BAYESIAN OPTIMIZATION WITH STRUCTURED MEASUREMENTS: A VECTOR-VALUED RKHS FRAMEWORK

Anonymous authors
Paper under double-blind review

ABSTRACT

Bayesian optimization (BO) is an efficient framework for optimizing expensive black-box functions. However, it is typically formulated as learning an end-to-end mapping from inputs to scalar objectives, thereby discarding the potentially rich information whenever a structured system output is available. In this work, we study Bayesian optimization over a vector-valued operator with structured measurements, where each measurement observes multidimensional or functional outputs, e.g., trajectories or spatial fields, rather than a single scalar value. The objective is then defined as a linear functional of these measurements. This allows each observation to reveal substantially richer information about the underlying system compared to scalar observations. Assuming the unknown operator lies in a vector-valued reproducing kernel Hilbert space (RKHS), we derive high-probability concentration bounds for the kernel ridge regression (KRR) estimator directly in the measurement space, characterizing uncertainty in a general Hilbert space. Building on these results, we propose an algorithm based on the upper confidence bound (UCB) acquisition function with regret guarantees under mild assumptions, recovering sublinear rates for common kernels. Empirically, we demonstrate that leveraging structured measurements leads to improved sample efficiency by enabling efficient transfer of information across objectives and adaptation to time-varying settings.

1 INTRODUCTION

Bayesian optimization (BO) is an efficient sequential strategy for optimizing expensive black-box functions Srinivas et al. (2012); Chowdhury & Gopalan (2017); Frazier (2018); Garnett (2023). It is commonly formulated as learning an end-to-end mapping from inputs to a scalar-valued objective using a Gaussian process surrogate model Chowdhury & Gopalan (2017); Srinivas et al. (2012). However, in many real-world problems, this scalar objective is not directly generated by the system, but instead arises as a functional of an underlying structured output, such as a trajectory, field, or time series. For example, in building operation, the cost is defined by integrating heating prices over realized heating power trajectories Shi et al. (2026); Wang et al. (2025). Standard BO treats the problem as learning a direct mapping from inputs to a scalar objective, effectively discarding the intermediate trajectory Yu et al. (2025); Xu et al. (2024); Menn et al. (2026); Hose et al. (2025); Brunzema et al. (2022). This limitation becomes particularly pronounced when multiple objectives correspond to different functionals of the same trajectory (e.g., different price signals), where existing methods must relearn each objective separately instead of sharing information through the common structure. While algorithms such as multi-task Bayesian optimization Chowdhury & Gopalan (2021); Evgeniou & Pontil (2004); Coutinho et al. (2026); Swersky et al. (2013) partially address this issue by modeling correlations across objectives, they typically rely on a finite number of scalar observations and a careful design of the correlation matrix. Another line of work models such correlations in the input space, such as contextual Bayesian optimization Krause & Ong (2011); Xu et al. (2024); Wang et al. (2025), which can be sample-inefficient for exploration across different contexts. When structured outputs are available for measurement, each observation can reveal more information about the underlying system Yuan & Cai (2010). This motivates leveraging such measurements in algorithm design, which can potentially improve learning efficiency.

To account for structured measurements, it is natural to consider Bayesian optimization with vector-valued outputs Kadri et al. (2016); Carmeli et al. (2010); Micchelli & Pontil (2005), where the black-box operator maps each input to outputs in a Hilbert space. This formulation enables modeling

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.

the underlying system through vector-valued observations. However, existing approaches remain limited. Most work focuses on finite-dimensional Euclidean output spaces Kudva & Paulson (2026); Chowdhury & Gopalan (2021); Swersky et al. (2013); Sessa et al. (2023), whereas many practical problems involve vector-valued trajectories and functions. Another line of work considers infinite-dimensional outputs but assumes that measurements are available only through linear functionals Brault et al. (2019); Huang et al. (2026); Lahr et al. (2026), effectively reducing each measurement to a scalar value. As a result, these methods fail to exploit the full structure of the underlying output and limit the amount of information obtained from each observation.

