu, e_{h,m})$ are obtained by the GP-regression techniques outlined in Section 2.1, while for internal vertices these are computed recursively from their children nodes as:

$$\mathrm{lcb}_m(u, e_{h,m}) := \sum_{\nu \in \mathcal{C}(u)} q_{h,\nu} \mathrm{lcb}_m(\nu, e_{h,m}) \,. \tag{14}$$

The movement cost is primarily controlled by this usage of Bregman divergence based mirror descent. Also, sampling from the conditional minimal distance distribution further restricts movement between alternate actions. Once the vector of conditional probabilities $q_h \in Q_{\mathcal{T}}$ has been updated, we can obtain the corresponding probability vector $z_{h,m}$ via the mapping $\Delta : Q_{\mathcal{T}} \to K_{\mathcal{T}}$ such that:

$$z = \Delta(q) \quad \Rightarrow \quad z_\nu = z_u q_\nu \quad \forall u \in V \setminus \mathcal{L},\ \nu \in \mathcal{C}(u). \tag{15}$$

3.1 Theoretical Guarantees

Our main theorem bounds the cumulative regret of GP-MD.

Theorem 1. *Let $\mathcal{X}$ be represented by a τ-HST space with $\tau > 4$ (Line 2 of Algorithm 1), and set $\delta \in (0,1)$. Then, with probability at least $1-\delta$, the regret of* GP-MD *over N_{ep} episodes is bounded by*

$$R_{N_{ep}}^{\alpha,\beta} = \mathcal{O}\Big(\beta_{N_{ep}} \big(N_{ep} H \exp(\gamma_H) \gamma_{H N_{ep}} + H \log(\tfrac{H}{\delta}) \big)^{\frac{1}{2}} + H(B+\psi) \log\big(\tfrac{N_{ep} \log(N_{ep})}{\delta} \big) \Big),$$

with approximation factors $\alpha = \mathcal{O}\big((\log n)^2\big)$ and $\beta = \mathcal{O}(1)$. Here, H is the episodes' length, $\beta_{N_{ep}}$ is the confidence level from Lemma 1, and $\gamma_{H N_{ep}}$ is the maximum information gain defined in Eq. (7).

For most of the popularly used kernels, Thm. 1 can be made more explicit by substituting bounds on $\gamma_{H N_{ep}}$ (e.g., in the case of linear kernel and compact domain, we have $\gamma_t = O(p \log t)$, while for squared-exponential kernel it holds $\gamma_t = O((\log t)^{p+1})$ [47]; see also Section 2.1). In such cases, we make the following two important observations regarding our result: i) The obtained regret bound is sublinear in the number of episodes N_{ep} and hence $\lim_{N_{ep} \to \infty} R_{N_{ep}}^{\alpha,\beta} / N_{ep} = 0$; ii) The bound is independent of the input space size, i.e., the number of actions n (although the approximation factor α depends logarithmically on n, similarly to [17]). These imply that GP-MD approaches α-competitive ratio performance of the MTS algorithm by [17], while learning about the service

