TIME-VARYING BAYESIAN OPTIMIZATION WITHOUT A METRONOME

Anonymous authors
Paper under double-blind review

ABSTRACT

Time-Varying Bayesian Optimization (TVBO) is the go-to framework for optimizing time-varying, expensive, noisy black-box functions. However, most asymptotic guarantees for TVBO algorithms assume observations are acquired at a constant frequency. Since streaming GP inference scales quadratically with dataset size, this assumption is unrealistic in the long run. We relax it and derive the first upper regret bound that explicitly accounts for changes in observation sampling frequency. This analysis yields practical recommendations for dataset sizes and stale-data policies. Experiments on synthetic and real-world problems show that an algorithm (BOLT) following these recommendations outperforms state-of-the-art TVBO methods.

1 INTRODUCTION

Bayesian Optimization (BO) uses a surrogate model of the objective function f, typically a Gaussian Process (GP), to efficiently learn and optimize it. Its sample efficiency has made BO the standard approach to black-box optimization across diverse applications (Marchant & Ramos, 2012; Bardou et al., 2025; Wang et al., 2014) and motivated extensions to challenging settings such as multi-objective optimization (Daulton et al., 2022) and high-dimensional input spaces (Bardou et al., 2024a). Yet only a handful of works study Time-Varying Bayesian Optimization (TVBO), where $f : \mathcal{S} \times \mathcal{T} \to \mathbb{R}$ also depends on time $t \in \mathcal{T} \subseteq \mathbb{R}$, despite the ubiquity of time-varying black-box optimization problems, with occurrences in unmanned aerial vehicles (Melo et al., 2021), online clustering (Aggarwal et al., 2004), and network management (Kim et al., 2019).

TVBO differs substantially from its static counterpart. In particular, recent work identifies the *response time* $R(n)$ (i.e., the time between consecutive iterations when the dataset has size n) as a key algorithmic quantity. Since $R(n)$ typically scales as $\mathcal{O}(n^2)$,[1] balancing a large dataset against a short response time can yield significant performance gains (Bardou et al., 2024b). However, to the best of our knowledge, no analysis links $R(n)$ to TVBO performance without assuming $R(n) \in \Theta(1)$. This paper fills this gap by:

- deriving the first asymptotic regret bound that explicitly relates TVBO performance to the response time $R(n)$ (Section 3.2);
- translating this analysis into practical recommendations and embedding them in a new algorithm, BOLT (Sections 3.2, 3.3, and 4);
- providing empirical evidence that BOLT outperforms state-of-the-art TVBO methods on synthetic and real-world problems (Section 5).

2 BACKGROUND

2.1 TIME-VARYING BAYESIAN OPTIMIZATION

Like most BO algorithms, a TVBO algorithm uses a centered GP with covariance function k (denoted by $\mathcal{GP}(0, k)$) as prior placed on the objective function f with observational noise σ_0^2. Conditioned on

[1] Inverting the covariance matrix $k(\mathcal{D}, \mathcal{D})$ requires $\mathcal{O}(n^3)$ operations, but rank-one updates allow the precision matrix to be updated in $\mathcal{O}(n^2)$ in the streaming setting.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

a dataset of n observations $\mathcal{D} = \{(\boldsymbol{x}_i, t_i, y_i)\}_{i\in[n]}$, where $y_i = f(\boldsymbol{x}_i, t_i) + \epsilon, \epsilon \sim \mathcal{N}(0, \sigma_0^2)$ for any $i \in [n]$, the posterior satisfies $f(\boldsymbol{x}, t)|\mathcal{D} \sim \mathcal{N}\left(\mu_{\mathcal{D}}(\boldsymbol{x}, t), \sigma^2_{\mathcal{D}}(\boldsymbol{x}, t)\right)$ for any $(\boldsymbol{x}, t) \in \mathcal{S} \times \mathcal{T}$, where

$$\mu_{\mathcal{D}}(\boldsymbol{x}, t) = \boldsymbol{k}((\boldsymbol{x}, t), \mathcal{D})\boldsymbol{\Delta}^{-1}\boldsymbol{y}, \tag{1}$$

$$\sigma^2_{\mathcal{D}}(\boldsymbol{x}, t) = k((\boldsymbol{x}, t), (\boldsymbol{x}, t)) - \boldsymbol{k}^\top ((\boldsymbol{x}, t), \mathcal{D})\, \boldsymbol{\Delta}^{-1}\boldsymbol{k}((\boldsymbol{x}, t), \mathcal{D}), \tag{2}$$

where $\boldsymbol{\Delta} = \boldsymbol{K} + \sigma_0^2\boldsymbol{I}_n$, $\boldsymbol{K} = \boldsymbol{k}(\mathcal{D}, \mathcal{D})$, $\boldsymbol{k}(\mathcal{X}, \mathcal{Y}) = (k((\boldsymbol{x}_i, t_i), (\boldsymbol{x}_j, t_j)))_{(\boldsymbol{x}_i,t_i)\in\mathcal{X},(\boldsymbol{x}_j,t_j)\in\mathcal{Y}}$, $\boldsymbol{y} = (y_i)_{y_i\in\mathcal{D}}$ and where $\boldsymbol{I}_n$ is the $n \times n$ identity matrix.

At the i-th iteration occurring at time t_i, a TVBO algorithm looks for a query $(\boldsymbol{x}_i, t_i)$ that achieves the best exploration-exploitation trade-off. In the BO framework, this problem is addressed by maximizing an acquisition function $\alpha_{\mathcal{D}} : \mathcal{S} \times \mathcal{T} \to \mathbb{R}$, so that $\boldsymbol{x}_i = \arg\max_{\boldsymbol{x}\in\mathcal{S}} \alpha_{\mathcal{D}}(\boldsymbol{x}, t_i)$. Many acquisition functions have been proposed, the most popular being GPUCB (Srinivas et al., 2012), Expected Improvement (Mockus, 1994) and Probability of Improvement (Jones et al., 1998).

At the i-th iteration, the instantaneous performance of a TVBO algorithm is measured by the instantaneous regret $r_i = f(\boldsymbol{x}_i^*, t_i) - f(\boldsymbol{x}_i, t_i)$, where $\boldsymbol{x}_i^* = \arg\max_{\boldsymbol{x}\in\mathcal{S}} f(\boldsymbol{x}, t_i)$. The performance up to an horizon T is usually measured by the cumulative regret $R_T = \sum_{i=1}^{T} r_i$.

2.2 The Three Time Scales of a TVBO Problem

A TVBO problem involves three interacting time scales (which are not given to the algorithm): (i) a temporal lengthscale l_T that characterizes the rate of variation of f, (ii) the time horizon H of the optimization task, and (iii) the response time $R(n)$ of the algorithm, which includes the time needed to observe f and to make a new query using a dataset of size n. In the following, we denote by $N_H = \lfloor H/R(0) \rfloor$ the maximal number of observations that could be collected if the computational overhead of GP inference were negligible.

Because $R(n)$ grows with n, every TVBO algorithm must implicitly address a trade-off between the information it retains and the frequency at which it can sample f. The literature resolves it in three ways: by retaining every observation (growing datasets), by discarding the dataset as a whole at chosen instants (resets), or by discarding selected observations while keeping the rest (removal). Which of these strategies is appropriate depends on the relative magnitudes of l_T, H and $R(n)$. We discuss three particular regimes here.

Near-Constant Response Time. Suppose that $R(N_H) \approx R(0)$. The computational overhead remains negligible even if the algorithm retains every observation that can be collected up to the horizon H. This regime typically occurs when objective evaluation is substantially more expensive than GP inference. Growing datasets are then computationally viable, and controlling the dataset size offers little benefit in terms of sampling frequency. Differences between algorithms are therefore driven primarily by the quality of statistical models rather than by their response times.

Quasi-Static Problems. Suppose that $l_T \gg R(n)$ for a reasonable dataset size $n \ll N_H$. The algorithm can collect many observations before the objective changes substantially. Resets may then remain competitive by treating the problem as a sequence of approximately static optimization tasks, although each reset may discard observations that are still informative. Growing datasets can also model the objective accurately, but their response times may eventually become prohibitive when the horizon is long.

General Long-Horizon Problems. The most challenging regime occurs when the horizon is long, the response time grows noticeably with the dataset size n, and l_T is comparable to $R(n)$ for practically useful values of n. Growing datasets may then issue observations at an increasingly low frequency, while resets may repeatedly discard the dataset before enough observations have been collected to build an informative surrogate. Removal is the only one among the three strategies that can, in principle, keep the response time bounded without sacrificing the observations that remain informative, but it raises two questions that the other two strategies answer trivially: how many observations to retain, and which ones to discard.

This paper targets this third regime, and Section 3 answers both questions in turn.

Table 1: Comparison of TVBO solutions. For each method, we report its principal model of temporal variation, how response time enters its theoretical analysis, its stale-data policy, and how its dataset size is controlled. "Not modeled" indicates that the method can be executed with a nonconstant response time, but does not explicitly account for it in its analysis.

TVBO Solution	Temporal Model or Assumption	Response Time in Analysis	Stale-Data Policy	Dataset-Size Rule
TV-GPUCB	Markov GP evolution	$R(n) \in \Theta(1)$	Growing	N/A
R-GPUCB	Markov GP evolution	$R(n) \in \Theta(1)$	Reset	Periodic
ABO	Spatio-temporal GP	Not modeled	Growing	N/A
SW-GPUCB	Finite variation budget	$R(n) \in \Theta(1)$	Sliding window	Fixed
W-GPUCB	Finite variation budget	$R(n) \in \Theta(1)$	Growing	N/A
ET-GPUCB	Markov GP evolution	$R(n) \in \Theta(1)$	Reset	Event-based
W-DBO	Spatio-temporal GP	Not modeled	Removal	Adaptive
BOLT	Spatio-temporal GP	Explicit $R(n)$	Removal	Adaptive

2.3 State-of-the-Art of TVBO

Table 1 summarizes the main TVBO solutions along the axes that the preceding discussion identifies as decisive: how the response time enters the analysis of each method, its stale-data policy, and how its dataset size is controlled.

Existing regret analyses have primarily considered two models of temporal variation. The first one assumes that the objective follows a persistent Markov evolution. In this setting, the expected cumulative regret is bounded from below linearly, and algorithms with upper bounds of the same asymptotic order, up to polylogarithmic factors, have been proposed (Bogunovic et al., 2016; Brunzema et al., 2025). The second one assumes that the cumulative temporal variation of the objective is sublinear. Sublinear regret can then be achieved when this variation grows sufficiently slowly with the horizon (Zhou & Shroff, 2021; Deng et al., 2022). Such condition excludes environments with a nonvanishing temporal variation.

We do not include these regret rates in Table 1, because they hold under fundamentally different models of nonstationarity. Moreover, in a bounded-reward setting, a linear upper bound alone does not distinguish between algorithms. The central distinction relevant to this work is instead whether an algorithm explicitly accounts for the dependence of its response time on its dataset size. Existing algorithms with regret guarantees assume $R(n) \in \Theta(1)$, which amounts to placing themselves in the near-constant regime, whereas methods without this assumption do not model the effect of $R(n)$ on their performance.

Mapping these methods onto the three regimes makes the gap visible. Growing datasets (TV-GPUCB, ABO, W-GPUCB) implicitly assume a near-constant response time, while resets (R-GPUCB, ET-GPUCB) and fixed-size sliding windows (SW-GPUCB) suit quasi-static problems. W-DBO is, to the best of our knowledge, the only method performing relevancy-based removal, but the size of its dataset follows from a user-specified removal budget rather than from an analysis of $R(n)$.

This paper closes that gap. Our analysis considers a GPUCB algorithm equipped with a FIFO policy, which isolates the trade-off between retaining more observations and sampling the evolving objective more frequently. It answers both questions raised above: how many observations to retain (Section 3.2) and which ones to discard (Section 3.3).

3 Main Results

3.1 Core Assumptions and Problem Setting

In this section, we introduce the core assumptions underpinning this work. Assumption 3.1 is made in all GP-based BO papers as it ensures the existence of the surrogate model. Assumption 3.2 is widely used in the TVBO literature (Bogunovic et al., 2016; Nyikosa et al., 2018; Bardou et al., 2024b).

It captures spatio-temporal dynamics with two correlation functions k_S and k_T. Assumption 3.3 requires that k_S and k_T verify some properties. It holds for most kernels used in practice (e.g., Matérn, RBF, Rational Quadratic, sinc) and the condition (3) is further discussed in Remark A.4. Assumption 3.4 introduces some regularity about the objective function in the spatial domain $\mathcal{S}$. It is used for most regret proofs in the literature (Srinivas et al., 2012; Bogunovic et al., 2016). This is satisfied by sufficiently smooth GPs (e.g., built with Matérn with $\nu > 2$, RBF) but it can fail with rougher GPs (e.g., Ornstein-Uhlenbeck processes built with Matérn kernels where $\nu = 1/2$).

Assumption 3.1 (Surrogate Model)**.** The objective function $f : \mathcal{S} \times \mathcal{T} \to \mathbb{R}$ is a sample path of $\mathcal{GP}(0, k)$, where $\mathcal{S} = [0, 1]^d$ and $k : (\mathcal{S} \times \mathcal{T})^2 \to \mathbb{R}$ is a covariance function.

Assumption 3.2 (Covariance Separability)**.** The covariance function k has the following form:

$$k((\boldsymbol{x}, t), (\boldsymbol{x}', t')) = \lambda k_S(\boldsymbol{x}, \boldsymbol{x}')k_T(t, t')$$

where, without loss of generality, $\lambda = 1$ is the signal variance while $k_S : \mathcal{S} \times \mathcal{S} \to [-1, 1]$ and $k_T : \mathcal{T} \times \mathcal{T} \to [-1, 1]$ are correlation functions in the spatial and temporal domain, respectively.