In this work, we study Bayesian optimization for black-box systems that produce structured, vector-valued outputs. We introduce a general framework in which the unknown operator lies in a vector-valued reproducing kernel Hilbert space (RKHS) and is observed through a known linear measurement operator, possibly mapping into a different Hilbert space. The objective is then defined as a linear functional of these measurements. This formulation unifies a broad class of existing problem settings Brault et al. (2019); Huang et al. (2026); Chowdhury & Gopalan (2021); Vien et al. (2018); Lahr et al. (2026); Kudva & Paulson (2026); Swersky et al. (2013); Sessa et al. (2023) and enables structured measurements to be incorporated in a principled manner. Under mild assumptions, we establish high-probability concentration bounds for the vector-valued kernel ridge regression (KRR) estimator under general measurement operators. From the derived uncertainty bound, we propose an algorithm based on the upper confidence bound (UCB) of the estimator and prove sublinear regret guarantees for commonly used kernels. Numerical experiments demonstrate that structured measurements significantly improve sample efficiency, permitting information transfer across objectives and adaptation in time-varying settings. Further discussion on related work can be found in Appendix A.

Contributions. The contributions of this work are as follows:

1. **Bayesian optimization with structured measurements.** We introduce a general formulation of Bayesian optimization in which measurements are obtained through linear measurement operators acting on a latent nonlinear vector-valued operator. This enables richer information to be incorporated at each observation and naturally provides the potential for improving learning efficiency.
2. **Theoretical guarantees under general measurement operators.** We establish high-probability concentration bounds for vector-valued KRR under bounded linear measurement operators in a general Hilbert space. Based on these results, we propose a UCB-type algorithm and prove sublinear regret guarantees, extending classical scalar-valued BO with structured observations.
3. **Improved efficiency in real-world applications.** Through synthetic benchmarks and real-world experiments, we demonstrate that the proposed method significantly improves sample efficiency by facilitating efficient information transfer across different objectives and achieving superior performance in optimizing time-varying objectives.

2 Problem formulation

Problem formulation. We consider the problem of optimizing a black-box system in which, for each input, the objective is defined as a linear functional of structured measurements of an underlying vector-valued operator. Specifically, we consider

$$\max_{x \in \mathcal{X}} \quad F(x) = \langle m, Mf(x) \rangle_{\mathcal{M}}, \tag{1}$$

where $f : \mathcal{X} \to \mathcal{Y}$ is an unknown vector-valued operator and where $\mathcal{Y}$ is the output Hilbert space. The operator $M : \mathcal{Y} \to \mathcal{M}$ is a known bounded linear measurement operator, where $\mathcal{M}$ is the Hilbert space of measurements, and $m \in \mathcal{M}$ defines a linear functional that evaluates the measured output. For each input x, the output $f(x)$ is not directly observable; instead, only its measurement $Mf(x)$ is available.

Practical implication. This formulation generalizes different observation regimes, ranging from full observation ($M = I$) to scalar observation (M is a linear functional). The inner product $\langle m, Mf(x) \rangle_{\mathcal{M}}$ naturally arises in many applications where the objective depends on vector-valued

outputs linearly, e.g., energy consumption computed from trajectories in building control Wang et al. (2025); Shi et al. (2026). It captures a broad class of common operations such as point evaluations, integrals, and averaging. Examples of common linear functionals in different Hilbert spaces are provided in Appendix B.

Comparison to scalar-valued BO. As illustrated in Fig. 1, scalar-valued BO models an end-to-end mapping from x to $F(x)$ using scalar feedback, thereby discarding the underlying structure in $f(x)$. This limits the ability to exploit richer observations and prevents information transfer across different objectives. In contrast, our formulation explicitly incorporates structured measurements $Mf(x)$, enabling more informative learning and adaptation across objectives.

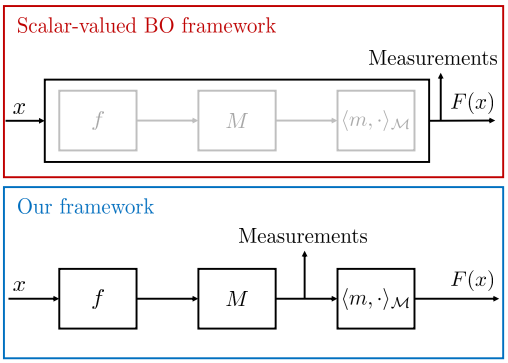


Figure 1: Vector-valued BO framework. Red box: classical BO framework. Blue box: vector-valued BO framework introduced in this work.