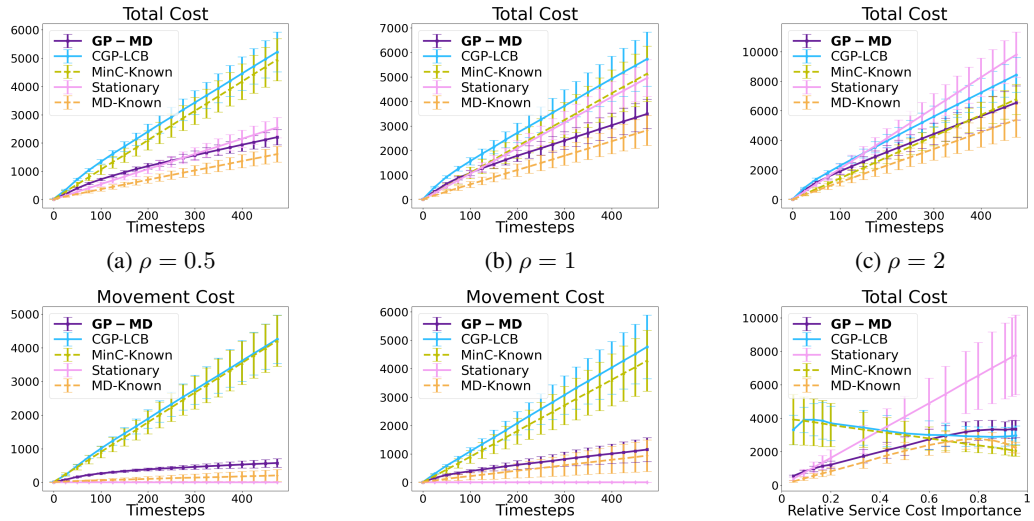

(a) $\rho = 0.5$ (b) $\rho = 1$ (c) $\rho = 2$

(d) $\rho = 0.5$ (e) $\rho = 2$ (f) Varying ρ

Figure 2: Total and movement cost performance of algorithms on synthetic functions for varying importance of movement/service cost (i.e., different ρ values). GP-MD outperforms CGP-LCB in terms of total incurred cost, and its performance closely follows one of the idealized benchmark MD-KNOWN. GP-MD also minimizes the movement cost while CGP-LCB suffers from significant movements (Figs. 2d and 2e). The performance of GP-MD remains robust when the movement cost importance in the total cost objective diminishes (Fig. 2f).

cost from noisy point evaluations (i.e., bandit feedback) only. Finally, in our analysis, we treat H as constant and note that a similar dependence on H is also present in the batch Bayesian optimization results and analysis that we also build upon (see [18, Theorem 2]). Improving this dependency in our bound is an interesting direction for future work.

Proof of Thm. 1 is detailed in Appendix E. Next, we outline some main steps. We make use of the competitive ratio guarantees for the used Mirror Descent algorithm from [17, Corollary 4] to bound the expected hallucinated service cost $\sum_{h=1}^{H} \langle \mathrm{lcb}_m(\cdot, e_{h,m}), l(z_{h,m}) \rangle$. Here, we use the fact that $\mathrm{lcb}_m(\cdot,\cdot)$ does not change within an episode. Since this is not the actual service cost, we bound $\sum_{h=1}^{H} \langle f(\cdot, e_{h,m}), l(z_{h,m}) \rangle$ by $\sum_{h=1}^{H} \langle \mathrm{lcb}_m(\cdot, e_{h,m}), l(z_{h,m}) \rangle$ with an additional learning error. The sum of the learning errors over all episodes can be rewritten as $\sum_{m=1}^{N_{ep}} \sum_{h=1}^{H} \langle \sigma_m^2(\cdot, e_{h,m}), l(z_{h,m}) \rangle$. We use the concentration of the conditional mean result from [27, Lemma 3] to upper bound it by the actual realizations $\sum_{m=1}^{N_{ep}} \sum_{h=1}^{H} \sigma_m^2(x_{h,m}, e_{h,m})$, and use the result of [18, Proposition 1], to further upper bound it with the maximum information gain quantity $\gamma_{N_{ep}H}$.

Finally, the movement cost can also be bounded similar to that of the previously mentioned expected hallucinated service cost $\sum_{h=1}^{H} \langle \mathrm{lcb}_m(\cdot, e_{h,m}), l(z_{h,m}) \rangle$ with an extra $\alpha = \mathcal{O}\big((\log n)^2\big)$ factor using [17, Corollary 4].

4 Experiments

This section provides numerical results on synthetic and real-world data. We compare the performance of our GP-MD algorithm with the following baselines:

- STATIONARY selects the stationary strategy $x_h = x_0$ for all h,
- CGP-LCB [29] neglects the movement cost and sets $x_h = \arg\min_x \mathrm{lcb}_h(x, e_h)$ for all h,
- MINC-KNOWN assumes $f(\cdot)$ is known and chooses $x_h = \arg\min_x f(x, e_h)$, and
- MD-KNOWN assumes $f(\cdot)$ is known and runs mirror descent from [17] on $f(\cdot, e_h)$.