Assumption 3.3 (Covariance Properties)**.** The correlation function k_S (resp., k_T) is real, even, isotropic, stationary and can therefore be rewritten as $k_S(\boldsymbol{x}, \boldsymbol{x}') = k_S(\boldsymbol{x} - \boldsymbol{x}')$ for all $\boldsymbol{x}, \boldsymbol{x}' \in \mathcal{S}$ (resp., $k_T(t, t') = k_T(|t - t'|)$ for all $t, t' \in \mathcal{T}$). Furthermore, there exists a nonnegative, nonincreasing function g_T such that $|k_T(\tau)| \leq g_T(\tau)$ for all $\tau \geq 0$ and such that

$$G = \sup_{m \geq 1} \frac{1}{R(m)} \int_0^{mR(m)} g_T(\tau) d\tau < \infty, \tag{3}$$

where R is the response time of the TVBO algorithm (Definition 3.5). Finally, $\min_{\boldsymbol{x}, \boldsymbol{x}' \in \mathcal{S}} k_S(\boldsymbol{x}, \boldsymbol{x}') > 0$.

Assumption 3.4 (Lipschitz Continuity in $\mathcal{S}$)**.** Let $g \sim \mathcal{GP}(0, k_S)$. Then, for any $\boldsymbol{x} \in \mathcal{S}$, any $L > 0$ and any $i \in [d]$,

$$\mathbb{P}\left[\left|\frac{\partial g(\boldsymbol{x})}{\partial x_i}\right| > L\right] \leq a e^{-(L/b)^2}.$$

Finally, let us properly define the response time $R(n)$ of a TVBO algorithm, which models the combined costs of observing f and of updating/conditioning the GP on a dataset of size n. It satisfies three natural properties: (i) the objective function cannot be observed at an arbitrarily high frequency because querying it takes a nonzero amount of time, (ii) the response time is an increasing function of the dataset size n and (iii) when the dataset size n diverges, the response time diverges as well.

Definition 3.5 (Response Time)**.** The response time $R : \mathbb{N} \to \mathbb{R}^+$ is a function that returns the time separating two consecutive iterations of a TVBO algorithm with a dataset of size $n \in \mathbb{N}$ such that (i) $R(0) > 0$, (ii) $\forall n \in \mathbb{N}, R(n + 1) \geq R(n)$ and (iii) $\lim_{n \to \infty} R(n) = +\infty$.

Based on Definition 3.5, we study TVBO algorithms in a realistic, wall-clock–constrained regime where obtaining an informative posterior and issuing the next query are not instantaneous because both function observation and GP inference require non-negligible amounts of time. Concretely, the algorithm queries as soon as it is able to, so that the query times satisfy

$$t_0 = 0, \qquad t_m = \sum_{i=1}^{m} R(n_i) \quad (m \geq 1),$$

where n_i is the dataset size of the TVBO algorithm at iteration i.

With this greedy-in-time policy (Nyikosa et al., 2018; Bardou et al., 2024b), we consider the *average regret* R_m/m as the metric of performance to minimize after m iterations. This formulation is natural and well-posed: it makes explicit the fundamental trade-off in TVBO between (i) spending more compute per query to obtain a high-quality posterior (larger $R(n)$) and (ii) issuing more frequent queries to track temporal changes in f. Any theoretical or empirical comparison of algorithms in this setting must therefore account for the growth and variability of $R(n)$ and for the intrinsic time-scale of variation of $f(\cdot, t)$. Ignoring either one of these properties yields conclusions that may be irrelevant in wall-clock constrained, time-varying applications.

3.2 Regret Analysis Without a Metronome

We now analyze the regret of a TVBO algorithm under the assumptions and definitions introduced in Section 3.1. As the cumulative regret of a TVBO algorithm is linear in the worst case when $R(n) \in \Theta(1)$ (Bogunovic et al., 2016), one can expect a similar result when $R(n) \in \omega(1)$.

Understanding how the temporal correlation structure k_T and the response time $R(n)$ jointly affect the regret is essential for designing TVBO algorithms. Existing regret analyses that explicitly manage stale observations have mainly considered policies that reset the entire dataset, either periodically (Bogunovic et al., 2016) or after an event is triggered (Brunzema et al., 2025), while assuming a constant response time $R(n) \in \Theta(1)$. Such resets are relatively aggressive, as they discard all past observations at once, including recent observations that may remain informative. We consider instead a gentler First-In, First-Out (FIFO) policy: once the dataset reaches a maximal size n, each new observation replaces the oldest one, so that the dataset always contains the n most recent observations. This setting allows us to study the trade-off between the information retained from past observations, governed by k_T, and the sampling frequency, governed by $R(n)$. This trade-off is formalized by the asymptotic regret bound established in the following theorem.

Theorem 3.6. *Let $\mathcal{A}$ be a TVBO algorithm that uses the GPUCB acquisition function. Let $R(s)$ be the response time of $\mathcal{A}$ for a dataset size $s \in \mathbb{N}$. Let n be the maximal dataset size of $\mathcal{A}$ and let $\|\boldsymbol{u}_n\|_2^2 = \sum_{i=1}^{n} k_T^2(iR(n))$. Let $\delta \in (0,1)$ and let R_T be the cumulative regret of $\mathcal{A}$ after T iterations. Then, with probability $1-\delta$,*

$$R_T \leq \sqrt{TC_1\beta_T\left(\gamma_n + \frac{1}{2\sigma_0^2}(T-n)(1-C_2\|\boldsymbol{u}_n\|_2^2)\right)} + \frac{\pi^2}{6}, \tag{4}$$

where C_1, $C_2 \in \mathcal{O}(1)$, $\gamma_n \in \mathcal{O}(n)$ and $\beta_T \in \mathcal{O}(\log T)$ are defined in Appendix A.

The proof of Theorem 3.6 is provided in Appendix A. The gist of the proof is to first bound the cumulative regret of the TVBO algorithm from above with an expression that involves the mutual information $\gamma_n = I(\boldsymbol{f}_n, \boldsymbol{y}_n)$, where $\boldsymbol{f}_n = (f(\boldsymbol{x}_1, t_1), \cdots, f(\boldsymbol{x}_n, t_n))$ and $\boldsymbol{y}_n = (y_1, \cdots, y_n)$, and the posterior variance of the surrogate model with maximal dataset size n. This is achieved with proof techniques introduced by Srinivas et al. (2012); Bogunovic et al. (2016). Next, an upper bound is derived for the posterior variance that involves $\|\boldsymbol{u}_n\|_2^2$. Combining these two upper bounds yields (4).

Dependence of the constants and non-negativity. The constants C_1 and C_2 are independent of n and T, but they are not universal. In particular, they depend on the observational noise σ_0^2 and on the spatial and temporal kernels k_S and k_T. Furthermore, the quantity under the square root in Theorem 3.6 is shown in Appendix A to be always non-negative (see Lemma A.3).

Interpretation. For fixed maximal dataset size n and $T \to \infty$, γ_n is constant and β_T grows polylogarithmically, yielding $R_T \in \widetilde{O}(T)$. This does not establish sublinear regret under persistent temporal variation, consistently with the TVBO literature (Bogunovic et al., 2016; Brunzema et al., 2025). The bound instead reveals how the response time $R(n)$ and temporal kernel k_T jointly affect the choice of maximal dataset size.

In the static case, $k_T(t,t') = 1$ and γ_n is the mutual information associated with k_S. Setting $n = T$ (no dataset cap) in (4) yields $R_T \leq \sqrt{TC_1\beta_T\gamma_T} + \pi^2/6$, recovering the standard sublinear GPUCB bound of Theorem 2 in Srinivas et al. (2012).

For fixed n and $T \gg n$, γ_n/T vanishes and $T - n \sim T$. The leading term in (4) therefore depends on n through $\sqrt{1 - C_2\|\boldsymbol{u}_n\|_2^2}$. Maximizing $\|\boldsymbol{u}_n\|_2^2$ thus minimizes the leading asymptotic term of the bound, motivating the following recommendation.

Recommendation 1. In the long-horizon regime, a TVBO algorithm should maintain a dataset of size n^*, where

$$n^* = \arg\max_{n \in \mathbb{N}} \|\boldsymbol{u}_n\|_2^2 = \arg\max_{n \in \mathbb{N}} \sum_{i=1}^{n} k_T^2(iR(n)). \tag{5}$$

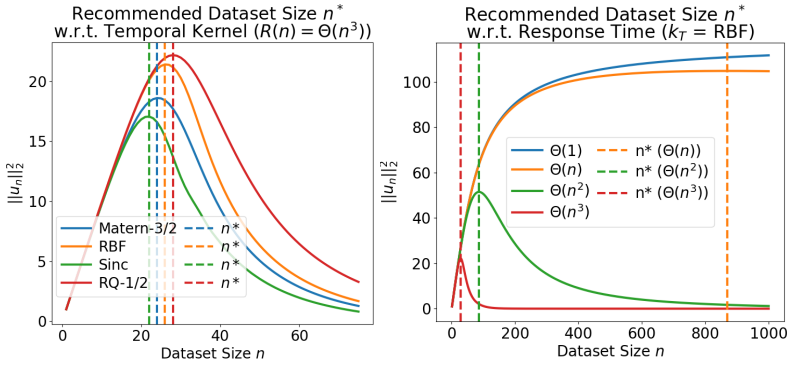


Figure 1: Recommended dataset size $n^* = \arg\max_{n \in \mathbb{N}} \|\boldsymbol{u}_n\|_2^2$. (Left) Recommended dataset size for several common temporal covariance functions k_T, under the assumption that the response time is $R(n) \in \Theta(n^3)$. (Right) The recommended dataset size for several response times $R(n)$, under the assumption that k_T is an RBF covariance function.

Computability of Recommendation 1. Assumptions 3.1-3.4 do not guarantee that $\|\boldsymbol{u}_n\|_2^2$ admits a maximizer over $\mathbb{N}$ (see the right panel of Figure 1 for an empirical example, and Appendix B for sufficient conditions). This has no practical consequence: at iteration T, the dataset contains at most T observations, so any maximal size larger than T has the same effect as $n^* = T$, and the search in (5) can be restricted to $n \in [T]$, over which the maximum is always attained. This range widens as the experiment goes on: it settles on the finite maximizer of $\|\boldsymbol{u}_n\|_2^2$ when one exists, and otherwise keeps recommending a larger dataset, that is, retaining every observation collected so far.

Figure 1 illustrates how $\|\boldsymbol{u}_n\|_2^2$ and n^* behave under different kernels and response times. The left panel shows that the smoothest temporal covariance functions (e.g., RBF or Rational Quadratic) also yield the largest n^*, as expected because a smoother GP surrogate can extract useful information from observations collected further in the past. The right panel shows that the response time of the TVBO algorithm has a significant impact on $\|\boldsymbol{u}_n\|_2^2$ and n^*: response times that scale slowly with the dataset size n lead to larger n^*. In the case where $R(n)$ does not scale with the dataset size n (i.e., $R(n) \in \Theta(1)$), $\|\boldsymbol{u}_n\|_2^2$ is always increasing; therefore, a dataset of size T is recommended.

3.3 Removal of Irrelevant Observations

In the previous section, we used a FIFO policy to study the joint effect of the temporal kernel k_T, the response time $R(n)$, and the maximal dataset size n. This analysis led to Recommendation 1. The simplicity of the FIFO policy makes it an analytically tractable reference, but it might be suboptimal because it needs not retain the most informative observations.

Given the maximal dataset size prescribed by Recommendation 1, the remaining question is therefore to select the observation to remove whenever this size is exceeded. Bardou et al. (2024b) propose to select $\boldsymbol{o}_i \in \mathcal{D}$ that minimizes the integrated 2-Wasserstein distance (Kantorovich, 1960) between the GP surrogate $\mathcal{GP}_{\mathcal{D}}\left(\mu_{\mathcal{D}}, \sigma^2_{\mathcal{D}}\right)$ conditioned on $\mathcal{D}$ and the surrogate $\mathcal{GP}_{\widetilde{\mathcal{D}}}\left(\mu_{\widetilde{\mathcal{D}}}, \sigma^2_{\widetilde{\mathcal{D}}}\right)$ conditioned on $\widetilde{\mathcal{D}} = \mathcal{D} \setminus \{\boldsymbol{o}_i\}$. This policy achieves good empirical performance (Bardou et al., 2024b), but its connection to the subsequent optimization decisions has not been formally established. In this section, we show that minimizing this distance controls the perturbation of the acquisition function induced by removing an observation, thereby providing a theoretical justification for the policy under the following assumption.

Assumption 3.7 (Lipschitz Acquisition)**.** At iteration $T \in \mathbb{N}$, the acquisition function $\alpha : \mathcal{S} \times \mathcal{T} \to \mathbb{R}$ is built from the posterior mean $\mu(\boldsymbol{x}, t)$ and the posterior standard deviation $\sigma(\boldsymbol{x}, t)$ of the GP surrogate, that is, $\alpha(\boldsymbol{x}, t) = g_T(\mu(\boldsymbol{x}, t), \sigma(\boldsymbol{x}, t))$ for some function $g_T : \mathbb{R}^2 \to \mathbb{R}$. Furthermore, g_T is a Lipschitz continuous function, that is,

$$|g_T(\boldsymbol{u}) - g_T(\boldsymbol{v})| \leq L_T ||\boldsymbol{u} - \boldsymbol{v}||_2, \ \forall \boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^2. \tag{6}$$

Algorithm 1 BOLT

Input: objective $f : \mathcal{S} \times \mathcal{T} \to \mathbb{R}$, acquisition function α, clock $\mathcal{C}$
Init dataset $\mathcal{D} = \emptyset$
while true do
  Get current time t from $\mathcal{C}$
  Find $x_t = \arg\max_{\boldsymbol{x} \in \mathcal{S}} \alpha_{\mathcal{D}}(\boldsymbol{x}, t)$
  Observe $y = f(\boldsymbol{x}_t, t) + \epsilon$
  $\mathcal{D} = \mathcal{D} \cup \{((\boldsymbol{x}_t, t), y)\}$
  Infer hyperparams and response time $R(n)$ from data
  Find the recommended n^* with (5)
  if $|\mathcal{D}| > n^*$ **then**
    Find $\boldsymbol{o}^* = \arg\min_{\boldsymbol{o} \in \mathcal{D}} W_2\left(\mathcal{GP}_{\mathcal{D}}, \mathcal{GP}_{\mathcal{D} \setminus \{\boldsymbol{o}\}}\right)$
    $\mathcal{D} = \mathcal{D} \setminus \{\boldsymbol{o}^*\}$
  end if
end while

Assumption 3.7 holds for many acquisition functions. For example, if α is GPUCB (Srinivas et al., 2012), that is, $\alpha(\boldsymbol{x}, t) = \mu(\boldsymbol{x}, t) + \beta_T^{1/2} \sigma(\boldsymbol{x}, t)$, then Assumption 3.7 holds with $L_T = \sqrt{1 + \beta_T}$.