Comparison to BO with augmented input. Observe that for any m and M, $\langle m, Mf(x)\rangle_{\mathcal{M}} = \langle M^*m, f(x)\rangle_{\mathcal{Y}}$. Thus, the objective can be viewed as applying a linear functional $M^*m \in \mathcal{Y}$ to $f(x)$. A common approach is to model this quantity using a scalar-valued function $\tilde{F}(x, y) := \langle y, f(x)\rangle_{\mathcal{Y}}$ with an augmented input space $\mathcal{X} \times \mathcal{Y}$, which is learned via a scalar kernel Lahr et al. (2026); Huang et al. (2026). While this enables standard scalar-valued analysis, it reduces the observation to a scalar projection of the output. In contrast, our formulation provides flexibility. Different choices of M correspond to different observation mechanisms, which can significantly affect the amount of information obtained about $f(x)$.

Comparison to grey-box BO. Our formulation is closely related to grey-box Bayesian optimization, which exploits known structure in the objective through intermediate system outputs or known mappings between latent quantities and the final objective Astudillo & Frazier (2022); Xu et al. (2023); Astudillo & Frazier (2019). In this sense, our approach can also be viewed as a structured optimization framework in which the objective is evaluated from an observable intermediate representation. A key distinction is that we explicitly model structured measurements in a general Hilbert space, allowing the intermediate output to be vector- or function-valued and potentially infinite-dimensional.

Performance metric. Since f is a black-box operator, solving equation 1 in a single step is generally intractable. Instead, we aim to construct a sequence of inputs $\{x_t\}_{t=1}^T$ under a finite budget T such that the resulting averaged performance approximates the optimal solution of equation 1. Specifically, at each iteration t, the learner queries f at a chosen point x_t and observes noisy measurements $y_t = Mf(x_t) + \delta_t$, where $\delta_t \in \mathcal{M}$ is additive noise. The goal is to design a policy that balances exploration and exploitation, which can be quantified via the cumulative regret

$$\mathcal{R}_T = \sum_{t=1}^{T} \langle m, Mf(x^*) - Mf(x_t)\rangle_{\mathcal{M}},$$

where x^* denotes an optimal solution to equation 1. Sublinear growth of $\mathcal{R}_T$ with respect to T indicates the problem is effectively solved Srinivas et al. (2012); Chowdhury & Gopalan (2017). In general, such results are not attainable for arbitrary operators and noise processes. To obtain meaningful results, we impose the following assumptions.

Assumption 1 (Observation model)**.** *Let $R \in \mathcal{L}(\mathcal{M})$ be a self-adjoint, positive semi-definite operator. Conditioned on the natural filtration $\mathcal{F}_{t-1}$, $\delta_t \in \mathcal{M}$ is a zero-mean R-subGaussian random element in $\mathcal{M}$, i.e., $\forall \alpha \in \mathcal{M}$, $\mathbb{E}[e^{\langle \alpha, \delta_t\rangle_{\mathcal{M}}} | \mathcal{F}_{t-1}] \leq e^{\|\alpha\|_R^2/2}$.*

Remark 1. *The sub-Gaussian noise assumption is commonly adopted in the literature Chowdhury & Gopalan (2021); Huang et al. (2026). It ensures that the noise has controlled (light-tailed) behavior. Depending on the measurement space $\mathcal{M}$, the noise δ can take values in an infinite-dimensional Hilbert space. Since δ is $\mathcal{M}$-valued, its coordinates in any orthonormal basis are square-summable almost surely. The stronger sub-Gaussian tail bound further controls the probability of large deviations, enabling high-probability confidence bounds for sequential estimation.*

Vector-valued RKHS. A vector-valued RKHS is a direct generalization of the scalar-valued RKHS Micchelli & Pontil (2005) and contains the operators with outputs in $\mathcal{Y}$. We adopt its definition from Micchelli & Pontil (2005).