MD-KNOWN and MINC-KNOWN unrealistically assume that $f(\cdot)$ is known and can be seen as upper-bound for the achievable performance of GP-MD and CGP-LCB, respectively. We use the same constant value $\beta = 2.0$ for the exploration parameter in both GP-MD and CGP-LCB (since the theoretical worst-case bounds are found to be overly pessimistic and can impede performance [47]). We run the algorithms over a single episode (as done in CGP-LCB). We also discover that updating the confidence bounds after every timestep in GP-MD leads to better performance in practice.

Synthetic experiments. We consider synthetic experiments, where the objective function is a random GP sample. The considered action space $\mathcal{X}$ is a subset of $[0,1]^2$ consisting of 400 points that form the uniform grid, while the context space $\mathcal{E}$ consists of 40 contexts that are uniformly sampled from $(0,1)$. We sample objective function (i.e., actual cost) $f : \mathcal{X} \times \mathcal{E} \to \mathbb{R}$ from a $GP(0,k)$, where k is a squared exponential kernel with lengthscale parameter set to $l = 0.2$. We use the Euclidean distance between the domain points as the movement cost, calculate the distance matrix (between every pair of points) and the average movement cost. We subtract the minimum value from f and scale it such that the average function value is equal to the average movement cost. We also set the noise parameter to 1% of the function range. We introduce the trade-off parameter $\rho \in \{0.25, 0.5, 1, 2, 4\}$ that only multiplies the service cost, i.e, $\rho f(x,e)$, but not the movement cost. This is to test the performance of the algorithms for varying importance of the service/movement costs. For each ρ we sample 25 different functions and run the algorithms for 500 timesteps wherein at each step the contexts are randomly sampled.

In Fig. 2a-Fig. 2e, we show the total cumulative cost as a function of timesteps for different ρ values. Then, in Fig. 2f, we show the performance of the algorithms (for known kernel parameters) when run for 800 timesteps for varying importance of the service and movement costs. In particular, we consider a convex combination of the service and movement costs, where we set the respective coefficients multiplying these two objectives as $\rho/(1+\rho)$ and $1/(1+\rho)$.

As shown in Fig. 2a-Fig. 2c, the performance of GP-MD is generally close to the one of the idealized, unrealistic benchmark MD-KNOWN, which, as expected, performs the best. The stationary baseline performs comparably when ρ is small, while its performance deteriorates for larger values. As expected, both MINC-KNOWN and CGP-LCB incur higher total costs than GP-MD when the movement cost is of the higher or same relative importance as the service cost (i.e., $\rho \in \{0.5, 1.0\}$), while the performance gap slowly decreases when the service cost becomes dominant ($\rho = 2.0$). In Figs. 2d and 2e, we also show the corresponding movement costs, and observe that movement cost ignorant CGP-LCB incurs significant movement costs, while our GP-MD successfully minimizes the movement costs. Finally, in Fig. 2f, we observe that the performance of GP-MD is robust, i.e., it clearly outperforms CGP-LCB whenever the movement cost dominates the total cost objective, while its performance remains comparable to the one of CGP-LCB (that is built to maximize service cost) when the movement cost becomes dominated by the service cost.

4.1 Altitude Optimization in AWE Systems

In airborne wind energy (AWE) systems, the turbine's operating altitude can be changed depending on the wind pattern. We follow the setup of Baheri et al. [7] that applied CGP-LCB [29] which ignores movement-costs, to learn this control task. In this section, we use a dataset from [9] which contains wind-speed information over various locations in Europe for a period ranging from 2011 to 2017, and also includes measurements at different altitudes per location. We consider the wind speed data from the second half of 2016. Our goal is to maximize the generated energy, while taking into account the energy loss due to moving the turbine from one altitude to another.