Theorem 3.8. *Let α be an acquisition function satisfying Assumption 3.7. Let us denote by $\alpha_{\mathcal{D}}$ (resp., $\alpha_{\tilde{\mathcal{D}}}$) the acquisition function α exploiting the surrogates $\mathcal{GP}_{\mathcal{D}}$ (resp., $\mathcal{GP}_{\tilde{\mathcal{D}}}$). Then, on any subset $\mathcal{S}' \times \mathcal{T}'$ of the spatio-temporal domain $\mathcal{S} \times \mathcal{T}$:*

$$||\alpha_{\mathcal{D}} - \alpha_{\tilde{\mathcal{D}}}||_2 \leq L_T W_2\left(\mathcal{GP}_{\mathcal{D}}, \mathcal{GP}_{\tilde{\mathcal{D}}}\right), \tag{7}$$

where $||\alpha_{\mathcal{D}} - \alpha_{\tilde{\mathcal{D}}}||_2$ is the L^2 distance between $\alpha_{\mathcal{D}}$ and $\alpha_{\tilde{\mathcal{D}}}$ on $\mathcal{S}' \times \mathcal{T}'$ and $W_2\left(\mathcal{GP}_{\mathcal{D}}, \mathcal{GP}_{\tilde{\mathcal{D}}}\right)$ is the integrated 2-Wasserstein distance between the Gaussian marginals of $\mathcal{GP}_{\mathcal{D}}$ and $\mathcal{GP}_{\tilde{\mathcal{D}}}$ on $\mathcal{S}' \times \mathcal{T}'$.

Theorem 3.8 is proven in Appendix C. Assumption 3.7 allows to upper bound the L^2 distance between $\alpha_{\mathcal{D}}$ and $\alpha_{\tilde{\mathcal{D}}}$ on any subset of $\mathcal{S} \times \mathcal{T}$ in terms of the posterior means and variances of $\mathcal{GP}_{\mathcal{D}}$ and $\mathcal{GP}_{\tilde{\mathcal{D}}}$, which in turn naturally leads to the integrated Wasserstein distance between the GP surrogates.

Theorem 3.8 provides an interesting connection between the discrepancy of the acquisition functions built from two surrogate models $\mathcal{GP}_{\mathcal{D}}$ and $\mathcal{GP}_{\tilde{\mathcal{D}}}$, and the integrated 2-Wasserstein distance between these two surrogates. This connection motivates the following stale data management policy: given a maximal dataset size n and a dataset $\mathcal{D}$ of size T (where $T > n$), one should build $\tilde{\mathcal{D}} = \mathcal{D} \setminus X^*$, where $X^* = \arg\min_{X \in 2^{\mathcal{D}}, |X| = T - n} W_2(\mathcal{GP}_{\mathcal{D}}, \mathcal{GP}_{\mathcal{D} \setminus X})$. Unfortunately, this policy would become intractable in the long run. Therefore, to deal with stale data, we recommend the same greedy heuristic as the one introduced in Bardou et al. (2024b), which can be run online.

Recommendation 2. Given $\mathcal{D} = \{\boldsymbol{o}_1, \ldots, \boldsymbol{o}_n\}$, let $\widetilde{\mathcal{D}}_i = \mathcal{D} \setminus \{\boldsymbol{o}_i\}$ and remove $\boldsymbol{o}_{i^*}$, where $i^* = \arg\min_{i \in [n]} W_2(\mathcal{GP}_{\mathcal{D}}, \mathcal{GP}_{\widetilde{\mathcal{D}}_i})$.

4 BOLT: Bayesian Optimization in the Long Term

In the previous section, we have made two recommendations about the dataset size and the observation removal policy for TVBO algorithms, which are only useful when the optimization task is a long-horizon problem (see Section 2.2), so that removing observations from $\mathcal{D}$ produces a noticeable effect on the response time $R(n)$ and, therefore, on performance.

We now introduce BOLT, a TVBO algorithm implementing both recommendations (Algorithm 1). At each iteration, it selects a query $(\boldsymbol{x}, t)$, observes its response $f(\boldsymbol{x}, t)$, and adds the observation to $\mathcal{D}$. Recommendation 1 determines the maximal dataset size from k_T and $R(n)$, while Recommendation 2 selects which observations to discard if removal is needed.

BOLT is based on W-DBO (Bardou et al., 2024b), but it replaces its user-specified removal budget with the maximal dataset size derived from Theorem 3.6. Through its dependence on $R(n)$, this size (5) adapts to differences in response times (caused by, e.g., different computing powers).

Estimating the Response Time. Algorithm 1 estimates $R(n)$ online, similarly to the GP hyperparameters (variance λ, spatial and temporal lengthscales, and noise variance σ_0^2). At each iteration, the response time $R(|\mathcal{D}|)$ is paired with the dataset size $|\mathcal{D}|$ and recorded in a dataset $\mathcal{R}$. Motivated by the $\mathcal{O}(n^3)$ complexity of GP inference, BOLT fits a cubic polynomial to $\mathcal{R}$ to estimate $R(n)$.

5 Numerical Results

In this section, we evaluate BOLT against the state-of-the-art of TVBO. All algorithms in the TVBO literature that use an infinite variational budget for the objective function f are included: ABO (Nyikosa et al., 2018), ET-GPUCB (Brunzema et al., 2025), W-DBO (Bardou et al., 2024b) (which implements Recommendation 2 but not Recommendation 1), TV-GPUCB and R-GPUCB (Bogunovic et al., 2016). We also include the vanilla GPUCB algorithm (Srinivas et al., 2012) and FIFO (which implements Recommendation 1 but not Recommendation 2).

5.1 Experimental Setting

We evaluate the TVBO algorithms on a set of 10 synthetic and real-world benchmarks, which are described in detail in Appendix D. Some can be viewed as either quasi-static or as having a near-constant response time (see Section 2.2 for definitions of these types of problems). Each benchmark is a $(d+1)$-dimensional function to optimize. The first d dimensions make up the spatial domain $\mathcal{S}$, which is normalized in $[0,1]^d$. The $(d+1)$th dimension is the time dimension.

Each experiment starts with a warm-up phase consisting of 15 random queries. Then, at each iteration, the TVBO algorithms must perform the following tasks in real-time:

- **Hyperparameters inference**: the variance λ, the spatial lengthscale l_S and the observational noise σ_0^2 must always be inferred. When applicable, the temporal lengthscale l_T and the response time $R(n)$ of the algorithm are also inferred at this stage. Unless requested otherwise by the algorithm, each solution uses a Matérn-5/2 spatial covariance function, and a Matérn-3/2 temporal covariance function.
- **Optimization of the acquisition function**: the acquisition function is always GPUCB (Srinivas et al., 2012), and the next observation $\boldsymbol{x}_t$ is found by using multi-start gradient ascent.
- **Function observation**: a noisy function value $y_i = f(\boldsymbol{x}_i, t_i) + \epsilon$ is observed, and the observation $(\boldsymbol{x}_i, t_i, y_i)$ is added to the dataset $\mathcal{D}$.
- **Dataset cleaning**: this is the stage in which some observations in $\mathcal{D}$ might be removed.

Each TVBO algorithm has been implemented with the same BO framework, namely BOTorch (Balandat et al., 2020). The critical, time-consuming stages (i.e., hyperparameters inference and acquisition function optimization) are performed using the same BOTorch routines. Finally, all experiments have been independently replicated 10 times on a laptop equipped with an Intel Core i9-9980HK @ 2.40 GHz with 8 cores (16 threads).

5.2 Experimental Results

Table 3 reports average regrets and standard errors for each algorithm and benchmark (Appendix D), while Figure 2 summarizes aggregate performance.

Figure 3 illustrates how BOLT adapts its dataset size on the synthetic benchmarks Eggholder and Powell. On Eggholder (left), it discards many observations, consistent with the short temporal correlations of this benchmark. Algorithms that remove observations (R-GPUCB, W-DBO, FIFO, and BOLT) consequently perform well. On Powell (right), observations remain informative over longer periods, and BOLT retains a growing dataset, similarly to TV-GPUCB. Despite their comparable dataset sizes, BOLT achieves substantially lower regret (Table 3). A possible explanation is its more flexible temporal model under Assumptions 3.1, 3.2, and 3.3: GPUCB ignores temporal variation, whereas TV-GPUCB assumes Markovian dynamics (Bogunovic et al., 2016).

The benchmarks also cover quasi-static problems (e.g., Temperature) and near-constant response times (e.g., Shekel), as listed in Table 2 of Appendix D. Performance is broadly similar in the near-constant regime, apart from ET-GPUCB, whose difficulties with online hyperparameter inference are

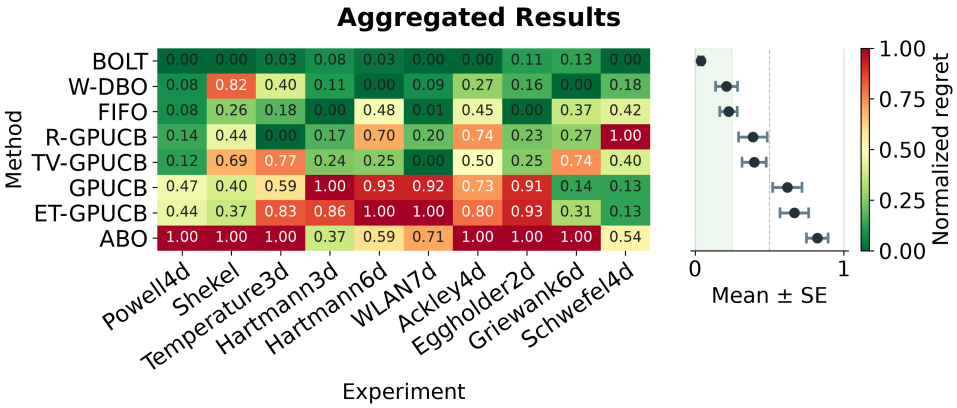


Figure 2: (Left) Heatmap of normalized regret (lower is better) for each benchmark and each method. For each benchmark, the best performing TVBO algorithm gets a normalized regret of 0, and the worst performing TVBO algorithm gets a normalized regret of 1. (Right) Averaged normalized regret across the benchmarks (lower is better).

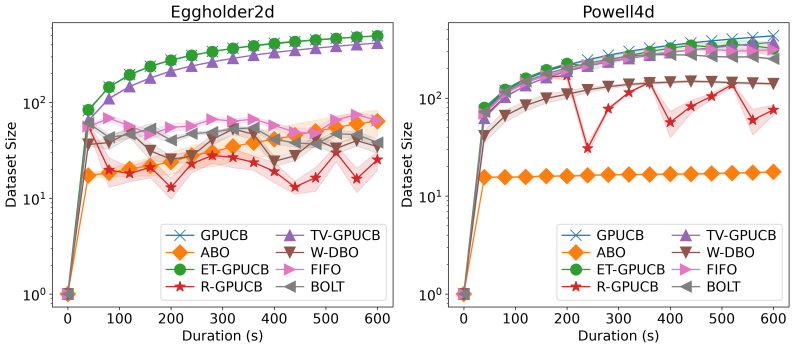


Figure 3: Dataset sizes on Eggholder (left) and Powell (right), with logarithmic y-axes.

discussed by Brunzema et al. (2025). On quasi-static problems, methods that remove observations (R-GPUCB, W-DBO, FIFO, and BOLT) perform best. Across these regimes and the remaining benchmarks, BOLT consistently achieves competitive or superior performance (Figure 2). This does not come at the cost of a significantly larger response time, which stays in the same range as for W-DBO, FIFO, and TV-GPUCB (Figures 4–13, Appendix D).

6 Conclusion

In this paper, we developed a regret analysis that explicitly accounts for the dependence of computational overhead and observation sampling frequency on the dataset size of a TVBO algorithm. Our main result, Theorem 3.6, establishes an asymptotic upper regret bound for GPUCB applied to a spatio-temporal GP with a FIFO policy. To the best of our knowledge, it is the first result to jointly relate TVBO regret to the temporal covariance structure and the algorithm's response time, yielding a principled recommendation for the maximal dataset size. Our second result, Theorem 3.8, bounds the acquisition-function perturbation induced by observation removal using the integrated Wasserstein distance between GP surrogates, providing theoretical support for the policy of Bardou et al. (2024b). Together, these results inform two complementary decisions: how many observations to retain and which to remove. We combined these recommendations in BOLT, which consistently achieves competitive or superior performance across diverse synthetic and real-world benchmarks.

Future work will extend this analysis beyond Assumptions 3.2 and 3.3. In particular, nonseparable spatio-temporal kernels could capture dynamics whose temporal behavior varies across the spatial domain and could be applied to more complex spatio-temporal black-box optimization problems.

AI USE STATEMENT

In this work, we used generative AI tools to critically review mathematical proofs, identify potential gaps and visualize the experimental results. We did not use generative AI tools to originate the research questions, theoretical framework, mathematical claims, methodology, or experiments; nor were they used to implement the methods, generate or process datasets. The remaining tasks requiring disclosure are not applicable to this work. Additionally, we used generative AI tools to review the manuscript's prose and polish some paragraphs for readability.

All AI-assisted suggestions were critically assessed by the authors. In particular, the authors retained full control over the final proofs, claims, and wording. We take responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

The assumptions underlying our theoretical results are stated and discussed in Sections 3.2 and 3.3, while complete proofs and additional technical results are provided in Appendices A-C. The full BOLT procedure is specified in Section 4. Section 5 describes the experimental protocol, evaluated baselines, implementation framework, computational environment, and number of independent repetitions. Detailed definitions of all synthetic and real-world benchmarks, including their domains, noise levels, evaluation costs, time horizons, and temporal characteristics, are provided in Appendix D. The source code and scripts required to reproduce the numerical experiments will be released publicly upon acceptance.