Definition 1. *The Hilbert space $\mathcal{H}$ is a reproducing kernel Hilbert space if for any $y \in \mathcal{Y}$ and $x \in \mathcal{X}$, the linear functional that maps $f \in \mathcal{H}$ to $\langle y, f(x)\rangle_{\mathcal{Y}}$ is continuous.*

According to the Riesz representation theorem Akhiezer (1981), we can define a mapping $\Phi : \mathcal{X} \to \mathcal{L}(\mathcal{Y}, \mathcal{H})$, where $\mathcal{L}(\mathcal{Y}, \mathcal{H})$ is the Banach space of bounded linear operators, such that for all $y \in \mathcal{Y}$, $\langle y, f(x)\rangle_{\mathcal{Y}} = \langle \Phi(x)y, f\rangle_{\mathcal{H}}$. Consequently, a mapping $K : \mathcal{X} \times \mathcal{X} \to \mathcal{L}(\mathcal{Y})$ defined by $K(x, s)y = (\Phi(s)y)(x)$ is called a vector-valued kernel. Further properties of such a kernel can be found in Caponnetto et al. (2008); Carmeli et al. (2010).

Assumption 2 (Regularity)**.** *The unknown operator f belongs to the vector-valued RKHS $\mathcal{H}$, induced by the operator-valued kernel K. The RKHS norm of f is bounded as $\|f\|_{\mathcal{H}} \leq \Gamma$. Furthermore, we assume that the kernel is uniformly bounded, i.e.,*

$$\sup_{x \in \mathcal{X}} \|K(x, x)\|_{\text{op}} \leq C_K.$$

The boundedness of the kernel and the RKHS norm ensure that the target operators lie in a well-structured hypothesis space, enabling efficient uncertainty reduction through sequential sampling.

Frequentist perspective. In this work, we adopt a frequentist perspective and assume that the unknown operator is fixed with bounded norm. The only source of randomness arises from the observation noise. Accordingly, we analyze worst-case regret over all possible operators, whereas the Bayesian formulation considers expected regret under a prior distribution.

3 Confidence bound in measurement space

Vector-valued KRR. Let $X_t := \{x_i\}_{i=1}^t$ be a sequence of inputs and $Y_t := \text{col}(y_1, \ldots, y_t) \in \mathcal{M}^t$ the corresponding noisy observations, where $y_i = Mf(x_i) + \delta_i$, and $\mathcal{M}^t$ is the direct sum of t copies of $\mathcal{M}$. An estimator $\mu_t \in \mathcal{H}$ of f is obtained by solving the regularized risk minimization problem

$$\mu_t = \arg\min_{\tilde{f} \in \mathcal{H}} \sum_{i=1}^{t} \|M\tilde{f}(x_i) - y_i\|_{\mathcal{M}}^2 + \lambda \|\tilde{f}\|_{\mathcal{H}}^2, \tag{2}$$

where $\lambda > 0$ is a regularization parameter controlling the complexity of the estimator. The explicit formulation of μ_t is provided in Appendix C. To analyze the prediction error $\|Mf(x) - M\mu_t(x)\|_{\mathcal{M}}$, a natural approach is to first bound the error in the output space $\mathcal{Y}$ and then propagate it through the measurement operator M. However, this can be overly conservative, as it requires controlling errors in all directions of $\mathcal{Y}$, including those that are not observable through M. This observation motivates modeling the measurable quantity $Mf(x)$ directly in the measurement space.

Induced kernel in the measurement space. We introduce the induced kernel $K^M : \mathcal{X} \times \mathcal{X} \to \mathcal{L}(\mathcal{M})$ defined by

$$K^M(x, s) = MK(x, s)M^*,$$

which is an operator-valued kernel taking values in $\mathcal{L}(\mathcal{M})$. The following lemma shows that K^M is a reproducing kernel in the measurement space.