We consider 25 different altitudes (ranging from 10m to 1600m) as the action space and the context space to be hours in the day (i.e., 24 values). We define our unknown service objective function to be $f(x,t) = \max_{x'}(E_S(x',t)) - E_S(x,t)$ where $E_S(x,t)$ denotes the energy generated based on the windspeed at altitude x and time t. Based on the discrete-time power generation formula from Baheri et al. [7] (Eq. (10)), we have

$$E_S(x,t) = \big(c_1(\min\{V_w(x,t), V_r\})^3 - c_2 V_w^2(x,t)\big)\Delta t. \tag{16}$$

Here $V_w(x,t)$ denotes the windspeed at altitude x and time t, and V_r denotes the rated windspeed of the turbine. The constants $c_1 = 0.0579$ and $c_2 = 0.09$ are system dependent. This corresponds to the energy generated at a particular altitude x for Δt time. Similarly to Baheri et al. [7], we use $\Delta t = 60$ since we consider intervals of one hour length. Next, we define the movement cost to be the energy lost in changing altitude (from x to x'):

$$E_M(x,x') = c_3 V_r^2 |x - x'|, \tag{17}$$

where $c_3 = 0.15$ (see Appendix F for more details).[2]

[2] According to the power equation from [7], $E_M(x,x')$ would depend on $V_w(x,t)$, whereas, we assume our movement cost is based on a fixed metric and is independent of contexts. Hence, we simply approximate $V_w(x,t)$ by V_r and consider time-independent movement costs.

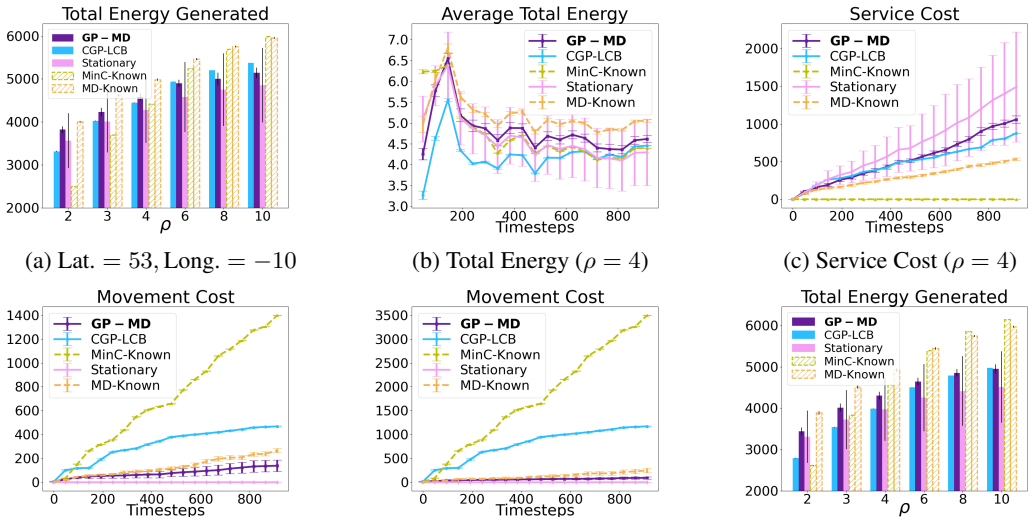

(a) Lat. $= 53$, Long. $= -10$ (b) Total Energy ($\rho = 4$) (c) Service Cost ($\rho = 4$)

(d) Movement Cost ($\rho = 4$) (e) Movement Cost ($\rho = 1$) (f) Lat. $= 53$, Long. $= -4$

Figure 3: AWE altitude optimization task; Fig. 3a: Total energy generated for 960 hours based on the wind data at a single location (Latitude$= 53$ and Longitude$= -10$). GP-MD outperforms previously used CGP-LCB (that optimizes for service costs only) for a range of ρ values that favor the service against the movement cost. Fig. 3b: The average total generated energy over 960 hours. Figs. 3c and 3d: The service and movement costs for $\rho = 4$. Fig. 3e: The movement costs for $\rho = 1$. The movement energy loss is slightly lower for GP-MD as compared to $\rho = 4$ due to higher importance towards movement cost reduction. Fig. 3f: Same as Fig. 3a, albeit by using wind data from a different location (Latitude$= 53$, Longitude$= -4$).