REFERENCES

Charu C Aggarwal, Jiawei Han, Jianyong Wang, and Philip S Yu. A framework for projected clustering of high dimensional data streams. In *Proceedings of the Thirtieth international conference on Very large data bases-Volume 30*, pp. 852–863, 2004.

Maximilian Balandat, Brian Karrer, Daniel R. Jiang, Samuel Daulton, Benjamin Letham, Andrew Gordon Wilson, and Eytan Bakshy. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. In *Advances in Neural Information Processing Systems 33*, 2020. URL http://arxiv.org/abs/1910.06403.

Anthony Bardou, Patrick Thiran, and Thomas Begin. Relaxing the additivity constraints in decentralized no-regret high-dimensional bayesian optimization. In *The Twelfth International Conference on Learning Representations*, 2024a.

Anthony Bardou, Patrick Thiran, and Giovanni Ranieri. This Too Shall Pass: Removing Stale Observations in Dynamic Bayesian Optimization. In *The 38th Annual Conference on Neural Information Processing Systems*, 2024b.

Anthony Bardou, Jean-Marie Gorce, and Thomas Begin. Assessing the performance of noma in a multi-cell context: A general evaluation framework. *IEEE Transactions on Wireless Communications*, 2025.

Ilija Bogunovic, Jonathan Scarlett, and Volkan Cevher. Time-varying gaussian process bandit optimization. In *Artificial Intelligence and Statistics*, pp. 314–323. PMLR, 2016.

Paul Brunzema, Alexander von Rohr, Friedrich Solowjow, and Sebastian Trimpe. Event-triggered time-varying bayesian optimization. *Transactions on Machine Learning Research*, 2025.

Samuel Daulton, David Eriksson, Maximilian Balandat, and Eytan Bakshy. Multi-objective bayesian optimization over high-dimensional search spaces. In *Uncertainty in Artificial Intelligence*, pp. 507–517. PMLR, 2022.

Yuntian Deng, Xingyu Zhou, Baekjin Kim, Ambuj Tewari, Abhishek Gupta, and Ness Shroff. Weighted gaussian process bandits for non-stationary environments. In *International Conference on Artificial Intelligence and Statistics*, pp. 6909–6932. PMLR, 2022.

Donald R. Jones, Matthias Schonlau, and William J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492, 1998.

Leonid V Kantorovich. Mathematical methods of organizing and planning production. *Management science*, 6(4):366–422, 1960.

JHB Kemperman. On the shannon capacity of an arbitrary channel. In *Indagationes Mathematicae (Proceedings)*, volume 77, pp. 101–115. North-Holland, 1974.

Seokhyun Kim, Kimin Lee, Yeonkeun Kim, Jinwoo Shin, Seungwon Shin, and Song Chong. Dynamic control for on-demand interference-managed wlan infrastructures. *IEEE/ACM Transactions on Networking*, 28(1):84–97, 2019.

Roman Marchant and Fabio Ramos. Bayesian optimisation for intelligent environmental monitoring. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, pp. 2242–2249. IEEE, 2012.

Aurelio G Melo, Milena F Pinto, Andre LM Marcato, Leonardo M Honório, and Fabrício O Coelho. Dynamic optimization and heuristics based online coverage path planning in 3d environment for uavs. *Sensors*, 21(4):1108, 2021.

Jonas Mockus. Application of bayesian approach to numerical methods of global and stochastic optimization. *Journal of Global Optimization*, 4:347–365, 1994.

Favour M Nyikosa, Michael A Osborne, and Stephen J Roberts. Bayesian optimization for dynamic problems. *arXiv preprint arXiv:1803.03432*, 2018.

Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias W. Seeger. Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, 58(5):3250–3265, 2012. doi: doi:10.1109/tit.2011.2182033.

Ziyu Wang, Babak Shakibi, Lin Jin, and Nando Freitas. Bayesian multi-scale optimistic optimization. In *Artificial Intelligence and Statistics*, pp. 1005–1014. PMLR, 2014.

Xingyu Zhou and Ness Shroff. No-regret algorithms for time-varying bayesian optimization. In *2021 55th Annual Conference on Information Sciences and Systems (CISS)*, pp. 1–6. IEEE, 2021.

A Upper Regret Bound without a Metronome

In this appendix, we provide all the details required to prove Theorem 3.6. For the sake of completeness, we start by deriving the usual instantaneous regret bound provided in most regret proofs that involve GPUCB (Srinivas et al., 2012; Bogunovic et al., 2016).

Lemma A.1. *Let $r_n = f(\boldsymbol{x}_n^*, t_n) - f(\boldsymbol{x}_n, t_n)$ be the instantaneous regret at the n-th iteration of the TVBO algorithm $\mathcal{A}$, where $\boldsymbol{x}_n^* = \arg\max_{\boldsymbol{x}\in\mathcal{S}} f(\boldsymbol{x}, t_n)$, $\boldsymbol{x}_n = \arg\max_{\boldsymbol{x}\in\mathcal{S}} \alpha_{\mathcal{D}}(\boldsymbol{x}, t_n)$, $\alpha_{\mathcal{D}}$ is the GPUCB acquisition function computed on $\mathcal{D}$ and where $\mathcal{D}$ is a dataset of observations collected with the distribution μ. Pick $\delta \in (0,1)$, then with probability at least $1-\delta$,*

$$r_n \leq 2\beta_n^{1/2}\sigma_{\mathcal{D}}(\boldsymbol{x}_n, t_n) + \frac{1}{n^2} \tag{8}$$

where $\sigma_{\mathcal{D}}(\boldsymbol{x}, t)$ is the posterior standard deviation (see (2)) and where

$$\beta_n = 2d\log\left(bdn^2\sqrt{\log(da\pi^2n^2/3\delta)}/6\delta\right) + 4\log(\pi n), \tag{9}$$

with $a, b > 0$ given by Assumption 3.4.

Proof. For the sake of consistency with the literature, we will reuse, whenever appropriate, the notations in Bogunovic et al. (2016).

Let us set a discretization $\mathcal{S}_n$ of the spatial domain $\mathcal{S} \subseteq [0,1]^d$. $\mathcal{S}_n$ is of size τ_n^d and satisfies

$$||\boldsymbol{x} - [\boldsymbol{x}]_n||_1 \leq \frac{d}{\tau_n},\ \forall \boldsymbol{x} \in \mathcal{S}, \tag{10}$$

where $[\boldsymbol{x}]_n = \arg\min_{\boldsymbol{s}\in\mathcal{S}_n} ||\boldsymbol{s} - \boldsymbol{x}||_1$ is the closest point in $\mathcal{S}_n$ to $\boldsymbol{x}$. Note that letting S_n be a uniform grid with discretization step $1/\tau_n$ on each dimension of $\mathcal{S}$ ensures (10).

Let us now fix $\delta > 0$ and condition on a high-probability event. If $\beta_n = 2\log\frac{\tau_n^d\pi^2n^2}{3\delta}$, then

$$|f(\boldsymbol{x}, t_n) - \mu_{\mathcal{D}}(\boldsymbol{x}, t_n)| \leq \beta_n^{1/2}\sigma_{\mathcal{D}}(\boldsymbol{x}, t_n),\ \forall n \in \mathbb{N}, \forall \boldsymbol{x} \in \mathcal{S}_n \tag{11}$$

with probability at least $1 - \delta/2$. This directly comes from $f(\boldsymbol{x}, t_n) \sim \mathcal{N}\left(\mu_{\mathcal{D}}(\boldsymbol{x}, t_n), \sigma_{\mathcal{D}}^2(\boldsymbol{x}, t_n)\right)$ and from the Chernoff concentration inequality applied to Gaussian tails $\mathbb{P}\left(|f(\boldsymbol{x}, t_n) - \mu_{\mathcal{D}}(\boldsymbol{x}, t_n)| \leq \sqrt{\beta_n}\sigma_{\mathcal{D}}(\boldsymbol{x}, t_n)\right) \geq 1 - e^{-\beta_n/2}$. Therefore, our choice of β_n ensures that for a given n and a given $\boldsymbol{x} \in \mathcal{S}_n$, $|f(\boldsymbol{x}, t_n) - \mu_{\mathcal{D}}(\boldsymbol{x}, t_n)| \leq \beta_n^{1/2}\sigma_{\mathcal{D}}(\boldsymbol{x}, t_n)$ occurs with a probability at least $1 - \frac{3\delta}{\pi^2n^2\tau_n^d}$. A union bound taken over $n \in \mathbb{N}$ and $\boldsymbol{x} \in \mathcal{S}_n$ establishes (11) with probability $1 - \delta/2$.

In particular, setting $\tau_n = Ldn^2$, where L is defined in Assumption 3.4, gives that for all $\boldsymbol{x} \in \mathcal{S}$ and all $t_n \in \mathcal{T}$,

$$|f(\boldsymbol{x}, t_n) - f([\boldsymbol{x}]_n, t_n)| \leq \frac{1}{n^2} \tag{12}$$

with probability at least $1 - \delta/2$.

Because $\tau_n = Ldn^2$, β_n becomes $2d\log\left(b\sqrt{\log(da\pi^2n^2/3\delta)}dn^2/6\delta\right) + 4\log(\pi n)$. Using the triangle inequality and combining (11) with (12), we get that with probability $1-\delta$, $\boldsymbol{x}_n^* = \arg\max_{\boldsymbol{x}\in\mathcal{S}} f(\boldsymbol{x}, t_n)$ satisfies

$$\begin{aligned}|f(\boldsymbol{x}_n^*, t_n) - \mu_{\mathcal{D}}([\boldsymbol{x}_n^*]_n, t_n)| &\leq |f([\boldsymbol{x}_n^*]_n, t_n) - \mu_{\mathcal{D}}([\boldsymbol{x}_n^*]_n, t_n)| + |f(\boldsymbol{x}_n^*, t_n) - f([\boldsymbol{x}_n^*]_n, t_n)| \\ &\leq \beta_n^{1/2}\sigma_{\mathcal{D}}([\boldsymbol{x}_n^*]_n, t_n) + \frac{1}{n^2}. \end{aligned} \tag{13}$$

We can now upper bound the instantaneous regret of the TVBO algorithm $\mathcal{A}$:

$$\begin{aligned} r_n &= f(\boldsymbol{x}_n^*, t_n) - f(\boldsymbol{x}_n, t_n) \\ &\leq \mu_{\mathcal{D}}([\boldsymbol{x}_n^*]_n, t_n) + \beta_n^{1/2}\sigma_{\mathcal{D}}([\boldsymbol{x}_n^*]_n, t_n) + \frac{1}{n^2} - f(\boldsymbol{x}_n, t_n) && (14) \\ &\leq \mu_{\mathcal{D}}(\boldsymbol{x}_n, t_n) + \beta_n^{1/2}\sigma_{\mathcal{D}}(\boldsymbol{x}_n, t_n) + \frac{1}{n^2} - f(\boldsymbol{x}_n, t_n) && (15) \\ &\leq 2\beta_n^{1/2}\sigma_{\mathcal{D}}(\boldsymbol{x}_n, t_n) + \frac{1}{n^2}, && (16) \end{aligned}$$

where (14) follows directly from (13), (15) from the definition of $\boldsymbol{x}_n = \arg\max_{\boldsymbol{x}\in\mathcal{S}} \alpha_{\mathcal{D}}(\boldsymbol{x}, t_n) = \arg\max_{\boldsymbol{x}\in\mathcal{S}} \mu_{\mathcal{D}}(\boldsymbol{x}, t_n) + \beta_n^{1/2}\sigma_{\mathcal{D}}(\boldsymbol{x}, t_n)$ and (16) from (11). □

Lemma A.1 shows that the posterior variance of the surrogate GP (2) plays a key role in the regret bound. Let us now derive a first bound on the cumulative regret R_T, by extending an original idea from Srinivas et al. (2012) to the time-varying setting.

Lemma A.2. *Let $\{\mathcal{D}_i\}_{i=1,\cdots,T}$ be a sequence of datasets of size $|\mathcal{D}_i| = \min(i, n)$ for $i = 1, \cdots, T$. Let $R_T = \sum_{i=1}^T r_i$ be the cumulative regret of a TVBO algorithm with maximal dataset size n. Then,*

$$R_T \leq \sqrt{C_1 T \beta_T \left(\gamma_n + \frac{1}{2}\sum_{i=n+1}^{T} \log(1 + \sigma_0^{-2}\sigma^2_{\mathcal{D}_i}(\boldsymbol{x}_i, t_i))\right)} + \frac{\pi^2}{6} \tag{17}$$

where

$$C_1 = \frac{8}{\log(1 + \sigma_0^{-2})}, \tag{18}$$

and where γ_n is the mutual information $\gamma_n = \frac{1}{2}\sum_{i=1}^n \log(1 + \sigma_0^{-2}\sigma^2_{\mathcal{D}_i}(\boldsymbol{x}_i, t_i))$.

Proof. We have

$$\begin{aligned} R_T &= \sum_{i=1}^T r_i \\ &\leq 2\left(\sum_{i=1}^n \beta_i^{1/2}\sigma_{\mathcal{D}_i}(\boldsymbol{x}_i, t_i) + \sum_{i=n+1}^T \beta_i^{1/2}\sigma_{\mathcal{D}_i}(\boldsymbol{x}_i, t_i)\right) + \frac{\pi^2}{6} \qquad (19) \\ &\leq \sqrt{4T\beta_T\left(\sum_{i=1}^n \sigma^2_{\mathcal{D}_i}(\boldsymbol{x}_i, t_i) + \sum_{i=n+1}^T \sigma^2_{\mathcal{D}_i}(\boldsymbol{x}_i, t_i)\right)} + \frac{\pi^2}{6} \qquad (20) \end{aligned}$$

where (19) uses Lemma A.1, the solution to the Basel problem $\sum_{i=1}^T i^{-2} \leq \sum_{i=1}^\infty i^{-2} = \pi^2/6$ and the maximal dataset size n of the TVBO algorithm, while (20) uses the Cauchy-Schwarz inequality and $\beta_i \leq \beta_T$ for any $i \leq T$.