Lemma 1 (Proposition 6, Carmeli et al. (2010))**.** *Let $M \in \mathcal{L}(\mathcal{Y}, \mathcal{M})$ and $K^M(x, s) = MK(x, s)M^*$. Then K^M is a reproducing kernel, and the corresponding RKHS $\mathcal{H}_M$ is related to $\mathcal{H}$ through a partial isometry $W \in \mathcal{L}(\mathcal{H}, \mathcal{H}_M)$ such that for all $x \in \mathcal{X}$ and $f \in \mathcal{H}$, $(Wf)(x) = Mf(x)$.*

Lemma 1 shows that composing elements of a vector-valued RKHS with a bounded linear operator yields a new vector-valued RKHS with kernel K^M. We further assume that this kernel is bounded above by C_{K^M}, i.e., $\sup_x \|K^M(x, x)\|_{\text{op}} \leq C_{K^M}$. Consequently, the induced estimator in the measurement space can equivalently be obtained by performing KRR in $\mathcal{H}_M$ using the observations

Y_t Kadri et al. (2016); Micchelli & Pontil (2005). Specifically, for any $x \in \mathcal{X}$, the estimator of $Mf(x)$ in $\mathcal{M}$ is given by

$$K^M_{xX_t}\big(K^M_{X_tX_t} + \lambda I\big)^{-1} Y_t, \tag{3}$$

where $K^M_{xX_t} \in \mathcal{L}(\mathcal{M}^t, \mathcal{M})$ is defined by $K^M_{xX_t} v = \sum_{i=1}^t K^M(x, x_i)v_i, \forall v = (v_1, \ldots, v_t) \in \mathcal{M}^t$, and $K^M_{X_tX_t} \in \mathcal{L}(\mathcal{M}^t, \mathcal{M}^t)$ is the block operator matrix with entries $(K^M_{X_tX_t})_{ij} = K^M(x_i, x_j)$. Importantly, equation 3 coincides with the measured output of the solution to equation 2, i.e., $M\mu_t(x)$ (see Appendix C). Thus, working in the measurement space does not change the estimator for the observable quantity $Mf(x)$, but provides a straightforward formulation. In particular, it enables uncertainty quantification directly in $\mathcal{M}$, avoiding the conservatism introduced by bounding errors in the full output space $\mathcal{Y}$.

Concentration bound. In this work, we consider one of the commonly used kernels, namely the uniform-block-diagonal (UBD) kernel. Here, we adapt its definition from Massiani et al. (2025).

Definition 2 (Massiani et al. (2025))**.** *The operator-valued kernel K^M is UBD if there exist separable Hilbert spaces $\tilde{\mathcal{M}}$ and $\mathcal{V}$, an isometry $\iota_{\mathcal{M}} \in \mathcal{L}(\tilde{\mathcal{M}} \otimes \mathcal{V}, \mathcal{M})$ and a trace-class kernel $K^M_{\mathrm{tr}} : \mathcal{X} \times \mathcal{X} \to \mathcal{L}(\tilde{\mathcal{M}})$ such that $K^M = \iota_{\mathcal{M}}(K^M_{\mathrm{tr}} \otimes I_{\mathcal{V}})\iota^{-1}_{\mathcal{M}}$.*

Remark 2 (UBD kernel)**.** *The UBD kernel was first introduced in Massiani et al. (2025) for analyzing kernel conditional tests. Assuming that K^M is UBD essentially assumes that $\mathcal{H}$ can be factorized as the tensor product of two subspaces, $\mathcal{H}_M \simeq \mathcal{H}_{K^M_{\mathrm{tr}}} \otimes \mathcal{V}$. The trace-class assumption on K^M_{tr} is weaker than assuming that K^M is a trace-class operator. The latter does not cover the case where K^M is composed of a scalar-valued kernel and the identity operator, which can make quantities such as the information gain unbounded. The UBD formulation allows the relevant quantities to remain bounded, thereby enabling proper regret analysis.*

To quantify the uncertainty of the estimator, we define

$$K^M_t(x, x) = K^M(x, x) - K^M_{xX_t}\big(K^M_{X_tX_t} + \lambda I\big)^{-1} K^M_{X_tx},$$

where $K^M_{X_tx}$ is the adjoint of $K^M_{xX_t}$. This operator can be interpreted as the posterior variance, which captures the remaining uncertainty at x after observing Y_t. Using this characterization, we can derive high-probability confidence bounds for the prediction error in the measurement space.