We assume that a wind speed gauge is attached to the turbine, and the operator knows the wind speed at the current altitude. Hence, instead of directly learning $f(x, t)$, we learn $V_w(x, t)$ and use its confidence bounds to calculate the confidence bounds of $f(x, t)$. To learn $V_w(x, t)$, we normalize the inputs, and fit a GP with RBF kernel (lengthscale=3.67, outputscale=6.85 and noise parameter=2.73).

We run the algorithms for different ρ for 960 timesteps, where again ρ is used to multiply E_S. For each ρ, the algorithms were initiated with every possible starting point (25 different altitudes), and ran for 3 iterations. Based on this we plot the total energy generated w.r.t. varying ρ in Fig. 3. In Figs. 3a and 3f, we show the performance of the algorithms at two different locations (we also consider additional locations and time periods in Appendix F). We use different values of $\rho > 1$ to show the robustness of our algorithm (as ρ increases, the importance of the service cost w.r.t. the movement cost in the overall objective increases). Our algorithm outperforms CGP-LCB for a range of ρ values. As ρ keeps increasing, we observe that MINC-KNOWN closes the performance gap to MD-KNOWN, and the same is happening with CGP-LCB w.r.t. GP-MD. In Fig. 3b, we focus on a particular $\rho = 4$, and notice that GP-MD performs better than CGP-LCB and STATIONARY algorithm at this location. In Fig. 3c, we plot the service cost and observe that both learning algorithms GP-MD and CGP-LCB have lower service cost than the STATIONARY baseline. We also note that due to the implicit service cost definition, the MINC-KNOWN baseline achieves zero service cost. In Figs. 3d and 3e, we compare the movement energy loss for $\rho = 4$ and $\rho = 1$. As expected, $\rho = 1$ results in slightly lower GP-MD movement energy loss due to the tradeoff shifting towards the movement cost.

5 Conclusions

We have considered the problem of optimizing an unknown cost function subject to time-varying contextual information, as well as *movement costs* of changing the selected action from round to round. Our problem formulation is motivated by Airborne Wind Energy systems, where one seeks to optimize the operating altitude of the wind turbine to maximize the amount of generated energy. We propose a novel algorithm, GP-MD, which makes use of GP confidence bounds and employs the mirror descent techniques from [17] for solving MTS problems. We analyze the theoretical performance of our algorithm by providing a rigorous regret bound. Moreover, we demonstrate its performance in synthetic experiments and on an AWE application by using real-world data. GP-MD carefully trades off service and movement costs while at the same time learning about the unknown objective function and yielding improved performance (i.e., generating more energy) compared to the considered baselines.

References

[1] Abbasi-Yadkori. Online learning for linearly parametrized control problems. 2013.

[2] Majid Abdolshah, Alistair Shilton, Santu Rana, Sunil Gupta, and Svetha Venkatesh. Cost-aware multi-objective Bayesian optimisation. *arXiv preprint arXiv:1909.03600*, 2019.

[3] Abdulelah Alkesaiberi, Fouzi Harrou, and Ying Sun. Efficient wind power prediction using machine learning methods: A comparative study. *Energies*, 2022.

[4] Leif Erik Andersson and Lars Imsland. Real-time optimization of wind farms using modifier adaptation and machine learning. *Wind Energy Science*, 2020.

[5] CJ Argue, Anupam Gupta, Ziye Tang, and Guru Guruganesh. Chasing convex bodies with linear competitive ratio. *Journal of the ACM (JACM)*, 2021.

[6] Ali Baheri and Chris Vermillion. Combined plant and controller design using batch bayesian optimization: a case study in airborne wind energy systems. *Journal of Dynamic Systems, Measurement, and Control*, 2019.