Going further, we get

$$\begin{aligned} R_T &\leq \sqrt{4\sigma_0^2 T\beta_T\left(\sum_{i=1}^n \sigma_0^{-2}\sigma^2_{\mathcal{D}_i}(\boldsymbol{x}_i, t_i) + \sum_{i=n+1}^T \sigma_0^{-2}\sigma^2_{\mathcal{D}_i}(\boldsymbol{x}_i, t_i)\right)} + \frac{\pi^2}{6} \\ &\leq \sqrt{\frac{8}{\log(1+\sigma_0^{-2})} T\beta_T\left(\frac{1}{2}\sum_{i=1}^n \log(1 + \sigma_0^{-2}\sigma^2_{\mathcal{D}_i}(\boldsymbol{x}_i, t_i)) + \frac{1}{2}\sum_{i=n+1}^T \log(1 + \sigma_0^{-2}\sigma^2_{\mathcal{D}_i}(\boldsymbol{x}_i, t_i))\right)} + \frac{\pi^2}{6} \qquad (21) \\ &= \sqrt{C_1 T\beta_T\left(\gamma_n + \frac{1}{2}\sum_{i=n+1}^T \log(1 + \sigma_0^{-2}\sigma^2_{\mathcal{D}_i}(\boldsymbol{x}_i, t_i))\right)} + \frac{\pi^2}{6} \qquad (22) \end{aligned}$$

where (21) uses the identity $z^2 \leq \sigma_0^{-2}\log(1 + z^2)/\log(1 + \sigma_0^{-2})$ for $z^2 \in [0, \sigma_0^{-2}]$ and where (22) uses (18) and the definition of mutual information γ_n. □

We now provide an upper bound for the posterior variance $\sigma^2_{\mathcal{D}}(\boldsymbol{x}, t)$.

Lemma A.3. *Let $\mathcal{A}$ be a TVBO algorithm, $\mathcal{D}$ be its dataset of observations with fixed maximal size n and $R(n)$ be its response time. Let*

$$G = \sup_{m\geq 1} \frac{1}{R(m)} \int_0^{mR(m)} g_T(\tau) d\tau, \tag{23}$$

where g_T is the function upper-bounding k_T in Assumption 3.3. Then,

$$\sigma^2_{\mathcal{D}}(\boldsymbol{x}, t) \leq 1 - C_2 \|\boldsymbol{u}_n\|_2^2, \tag{24}$$

where $C_2 = \min_{\boldsymbol{x},\boldsymbol{x}' \in \mathcal{S}} k_S^2(\boldsymbol{x}, \boldsymbol{x}')/(1 + 2G + \sigma_0^2)$ *and where*

$$\boldsymbol{u}_n = \left(k_T(R(n)), k_T(2R(n)), \cdots, k_T(nR(n))\right).$$

Proof. Let us start by recalling the definition of the posterior variance in (2):

$$\begin{aligned}\sigma^2_{\mathcal{D}}(\boldsymbol{x}, t) &= k((\boldsymbol{x}, t), (\boldsymbol{x}, t)) - \boldsymbol{k}^\top((\boldsymbol{x}, t), \mathcal{D})\boldsymbol{\Delta}^{-1}\boldsymbol{k}((\boldsymbol{x}, t), \mathcal{D}) \\ &= 1 - \boldsymbol{k}^\top((\boldsymbol{x}, t), \mathcal{D})\boldsymbol{\Delta}^{-1}\boldsymbol{k}((\boldsymbol{x}, t), \mathcal{D}) \end{aligned} \tag{25}$$

where (25) follows from Assumptions 3.2 and 3.3.

Finding an upper bound on $\sigma^2_{\mathcal{D}}(\boldsymbol{x}, t)$ boils down to finding a lower bound on the quadratic form $q_n(\boldsymbol{x}, t) = \boldsymbol{k}^\top((\boldsymbol{x}, t), \mathcal{D})\boldsymbol{\Delta}^{-1}\boldsymbol{k}((\boldsymbol{x}, t), \mathcal{D})$. Let us first consider the kernel vector $\boldsymbol{k}((\boldsymbol{x}, t), \mathcal{D})$. Because the dataset $\mathcal{D} = \{(\boldsymbol{x}_1, t_1, y_1), \cdots, (\boldsymbol{x}_n, t_n, y_n)\}$ has a fixed maximal dataset size n, for every iteration $m > 2n$:

$$\boldsymbol{k}^\top((\boldsymbol{x}, t), \mathcal{D}) = \left(k_S(\boldsymbol{x}, \boldsymbol{x}_n)k_T(R(n)), \cdots, k_S(\boldsymbol{x}, \boldsymbol{x}_1)k_T(nR(n))\right). \tag{26}$$

The expression in (26) is due to $\mathcal{A}$ having a fixed dataset size n, and therefore a constant response time $R(n)$. In such a dataset $\mathcal{D} = \{(\boldsymbol{x}_i, t_i, y_i)\}_{i\in[n]}$, the temporal input t_i is deterministic. Because $\mathcal{A}$ follows a FIFO removal policy, it first fills its dataset during the first n iterations. Once the maximal dataset size has been reached, its response time remains equal to $R(n)$, and the oldest observation is removed after every new query. Consequently, for every iteration $m > 2n$, the dataset consists of the n most recent observations, collected at constant intervals $R(n)$, and the kernel vector $\boldsymbol{k}^\top((\boldsymbol{x}, t), \mathcal{D})$ verifies (26).

Next, consider $\boldsymbol{\Delta} = \boldsymbol{k}(\mathcal{D}, \mathcal{D}) + \sigma_0^2\boldsymbol{I}$. Because it is a positive definite matrix, its spectral decomposition $\boldsymbol{Q}\boldsymbol{\Lambda}\boldsymbol{Q}^\top$ exists, where $\boldsymbol{Q} = (\boldsymbol{\phi}_1, \cdots, \boldsymbol{\phi}_n)$ is the orthogonal matrix whose columns are the eigenvectors of $\boldsymbol{\Delta}$, and $\boldsymbol{\Lambda} = \text{diag}\left(\lambda_1 + \sigma_0^2, \cdots, \lambda_n + \sigma_0^2\right)$ is the diagonal matrix comprising the associated eigenvalues. Note that, by convention, the eigenvalues are sorted in decreasing order $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$.

The quadratic form $q_n(\boldsymbol{x}, t)$ then becomes

$$\begin{aligned} q_n(\boldsymbol{x}, t) &= \boldsymbol{k}^\top((\boldsymbol{x}, t), \mathcal{D})\boldsymbol{Q}\boldsymbol{\Lambda}^{-1}\boldsymbol{Q}^\top\boldsymbol{k}((\boldsymbol{x}, t), \mathcal{D}) \\ &= \boldsymbol{v}^\top\boldsymbol{\Lambda}^{-1}\boldsymbol{v} \\ &= \sum_{i=1}^{n} \frac{v_i^2}{\lambda_i + \sigma_0^2} \\ &\geq \frac{1}{\lambda_1 + \sigma_0^2} \sum_{i=1}^{n}\sum_{j=1}^{n}\sum_{k=1}^{n} k((\boldsymbol{x}, t), (\boldsymbol{x}_j, t_j))\phi_{ij}\phi_{ik}k((\boldsymbol{x}, t), (\boldsymbol{x}_k, t_k)) \end{aligned} \tag{27}$$

where $\boldsymbol{v} = \boldsymbol{Q}^\top\boldsymbol{k}((\boldsymbol{x}, t), \mathcal{D})$ while (27) is the quadratic form $q_n(\boldsymbol{x}, t)$ rewritten with sums instead of matrix products and using the fact that $\lambda_1 \geq \lambda_i$ for all $i \in [n]$. Furthermore, note that the largest eigenvalue λ_1 of the Gram matrix $\boldsymbol{K}$ built from observations collected every $R(n)$ unit of time can be bounded from above:

$$\lambda_1 \leq \max_{i\in[n]} \boldsymbol{K}_{ii} + \sum_{j\neq i} |\boldsymbol{K}_{ij}| \tag{28}$$

$$\leq 1 + 2\sum_{j=1}^{n-1} |k_T(jR(n))| \tag{29}$$

$$\leq 1 + 2\sum_{j=1}^{n-1} g_T(jR(n)) \tag{30}$$

$$\leq 1 + \frac{2}{R(n)} \int_0^{(n-1)R(n)} g_T(\tau)d\tau \tag{31}$$

$$\leq 1 + 2G \tag{32}$$

where (28) is an application of Gershgorin's theorem for real and symmetric matrices, (29) holds because all the diagonal entries of $\boldsymbol{K}$ are equal to one, and because for all $j \in [n]$, there are at most two entries $k_T(jR(n))$ at each temporal lag $jR(n)$, and (30) holds under Assumption 3.3. Inequality (31) is an application of the quadrature rule on the truncated range $[0, (n-1)R(n)]$: because g_T is nonnegative and nonincreasing, $R(n)g_T(jR(n)) \leq \int_{(j-1)R(n)}^{jR(n)} g_T(\tau)d\tau$ for every $j \geq 1$, and summing these inequalities over $j = 1, \cdots, n-1$ yields (31). Finally, (32) holds because g_T is nonnegative, so that $\int_0^{(n-1)R(n)} g_T(\tau)d\tau \leq \int_0^{nR(n)} g_T(\tau)d\tau$, and because of the definition (23) of G.

Let $\bar{\lambda} = 1 + 2G$. Then, rearranging the sums in (27), we have

$$q_n(\boldsymbol{x}, t) \geq \frac{1}{\bar{\lambda} + \sigma_0^2} \sum_{j=1}^{n} \sum_{k=1}^{n} k((\boldsymbol{x}, t), (\boldsymbol{x}_j, t_j)) k((\boldsymbol{x}, t), (\boldsymbol{x}_k, t_k)) \underbrace{\sum_{i=1}^{n} \phi_{ij}\phi_{ik}}_{\delta_{jk}}$$

$$= \frac{1}{\bar{\lambda} + \sigma_0^2} \sum_{i=1}^{n} k^2((\boldsymbol{x}, t), (\boldsymbol{x}_i, t_i)) \tag{33}$$

$$= \frac{1}{\bar{\lambda} + \sigma_0^2} \sum_{i=1}^{n} k_S^2(\boldsymbol{x}, \boldsymbol{x}_i) k_T^2(t, t_i) \tag{34}$$

where δ_{jk} is the Kronecker delta with value 1 if $j = k$ and 0 otherwise, (33) is due to the orthonormality of the eigenvectors of $\boldsymbol{K}$ and (34) is due to Assumption 3.2.

Let $\kappa = \min_{\boldsymbol{x}, \boldsymbol{x}' \in \mathcal{S}} k_S(\boldsymbol{x}, \boldsymbol{x}')$. Then, $\kappa > 0$ as per Assumption 3.3 and we finally have

$$q_n(\boldsymbol{x}, t) \geq \frac{\kappa^2}{\bar{\lambda} + \sigma_0^2} \sum_{i=1}^{n} k_T^2(t, t_i)$$
$$= C_2 \|\boldsymbol{u}_n\|_2^2,$$

where $C_2 = \kappa^2/(\bar{\lambda} + \sigma_0^2)$.

Noting that $\sigma_{\mathcal{D}}^2(\boldsymbol{x}, t) = 1 - q_n(\boldsymbol{x}, t) \leq 1 - C_2\|\boldsymbol{u}_n\|_2^2$ concludes the proof. Recalling that $\sigma_{\mathcal{D}}^2(\boldsymbol{x}, t) \geq 0$ also yields that $0 \leq 1 - C_2\|\boldsymbol{u}_n\|_2^2$, confirming that $1 - C_2\|\boldsymbol{u}_n\|_2^2$ is non-negative. □

Remark A.4. The constant G can be finite even when g_T is not integrable. Suppose that $g_T(\tau) \leq c(1+\tau)^{-p}$ and $R(m) \geq c_R m^q$ for some $c, c_R, p > 0$ and $q \geq 0$. Then $G < \infty$ whenever $p(1+q) > 1$, a condition that does not require $p > 1$. In particular, temporal correlation functions whose tails decay as $\Theta(\tau^{-1})$, such as sinc or a Rational Quadratic with $\alpha = 1/2$, satisfy it for every $q > 0$, although no integrable g_T dominates them (see Appendix B for a detailed discussion on these two cases).

Finally, let us prove Theorem 3.6.

Proof. Combining Lemmas A.2 and A.3, we get

$$R_T \leq \sqrt{C_1 T \beta_T \left(\gamma_n + \frac{1}{2} \sum_{i=n+1}^{T} \log(1 + \sigma_0^{-2}(1 - C_2\|\boldsymbol{u}_n\|_2^2)) \right)} + \frac{\pi^2}{6}$$
$$\leq \sqrt{C_1 T \beta_T \left(\gamma_n + \frac{\sigma_0^{-2}}{2}(T-n)(1 - C_2\|\boldsymbol{u}_n\|_2^2) \right)} + \frac{\pi^2}{6}, \tag{35}$$

where (35) uses $\log(1+x) \leq x$ for any $x > 0$. □

For the sake of completeness, we conclude this appendix by showing that the information gain γ_n scales linearly with n in the time-varying setting.

Lemma A.5. *For all $n \in \mathbb{N}$, $\gamma_n \in \mathcal{O}(n)$.*

Proof. Recall that γ_n can be recovered from the covariance matrix $k(\mathcal{D}, \mathcal{D})$ and its spectrum. More specifically, denoting $\{\lambda_i\}_{1 \leq i \leq n}$ the ordered spectrum of $k(\mathcal{D}, \mathcal{D})$, we have

$$\begin{aligned}
\gamma_n &= \frac{1}{2} \sum_{i=1}^{n} \log(1 + \sigma_0^{-2} \lambda_i) \\
&\leq \frac{1}{2} \sigma_0^{-2} \sum_{i=1}^{n} \lambda_i && (36) \\
&= \frac{1}{2} \sigma_0^{-2} \operatorname{Tr}(k(\mathcal{D}, \mathcal{D})) && (37) \\
&= \frac{1}{2} \sigma_0^{-2} n, && (38)
\end{aligned}$$

where (36) uses the property $\log(1 + x) \leq x$ for all $x \in \mathbb{R}_{>0}$, (37) uses the property that the trace of a matrix is the sum of its eigenvalues, and (38) holds for any kernel that is stationary and normalized (i.e., that satisfies $k((\boldsymbol{x}, t), (\boldsymbol{x}, t)) = 1$ for all $(\boldsymbol{x}, t) \in \mathcal{S} \times \mathcal{T}$).