[7] Ali Baheri, Shamir Bin-Karim, Alireza Bafandeh, and Christopher Vermillion. Real-time control using Bayesian optimization: A case study in airborne wind energy systems. *Control Engineering Practice*, 2017.

[8] Yair Bartal. Probabilistic approximation of metric spaces and its algorithmic applications. *Proceedings of 37th Conference on Foundations of Computer Science*, 1996.

[9] Philip Bechtle, Mark Schelbergen, Roland Schmehl, Udo Zillmann, and Simon Watson. Airborne wind energy resource analysis. *Renewable energy*, 2019.

[10] Ilija Bogunovic, Jonathan Scarlett, Andreas Krause, and Volkan Cevher. Truncated variance reduction: A unified approach to Bayesian optimization and level-set estimation. *Conference on Neural Information Processing Systems (NeurIPS)*, 2016.

[11] Allan Borodin, Nathan Linial, and Michael E Saks. An optimal on-line algorithm for metrical task system. *Journal of the ACM (JACM)*, 1992.

[12] Sébastien Bubeck, Yin Tat Lee, Yuanzhi Li, and Mark Sellke. Competitively chasing convex bodies. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, 2019.

[13] Sébastien Bubeck, Bo'az Klartag, Yin Tat Lee, Yuanzhi Li, and Mark Sellke. Chasing nested convex bodies nearly optimally. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2020.

[14] Sébastien Bubeck, Michael B Cohen, James R Lee, and Yin Tat Lee. Metrical task systems on trees via mirror descent and unfair gluing. *SIAM Journal on Computing*, 2021.

[15] Ian Char, Youngseog Chung, Willie Neiswanger, Kirthevasan Kandasamy, Andrew O Nelson, Mark Boyer, Egemen Kolemen, and Jeff Schneider. Offline contextual Bayesian optimization. *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.

[16] Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. *International Conference on Machine Learning (ICML)*, 2017.

[17] Christian Coester and James R Lee. Pure entropic regularization for metrical task systems. *Conference on Learning Theory (COLT)*, 2019.

[18] Thomas Desautels, Andreas Krause, and Joel W Burdick. Parallelizing exploration-exploitation tradeoffs in Gaussian process bandit optimization. *Journal of Machine Learning Research (JMLR)*, 2014.

[19] Bart M Doekemeijer, Daan C Van Der Hoek, and Jan-Willem van Wingerden. Model-based closed-loop wind farm control for power maximization using bayesian optimization: a large eddy simulation study. In *2019 IEEE Conference on Control Technology and Applications (CCTA)*. IEEE, 2019.

[20] Dave Elliot. Flights of fancy: airborne wind turbines, 2014.

[21] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *Journal of Computer and System Sciences*, 2004.

[22] Gautam Goel and Adam Wierman. An online algorithm for smoothed regression and lqr control. In *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR, 2019.

[23] Gautam Goel, Yiheng Lin, Haoyuan Sun, and Adam Wierman. Beyond online balanced descent: An optimal algorithm for smoothed online optimization. *Advances in Neural Information Processing Systems*, 2019.

[24] Javier González, Zhenwen Dai, Philipp Hennig, and Neil Lawrence. Batch Bayesian optimization via local penalization. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2016.

[25] Kirthevasan Kandasamy, Gautam Dasarathy, Jeff Schneider, and Barnabás Póczos. Multi-fidelity Bayesian optimisation with continuous approximations. *International Conference on Machine Learning (ICML)*, 2017.

[26] Tarun Kathuria, Amit Deshpande, and Pushmeet Kohli. Batched Gaussian process bandit optimization via determinantal point processes. *Conference on Neural Information Processing Systems (NeurIPS)*, 2016.

[27] Johannes Kirschner and Andreas Krause. Information directed sampling and bandits with heteroscedastic noise. *Conference On Learning Theory (COLT)*, 2018.

[28] Johannes Kirschner, Ilija Bogunovic, Stefanie Jegelka, and Andreas Krause. Distributionally robust Bayesian optimization. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.

[29] Andreas Krause and Cheng Soon Ong. Contextual Gaussian process bandit optimization. *Conference on Neural Information Processing Systems (NeurIPS)*, 2011.

[30] Remi Lam and Karen Willcox. Lookahead Bayesian optimization with inequality constraints. *Conference on Neural Information Processing Systems (NeurIPS)*, 2017.

[31] Remi Lam, Karen Willcox, and David H Wolpert. Bayesian optimization with a finite budget: An approximate dynamic programming approach. *Conference on Neural Information Processing Systems (NeurIPS)*, 2016.

[32] Chi-Heng Lin, Joseph D Miano, and Eva L Dyer. Bayesian optimization for modular black-box systems with switching costs. *Uncertainty in Artificial Intelligence (UAI)*, 2021.

[33] Minghong Lin, Adam Wierman, Lachlan LH Andrew, and Eno Thereska. Dynamic right-sizing for power-proportional data centers. *IEEE/ACM Transactions on Networking*, 2012.

[34] Yiheng Lin, Gautam Goel, and Adam Wierman. Online optimization with predictions and non-convex losses. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 2020.

[35] Jonas Močkus. On Bayesian methods for seeking the extremum. In *Optimization Techniques*, 1975.

[36] Nikolaos Moustakis, Sebastiaan Paul Mulders, Jens Kober, and Jan-Willem van Wingerden. A practical bayesian optimization approach for the optimal estimation of the rotor effective wind speed. In *2019 American Control Conference (ACC)*. IEEE, 2019.

[37] Jinkyoo Park. Contextual Bayesian optimization with trust region (cbotr) and its application to cooperative wind farm control in region 2. *Sustainable Energy Technologies and Assessments*, 2020.

[38] Jinkyoo Park and Kincho H Law. A bayesian optimization approach for wind farm power maximization. In *Smart Sensor Phenomena, Technology, Networks, and Systems Integration 2015*. International Society for Optics and Photonics, 2015.

[39] C. E. Rasmussen and C. K. Williams. Gaussian processes for machine learning. *volume 1. MIT press Cambridge*, 2006.

[40] Mark Sellke. Chasing convex bodies optimally. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2020.

[41] Pier Giuseppe Sessa, Ilija Bogunovic, Maryam Kamgarpour, and Andreas Krause. No-regret learning in unknown games with correlated payoffs. *Conference on Neural Information Processing Systems (NeurIPS)*, December 2019.

[42] Pier Giuseppe Sessa, Ilija Bogunovic, Maryam Kamgarpour, and Andreas Krause. Mixed strategies for robust optimization of unknown objectives. In *Proc. International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.

[43] Pier Giuseppe Sessa, Ilija Bogunovic, Maryam Kamgarpour, and Andreas Krause. Learning to play sequential games versus unknown opponents. *Conference on Neural Information Processing Systems (NeurIPS)*, December 2020.

[44] Pier Giuseppe Sessa, Ilija Bogunovic, Andreas Krause, and Maryam Kamgarpour. Contextual games: Multi-agent learning with side information. *Conference on Neural Information Processing Systems (NeurIPS)*, December 2020.

[45] Guanya Shi, Yiheng Lin, Soon-Jo Chung, Yisong Yue, and Adam Wierman. Online optimization with memory and competitive control. *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.

[46] Jialin Song, Yuxin Chen, and Yisong Yue. A general framework for multi-fidelity Bayesian optimization with Gaussian processes. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.

[47] N. Srinivas, A. Krause, S. M. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *International Conference on Machine Learning (ICML)*, 2010.

[48] Kübra Yazici and Semra Boran. Short-term wind power prediction approach based on bayesian optimization and ensemble learning. *Journal of Intelligent Systems: Theory and Applications*.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes] See Section **??**.
- Did you include the license to the code and datasets? [No] The code and the data are proprietary.
- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]

2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Added in supplementary material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] Internal Cluster was used
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]