One immediately concludes from (38) that $\gamma_n \in \mathcal{O}(n)$, which also concludes the proof. □

B Sufficient Conditions for Benefiting from a Maximal Dataset Size

In this appendix, we derive sufficient conditions for $\|\boldsymbol{u}_n\|_2^2$ to admit a finite maximizer n^*. According to Recommendation 1, these become sufficient conditions for characterizing when capping the dataset size of a TVBO algorithm is beneficial. Throughout, p denotes the polynomial decay exponent of the temporal correlation function k_T and q the polynomial growth exponent of the response time R, while c_T and c_R denote the associated multiplicative constants.

Proposition B.1. *Suppose that there exist $p, q, c_T, c_R > 0$ such that*

$$|k_T(\tau)| \leq c_T (1 + \tau)^{-p} \quad \textit{and} \quad R(n) \geq c_R n^q$$

for every $\tau \geq 0$ and $n \geq 1$. Then

$$\|\boldsymbol{u}_n\|_2^2 \leq c\, n^{-2pq} \sum_{i=1}^{n} i^{-2p}$$

for some constant $c > 0$. Consequently,

$$\|\boldsymbol{u}_n\|_2^2 = \begin{cases} \mathcal{O}\left(n^{-2pq}\right), & p > 1/2, \\ \mathcal{O}\left(n^{-q} \log n\right), & p = 1/2, \\ \mathcal{O}\left(n^{1-2p(1+q)}\right), & p < 1/2. \end{cases}$$

In particular, if $2p(1 + q) > 1$, then $\|\boldsymbol{u}_n\|_2^2 \to 0$ and the sequence admits a finite maximizer.

Proof. Since $\tau \mapsto c_T(1 + \tau)^{-p}$ is nonincreasing and $R(n) \geq c_R n^q$, we have, for every $n \geq 1$ and every $i \in [n]$,

$$k_T^2(iR(n)) \leq c_T^2 \left(1 + i c_R n^q\right)^{-2p} \leq c_T^2 \left(i c_R n^q\right)^{-2p} = c\, n^{-2pq} i^{-2p},$$

where $c = c_T^2 c_R^{-2p}$. Summing this inequality over $i = 1, \ldots, n$ gives the first claim. The three asymptotic regimes then follow from the standard behavior of the partial sums of i^{-2p}, which are bounded when $2p > 1$, $\mathcal{O}(\log n)$ when $2p = 1$, and $\mathcal{O}(n^{1-2p})$ when $2p < 1$.

For the last claim, observe that $2p \geq 1$ implies $2p(1 + q) \geq 1 + q > 1$, so that the single condition $2p(1 + q) > 1$ includes the first two regimes. In those two regimes, $\|\boldsymbol{u}_n\|_2^2$ is $\mathcal{O}(n^{-2pq})$ and $\mathcal{O}(n^{-q} \log n)$ respectively, both of which vanish because $q > 0$. In the third regime, $2p(1 + q) > 1$ makes the exponent $1 - 2p(1 + q)$ negative, so $\|\boldsymbol{u}_n\|_2^2 \to 0$ as well. A nonnegative sequence converging to 0 attains its supremum at a finite index, which concludes the proof. □

The condition $2p(1+q) > 1$ admits a simple interpretation: capping the dataset size is beneficial when the temporal correlation decays fast enough (large p) relative to how fast the response time grows (large q). It fails only when $2p(1+q) \leq 1$, which, since $q > 0$, forces $p < 1/2$. The next proposition shows that, in this regime, $\|\boldsymbol{u}_n\|_2^2$ diverges when the polynomial decay of k_T is tight.

Proposition B.2. *Suppose that there exist $p, q, c'_T, c'_R > 0$ such that*

$$k_T(\tau) \geq c'_T(1+\tau)^{-p} \quad \textit{and} \quad R(n) \leq c'_R n^q$$

for every $\tau \geq 0$ and $n \geq 1$. If $2p(1+q) < 1$, then $\|\boldsymbol{u}_n\|_2^2 \to +\infty$, and the sequence admits no finite maximizer.

Proof. For every $n \geq 1$ and every $i \in [n]$, we have $1 \leq in^q$, hence $1 + iR(n) \leq 1 + ic'_R n^q \leq (1 + c'_R) in^q$. Therefore,

$$k_T^2(iR(n)) \geq c_T'^2 \left(1 + iR(n)\right)^{-2p} \geq c'\, n^{-2pq} i^{-2p},$$

where $c' = c_T'^2 (1 + c'_R)^{-2p}$. Since $q > 0$, the assumption $2p(1+q) < 1$ implies $2p < 1$, so that

$$\sum_{i=1}^{n} i^{-2p} \geq \int_1^{n+1} x^{-2p} dx = \frac{(n+1)^{1-2p} - 1}{1-2p}.$$

Summing the pointwise bound over $i = 1, \ldots, n$ therefore gives $\|\boldsymbol{u}_n\|_2^2 \in \Omega\left(n^{1-2p(1+q)}\right)$, which diverges because $1 - 2p(1+q) > 0$. A diverging sequence admits no finite maximizer. □

Remark B.3. Propositions B.1 and B.2 leave out the edge case $2p(1+q) = 1$, at which the upper bound of the former is $\mathcal{O}(1)$ and the lower bound of the latter is $\Omega(1)$. At this order the analysis is inconclusive, and whether n^* is finite or not has to be determined by another method.

The remaining regime $2p(1+q) < 1$ is, however, unreachable under Assumption 3.3. Indeed, suppose that the envelope satisfies $g_T(\tau) = \Theta(\tau^{-p})$ and that the response time satisfies $R(m) = \Theta(m^q)$ with $q > 0$. Then the finiteness of G required by Assumption 3.3 already imposes $p(1+q) \geq 1$, hence $2p(1+q) \geq 2$. Assumption 3.3 is therefore more demanding than the condition $2p(1+q) > 1$ of Proposition B.1, so that the divergence regime of Proposition B.2 and the edge case $2p(1+q) = 1$ both fall outside the setting of this paper.

Consequently, under Assumptions 3.1-3.4 and a response time growing polynomially with $q > 0$, every temporal correlation function with a polynomially decaying envelope admits a finite n^*: capping the dataset size is always beneficial in this setting. This does not, however, confine the analysis to the first regime of Proposition B.1. Because Assumption 3.3 constrains the product $p(1+q)$ rather than the decay exponent p alone, all three regimes remain reachable: an envelope with $p = 1/5$ combined with a response time with $q = 4$, for instance, satisfies $p(1+q) = 1$, falls in the third regime, and yields $\|\boldsymbol{u}_n\|_2^2 \in \mathcal{O}(n^{-1})$.

Two borderline examples: sinc and RQ-1/2. Let us instantiate the above discussion on the two temporal correlation functions of Figure 1 that sit exactly at the boundary of validity of Assumption 3.3. The sinc correlation function $k_T(\tau) = \sin(\tau/l_T)/(\tau/l_T)$ satisfies $|k_T(\tau)| \leq \min(1, l_T/\tau)$, while the Rational Quadratic correlation function $k_T(\tau) = \left(1 + \tau^2/(2\alpha l_T^2)\right)^{-\alpha}$ decays as $\Theta(\tau^{-2\alpha})$. Both therefore admit a polynomial envelope with exponent $p = 1$, the former exactly and the latter for $\alpha = 1/2$. Neither is integrable, so neither is dominated by an integrable envelope g_T. This is precisely why Assumption 3.3 is stated through the quantity G rather than through $\int_0^\infty g_T(\tau) d\tau$.

With $p = 1$, Assumption 3.3 holds as soon as $p(1+q) = 1 + q > 1$, that is, as soon as $q > 0$, and Proposition B.1 applies because $2p(1+q) = 2(1+q) > 1$. The resulting rate is $\|\boldsymbol{u}_n\|_2^2 \in \mathcal{O}(n^{-2q})$, so both correlation functions admit a finite n^* as soon as the response time grows polynomially with the dataset size. For the response time $R(n) \in \Theta(n^3)$ used in the left panel of Figure 1, this gives $\|\boldsymbol{u}_n\|_2^2 \in \mathcal{O}(n^{-6})$, which is consistent with the finite maxima plotted there. More generally, for a RQ correlation function with parameter α, Assumption 3.3 requires $\alpha > 1/(2(1+q))$ while Proposition B.1 requires $\alpha > 1/(4(1+q))$. For $q = 3$, these thresholds are $\alpha > 1/8$ and $\alpha > 1/16$ respectively, so that $\alpha = 1/2$ lies comfortably inside both.

The only case in which these two correlation functions fall outside Assumption 3.3 is $q = 0$, that is, a response time that does not grow with the dataset size. Then $\int_0^{mR(m)} g_T(\tau) d\tau$ grows

logarithmically in m while $R(m)$ remains constant, so that G diverges. However, in that setting, $\|\boldsymbol{u}_n\|_2^2$ is nondecreasing and Recommendation 1 returns $n^* = T$ irrespective of the temporal correlation function. Capping the dataset size is not beneficial in that regime, so nothing is lost by excluding it.

C Recommended Stale Data Management Policy

In this appendix, we prove Theorem 3.8 by connecting the L^2-distance between Lipschitz-continuous acquisition functions and the integrated 2-Wasserstein distance between GP posteriors.

Proof. Let us start by considering the effect of Assumption 3.7 on the distance at a single point $(\boldsymbol{x}, t) \in \mathcal{S} \times \mathcal{T}$.

$$\begin{aligned} |\alpha_{\mathcal{D}}(\boldsymbol{x}, t) - \alpha_{\tilde{\mathcal{D}}}(\boldsymbol{x}, t)| &\leq L_T || \left(\mu_{\mathcal{D}}(\boldsymbol{x}, t) - \mu_{\tilde{\mathcal{D}}}(\boldsymbol{x}, t), \sigma_{\mathcal{D}}(\boldsymbol{x}, t) - \sigma_{\tilde{\mathcal{D}}}(\boldsymbol{x}, t)\right) ||_2 \\ &= L_T \sqrt{\left(\mu_{\mathcal{D}}(\boldsymbol{x}, t) - \mu_{\tilde{\mathcal{D}}}(\boldsymbol{x}, t)\right)^2 + \left(\sigma_{\mathcal{D}}(\boldsymbol{x}, t) - \sigma_{\tilde{\mathcal{D}}}(\boldsymbol{x}, t)\right)^2}. \end{aligned} \tag{39}$$

We can now bound the squared L^2 distance between the two acquisition functions $\alpha_{\mathcal{D}}$ and $\alpha_{\tilde{\mathcal{D}}}$ on any subset $\mathcal{S}' \times \mathcal{T}'$ of the spatio-temporal domain $\mathcal{S} \times \mathcal{T}$:

$$\begin{aligned} ||\alpha_{\mathcal{D}} - \alpha_{\tilde{\mathcal{D}}}||_2^2 &= \oint_{\mathcal{S}'} \int_{\mathcal{T}'} \left(\alpha_{\mathcal{D}}(\boldsymbol{x}, t) - \alpha_{\tilde{\mathcal{D}}}(\boldsymbol{x}, t)\right)^2 d\boldsymbol{x} dt \\ &\leq \oint_{\mathcal{S}'} \int_{\mathcal{T}'} L_T^2 \left(\left(\mu_{\mathcal{D}}(\boldsymbol{x}, t) - \mu_{\tilde{\mathcal{D}}}(\boldsymbol{x}, t)\right)^2 + \left(\sigma_{\mathcal{D}}(\boldsymbol{x}, t) - \sigma_{\tilde{\mathcal{D}}}(\boldsymbol{x}, t)\right)^2\right) d\boldsymbol{x} dt \qquad (40) \\ &= L_T^2 \oint_{\mathcal{S}'} \int_{\mathcal{T}'} \left(\left(\mu_{\mathcal{D}}(\boldsymbol{x}, t) - \mu_{\tilde{\mathcal{D}}}(\boldsymbol{x}, t)\right)^2 + \left(\sigma_{\mathcal{D}}(\boldsymbol{x}, t) - \sigma_{\tilde{\mathcal{D}}}(\boldsymbol{x}, t)\right)^2\right) d\boldsymbol{x} dt \\ &= L_T^2 W_2^2(\mathcal{GP}_{\mathcal{D}}, \mathcal{GP}_{\tilde{\mathcal{D}}}), \qquad (41) \end{aligned}$$

where (40) follows from (39) and (41) from the definition of the integrated 2-Wasserstein distance on the domain $\mathcal{S}' \times \mathcal{T}'$. Applying a square-root on both sides yields the desired result. □

D Numerical Results

Here, we provide a detailed description of each implemented benchmark and the associated figures. There are two figures associated with each benchmark, showing their average regrets and the size of their datasets throughout the experiment.

In the following, the synthetic benchmarks will be described as functions of a point $\boldsymbol{z}$ in the $(d+1)$-dimensional spatio-temporal domain $\mathcal{S} \times \mathcal{T}$. More precisely, the point $\boldsymbol{z}$ is explicitly given by $\boldsymbol{z} = (x_1, \cdots, x_d, t)$. Also, we will write $d' = d + 1$ for the sake of brevity.

Each benchmark comes with a time horizon H. To make the benchmarks noisy, each call to the objective function is perturbed with a centered Gaussian noise of variance σ_0^2, equal to 1% of the signal variance. Moreover, to control the expensiveness of the benchmarks, each call to the objective function takes $R(0)$ seconds to complete. The values of H, σ_0^2 and $R(0)$ are unknown to the evaluated TVBO algorithms but are provided in Table 2 for the sake of completeness.

The average regrets and their standard errors are reported for each method and each benchmark in Table 3.

Schwefel. The Schwefel function is d'-dimensional, and has the form

$$f(\boldsymbol{z}) = 418.9829 d' - \sum_{i=1}^{d'} z_i \sin\left(\sqrt{|z_i|}\right).$$

For the numerical evaluation, we set $d' = 4$ and we optimized the function on the domain $[-500, 500]^{d'}$. The results are provided in Figure 4.

Table 2: Noise variance σ_0^2, cost of single call $R(0)$ in seconds, time horizon H in seconds and approximated temporal lengthscale l_T in minutes for each benchmark. The table also indicates if each benchmark can be viewed as quasi-static ($l_T \gg R(0)$) or if it can be viewed as near-constant ($R(0) \sim R(\lfloor H/R(0) \rfloor)$).

Benchmark $(d+1)$	Observ. Noise σ_0^2	Cost $R(0)$ (s)	Horizon H (s)	Lengthscale l_T (s)	Quasi-Static	Near-Constant
Shekel (4)	0.02	8.00	600	174	No	Yes
Hartmann (3)	0.05	8.00	600	164	No	Yes
Ackley (4)	0.05	0.05	600	216	Yes	No
Griewank (6)	0.30	0.05	600	55	No	No
Eggholder (2)	0.10	0.05	600	22	No	No
Schwefel (4)	0.25	0.05	600	35	No	No
Hartmann (6)	0.05	0.10	600	270	Yes	No
Powell (4)	2.50	0.01	600	552	Yes	No
Temperature (3)	0.16	0.01	1800	396	Yes	No
WLAN (7)	1.50	0.10	600	120	No	No

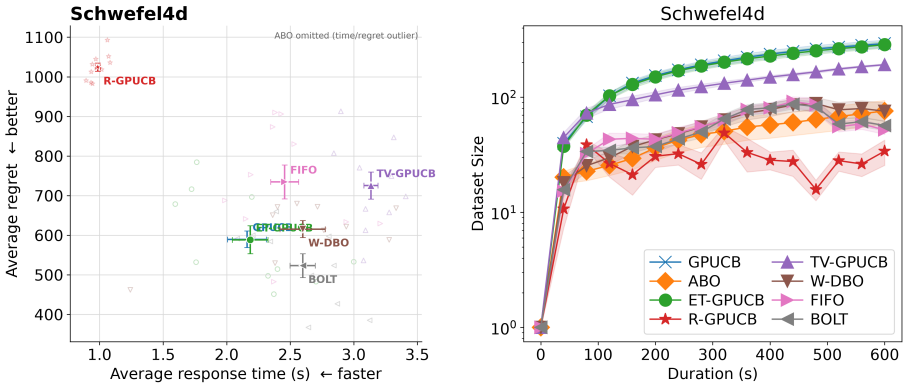


Figure 4: (Left) Average response time and average regrets of the TVBO solutions during the optimization of the Schwefel synthetic function. (Right) Dataset sizes of the TVBO solutions during the optimization of the Schwefel synthetic function.

Eggholder. The Eggholder function is 2-dimensional, and has the form

$$f(\boldsymbol{z}) = -(z_2 + 47)\sin\left(\sqrt{\left|z_2 + \frac{z_1}{2} + 47\right|}\right) - z_1 \sin\left(\sqrt{|z_1 - z_2 - 47|}\right).$$

For the numerical evaluation, we optimized the function on the domain $[-512, 512]^2$. The results are provided in Figure 5.

Ackley. The Ackley function is d'-dimensional, and has the form

$$f(\boldsymbol{z}) = -a\exp\left(-b\sqrt{\frac{1}{d'}\sum_{i=1}^{d'} z_i^2}\right) - \exp\left(\frac{1}{d'}\sum_{i=1}^{d'}\cos(cz_i)\right) + a + \exp(1).$$

For the numerical evaluation, we set $a = 20$, $b = 0.2$, $c = 2\pi$, $d' = 4$ and we optimized the function on the domain $[-32, 32]^{d'}$. The results are provided in Figure 6.

Table 3: Comparison of BOLT against state-of-the-art TVBO algorithms. For each experiment and each algorithm, the average regret and its standard error over 10 independent experiments is provided (lower is better). For each experiment, the performance of the best algorithm is written **in bold**, and the performance of algorithms whose confidence intervals overlap the confidence interval of the best performing algorithm are underlined. Due to space constraints, each instance of "GPUCB" in the methods' names has been replaced by "G", e.g., "TV-G" is "TV-GPUCB".

BENCH. $(d+1)$	G	ET-G	R-G	TV-G	ABO	W-DBO	FIFO	BOLT
SHEKEL (4)	2.57 ±0.05	2.57 ±0.04	2.58 ±0.02	2.62 ±0.05	2.67 ±0.11	2.64 ±0.04	2.55 ±0.04	**2.51 ±0.02**
HARTMANN (3)	1.58 ±0.08	1.43 ±0.12	0.70 ±0.05	0.77 ±0.10	0.91 ±0.12	0.63 ±0.04	**0.51 ±0.06**	0.60 ±0.04
ACKLEY (4)	4.26 ±0.50	4.39 ±0.45	4.28 ±0.20	3.83 ±0.57	4.75 ±0.53	3.42 ±0.46	3.74 ±0.45	**2.92 ±0.35**
GRIEWANK (6)	0.59 ±0.01	0.61 ±0.01	0.61 ±0.01	0.66 ±0.02	0.69 ±0.03	**0.57 ±0.03**	0.62 ±0.03	0.58 ±0.02
EGGHOLDER (2)	517 ±16.4	522 ±13.2	298 ±4.3	305 ±13.3	546 ±109.9	277 ±12.4	**225 ±6.8**	262 ±7.1
SCHWEFEL (4)	590 ±20.9	589 ±35.3	1024 ±10.3	726 ±34.1	792 ±61.3	615 ±21.5	735 ±43.0	**523 ±30.3**
HARTMANN (6)	1.67 ±0.01	1.75 ±0.07	1.44 ±0.02	0.96 ±0.24	1.32 ±0.17	**0.70 ±0.03**	1.20 ±0.14	0.72 ±0.06
POWELL (4)	3429 ±153	3283 ±240	1422 ±64	1301 ±155	6721 ±2113	1066 ±149	1070 ±202	**547 ±166**
TEMPERATURE (3)	1.12 ±0.05	1.29 ±0.11	**0.70 ±0.01**	1.25 ±0.10	1.41 ±0.26	0.98 ±0.06	0.83 ±0.09	0.72 ±0.03
WLAN (7)	20.2 ±1.58	21.2 ±1.05	10.4 ±0.16	7.8 ±0.34	17.3 ±1.93	8.9 ±0.20	7.9 ±0.17	**7.7 ±0.15**
OVERALL	0.62 ±0.10	0.67 ±0.10	0.39 ±0.10	0.40 ±0.08	0.82 ±0.07	0.21 ±0.07	0.23 ±0.06	**0.04 ±0.02**

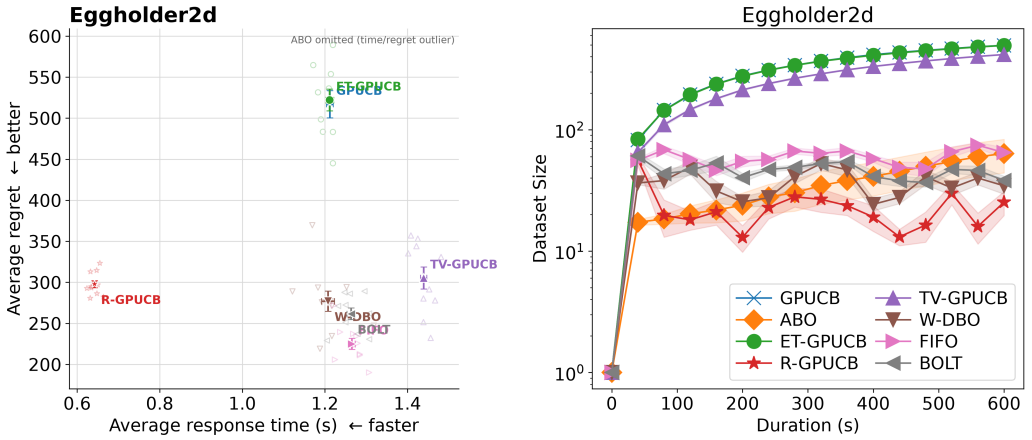


Figure 5: (Left) Average response time and average regrets of the TVBO solutions during the optimization of the Eggholder synthetic function. (Right) Dataset sizes of the TVBO solutions during the optimization of the Eggholder synthetic function.

Shekel. The Shekel function is 4-dimensional, and has the form

$$f(\boldsymbol{z}) = -\sum_{i=1}^{m}\left(\sum_{j=1}^{4}(z_j - C_{ji})^2 + \beta_i\right)^{-1}.$$

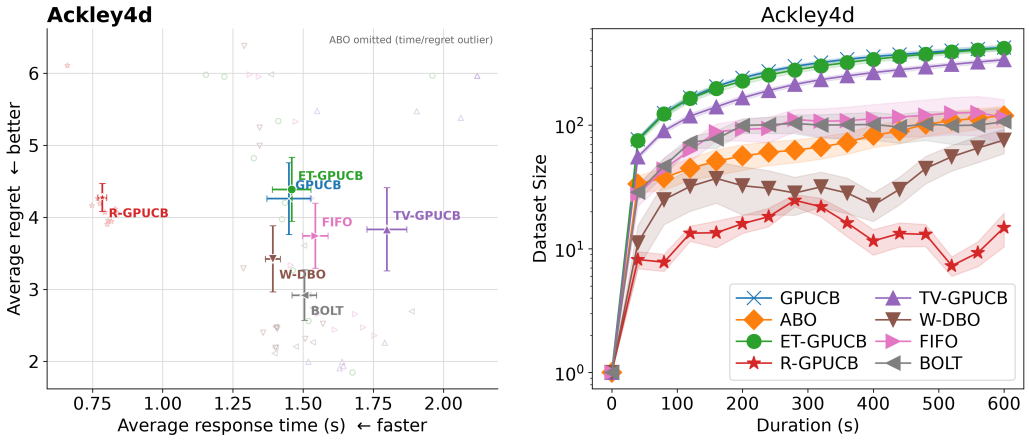


Figure 6: (Left) Average response time and average regrets of the TVBO solutions during the optimization of the Ackley synthetic function. (Right) Dataset sizes of the TVBO solutions during the optimization of the Ackley synthetic function.

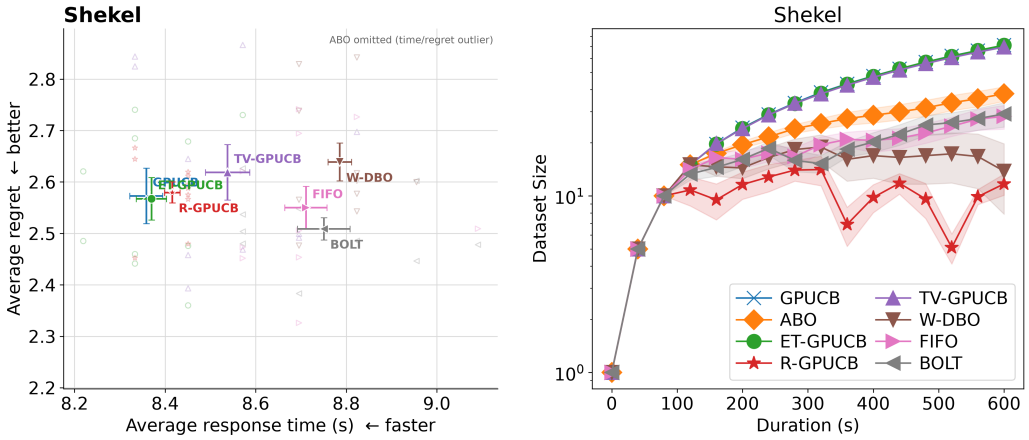


Figure 7: (Left) Average response time and average regrets of the TVBO solutions during the optimization of the Shekel synthetic function. (Right) Dataset sizes of the TVBO solutions during the optimization of the Shekel synthetic function.

For the numerical evaluation, we set $m = 10$, $\boldsymbol{\beta} = \frac{1}{10}(1, 2, 2, 4, 4, 6, 3, 7, 5, 5)$,

$$\boldsymbol{C} = \begin{pmatrix} 4 & 1 & 8 & 6 & 3 & 2 & 5 & 8 & 6 & 7 \\ 4 & 1 & 8 & 6 & 7 & 9 & 3 & 1 & 2 & 3.6 \\ 4 & 1 & 8 & 6 & 3 & 2 & 5 & 8 & 6 & 7 \\ 4 & 1 & 8 & 6 & 7 & 9 & 3 & 1 & 2 & 3.6 \end{pmatrix},$$

and we optimized the function on the domain $[0, 10]^4$. The results are provided in Figure 7.

Griewank. The Griewank function is d'-dimensional, and has the form

$$f(\boldsymbol{z}) = \sum_{i=1}^{d'} \frac{x_i^2}{4000} - \prod_{i=1}^{d'} \cos\left(\frac{x_i}{\sqrt{i}}\right) + 1$$

For the numerical evaluation, we set $d' = 6$ and we optimized the function on the domain $[-600, 600]^{d'}$. The results are provided in Figure 8.

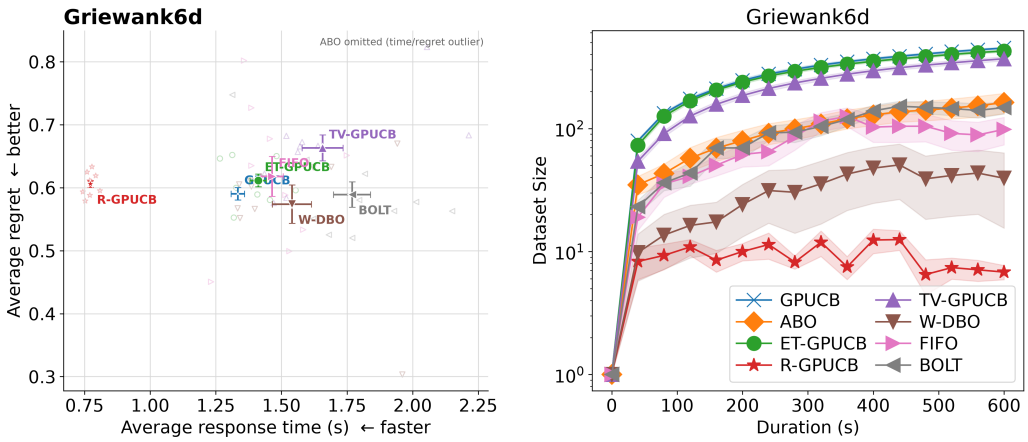


Figure 8: (Left) Average response time and average regrets of the TVBO solutions during the optimization of the Griewank synthetic function. (Right) Dataset sizes of the TVBO solutions during the optimization of the Griewank synthetic function.

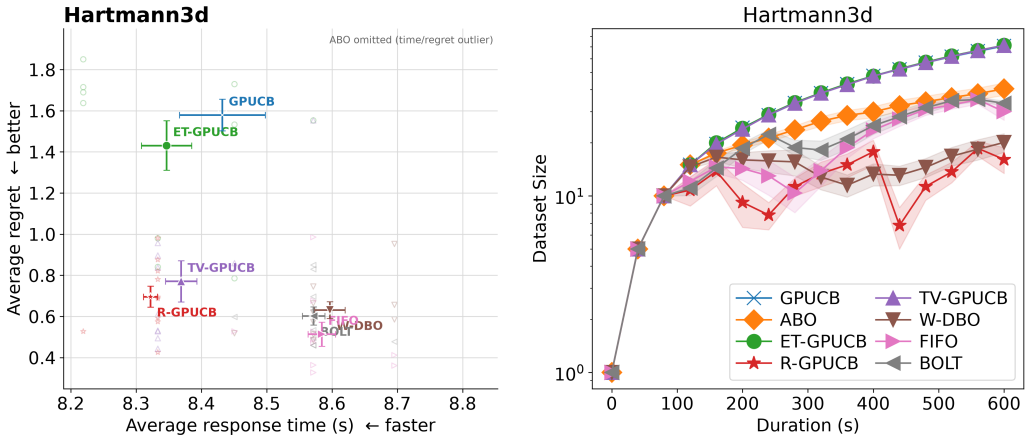


Figure 9: (Left) Average response time and average regrets of the TVBO solutions during the optimization of the Hartmann-3 synthetic function. (Right) Dataset sizes of the TVBO solutions during the optimization of the Hartmann-3 synthetic function.

Hartmann-3. The Hartmann-3 function is 3-dimensional, and has the form

$$f(\boldsymbol{z}) = -\sum_{i=1}^{4} \alpha_i \exp\left(-\sum_{j=1}^{3} A_{ij}(z_j - P_{ij})^2\right).$$

For the numerical evaluation, we set $\boldsymbol{\alpha} = (1.0, 1.2, 3.0, 3.2)$,

$$\boldsymbol{A} = \begin{pmatrix} 3 & 10 & 30 \\ 0.1 & 10 & 35 \\ 3 & 10 & 30 \\ 0.1 & 10 & 35 \end{pmatrix}, \boldsymbol{P} = 10^{-4} \begin{pmatrix} 3689 & 1170 & 2673 \\ 4699 & 4387 & 7470 \\ 1091 & 8732 & 5547 \\ 381 & 5743 & 8828 \end{pmatrix},$$

and we optimized the function on the domain $[0, 1]^3$. The results are provided in Figure 9.

Hartmann-6. The Hartmann-6 function is 6-dimensional, and has the form

$$f(\boldsymbol{z}) = -\sum_{i=1}^{4} \alpha_i \exp\left(-\sum_{j=1}^{6} A_{ij}(z_j - P_{ij})^2\right).$$

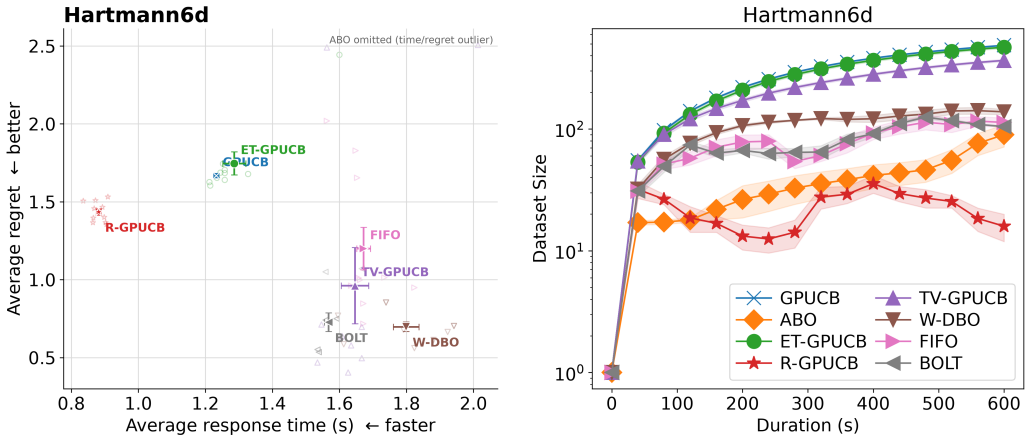


Figure 10: (Left) Average response time and average regrets of the TVBO solutions during the optimization of the Hartmann-6 synthetic function. (Right) Dataset sizes of the TVBO solutions during the optimization of the Hartmann-6 synthetic function.

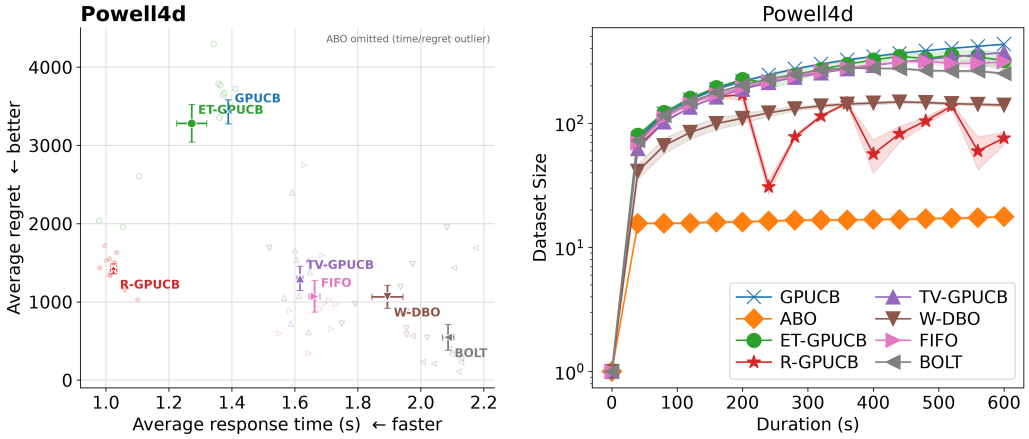


Figure 11: (Left) Average response time and average regrets of the TVBO solutions during the optimization of the Powell synthetic function. (Right) Dataset sizes of the TVBO solutions during the optimization of the Powell synthetic function.

For the numerical evaluation, we set $\boldsymbol{\alpha} = (1.0, 1.2, 3.0, 3.2)$,

$$\boldsymbol{A} = \begin{pmatrix} 10 & 3 & 17 & 3.50 & 1.7 & 8 \\ 0.05 & 10 & 17 & 0.1 & 8 & 14 \\ 3 & 3.5 & 1.7 & 10 & 17 & 8 \\ 17 & 8 & 0.05 & 10 & 0.1 & 14 \end{pmatrix}, \boldsymbol{P} = 10^{-4} \begin{pmatrix} 1312 & 1696 & 5569 & 124 & 8283 & 5886 \\ 2329 & 4135 & 8307 & 3736 & 1004 & 9991 \\ 2348 & 1451 & 3522 & 2883 & 3047 & 6650 \\ 4047 & 8828 & 8732 & 5743 & 1091 & 381 \end{pmatrix},$$

and we optimized the function on the domain $[0, 1]^6$. The results are provided in Figure 10.

Powell. The Powell function is d'-dimensional, and has the form

$$f(\boldsymbol{z}) = \sum_{i=1}^{d'/4} (z_{4i-3} + 10 z_{4i-2})^2 + 5(z_{4i-1} - z_{4i})^2 + (z_{4i-2} - 2z_{4i-1})^4 + 10(z_{4i-3} - z_{4i})^4.$$

For the numerical evaluation, we set $d' = 4$ and we optimized the function on the domain $[-4, 5]^{d'}$. The results are provided in Figure 11.

Temperature. This benchmark comes from the temperature dataset collected from 46 sensors deployed at Intel Research Berkeley. It is a famous benchmark, used in other works such as Bogunovic

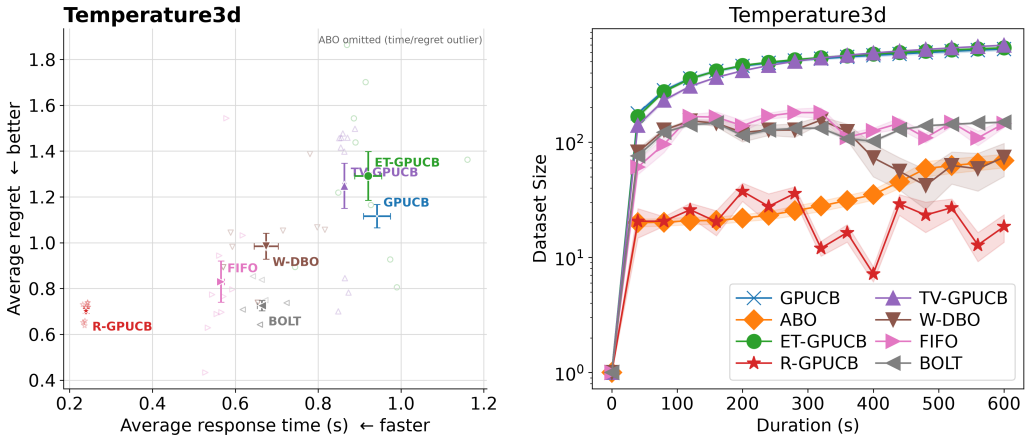


Figure 12: (Left) Average response time and average regrets of the TVBO solutions during the Temperature real-world experiment. (Right) Dataset sizes of the TVBO solutions during the Temperature real-world experiment.

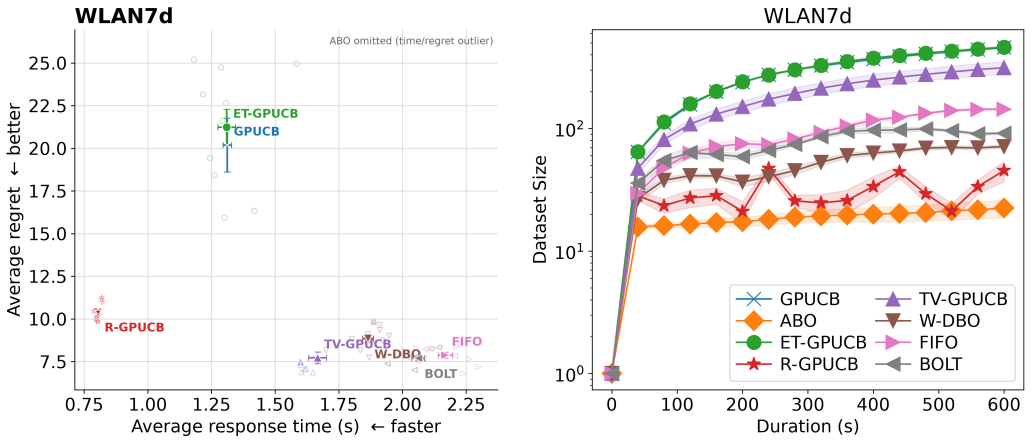


Figure 13: (Left) Average response time and average regrets of the TVBO solutions during the WLAN real-world experiment. (Right) Dataset sizes of the TVBO solutions during the WLAN real-world experiment.

et al. (2016); Brunzema et al. (2025). The goal of the TVBO task is to activate the sensor with the highest temperature, which will vary with time. To make the benchmark more interesting, we interpolate the data in space-time. With this interpolation, the algorithms can activate any point in space-time, making it a 3-dimensional benchmark (2 spatial dimensions for a location in Intel Research Berkeley, 1 temporal dimension).

For the numerical evaluation, we used the first day of data. The results are provided in Figure 12.

WLAN. This benchmark aims at maximizing the throughput of a Wireless Local Area Network (WLAN). 27 moving end-users are associated with one of 6 fixed nodes and continuously stream a large amount of data. As they move in space, they change the radio environment of the network, which should adapt accordingly to improve its performance. To do so, each node has a power level that can be tuned for the purpose of reaching the best trade-off between serving all its users and not causing interference for the neighboring nodes.

The performance of the network is computed as the sum of the Shannon capacities for each pair of node and associated end-users. The Shannon capacity (Kemperman, 1974) sets a theoretical upper bound on the throughput of a wireless communication. We denote it $C(i, j)$, we express it in bits per second (bps). It depends on S_{ij} the Signal-to-Interference plus Noise Ratio (SINR) of the

communication between node i and end-user j, as well as on W, the bandwidth of the radio channel (in Hz):

$$C_{ij}(\boldsymbol{x}, t) = W \log_2(1 + S_{ij}(\boldsymbol{x}, t)).$$

Then, the objective function is

$$f(\boldsymbol{x}, t) = \sum_{i=1}^{6} \sum_{j \in \mathcal{N}_i} C_{ij}(\boldsymbol{x}, t),$$

where $\mathcal{N}_i$ are the end-users associated with node i.

For the numerical evaluation, we optimized the power levels $\boldsymbol{x}$ in the domain $[10^{0.1}, 10^{2.5}]^6$. For this experiment, the TVBO solutions were evaluated with a Matérn-5/2 for the spatial covariance function and a Matérn-1/2 for the temporal covariance function. The results are provided in Figure 13.