Theorem 1. *Let Assumptions 1-2 hold. Assume that K^M is a UBD kernel with trace-class block kernel K^M_{tr} and isometry $\iota_{\mathcal{M}} \in \mathcal{L}(\tilde{M} \otimes \mathcal{V}, \mathcal{M})$, and there exists a $\sigma > 0$ such that $R = \sigma^2 \iota_{\mathcal{M}}(I_{\tilde{\mathcal{M}}} \otimes R_{\mathcal{V}})\iota^{-1}_{\mathcal{M}}$ where $R_{\mathcal{V}}$ is a trace-class operator. Given X_t and Y_t, with probability at least $1 - \zeta$, for any $x \in \mathcal{X}$, the prediction error is bounded by*

$$\|Mf(x) - M\mu_t(x)\|_{\mathcal{M}} \leq \left(\Gamma + \frac{\sigma\sqrt{\mathrm{Tr}(R_{\mathcal{V}})}}{\sqrt{\lambda}}\sqrt{2\log(1/\zeta) + \log\det\big(I + \lambda^{-1}K^M_{\mathrm{tr},X_tX_t}\big)}\right)\|K^M_t(x, x)\|^{1/2}_{\mathrm{op}}.$$

The proof of Theorem 1 can be found in Appendix D. The result builds on the confidence bounds established in Massiani et al. (2025), and adapts their analysis to our setting in which the kernel is composed with an additional measurement operator. The assumption that K^M is UBD is mild, as the factorization of K^M into a trace-class operator and an identity operator provides sufficient flexibility to model various kernels. For instance, both the case where K^M is trace class and the case where K^M is composed of a scalar-valued kernel and the identity operator are included in this framework. The bound consists of two components. The first term, $\log\det\big(I + \lambda^{-1}K^M_{\mathrm{tr},X_tX_t}\big)$, depending only on K^M_{tr}, resembles the mutual information in the scalar-valued setting. The second term, $\|K^M_t(x, x)\|^{1/2}_{\mathrm{op}}$, captures the uncertainty at the query point x. In practice, this quantity can be evaluated via the largest eigenvalue of $K^M_t(x, x)$. In Appendix I, we provide finite-dimensional reformulations to enable efficient numerical evaluation.

3.1 Special cases

Full observation ($M = I$). When $M = I$, the learner observes the full vector-valued output $f(x)$, and the induced kernel reduces to $K^M(x, s) = K(x, s)$. The objective becomes $\langle m, f(x)\rangle_{\mathcal{Y}}$ with

Algorithm 1 Vector-valued Bayesian optimization (vvBO)

Require: Total budget T, regularization parameter $\lambda > 0$, confidence level $\zeta \in (0, 1)$
1: Initialize dataset $X_0 = \emptyset$, $Y_0 = \emptyset$ with $M\mu_0(x) = 0$, $K_0^M(x, x) = K^M(x, x)$
2: **for** $t = 1, \ldots, T$ **do**
3: Compute $M\mu_{t-1}(x)$, $K_{t-1}^M(x, x)$
4: Choose the next sample via

$$x_t \in \arg\max_{x \in \mathcal{X}} \langle m, M\mu_{t-1}(x)\rangle_{\mathcal{M}} + \beta_{t-1}\|m\|_{\mathcal{M}}\|K_{t-1}^M(x, x)\|_{\text{op}}^{1/2}$$

5: Query f at x_t and observe $y_t = Mf(x_t) + \delta_t$
6: Update the dataset X_t, Y_t
7: **end for**

$m \in \mathcal{Y}$. This corresponds to the most informative measurements, where the entire output is observed. In this case, the confidence bound in Theorem 1 provides uncertainty quantification directly over the vector-valued output $f(x)$.

Scalar linear measurements ($M : \mathcal{Y} \to \mathbb{R}$). When M is a bounded linear functional, the observations become scalar-valued. In this case, $K^M(x, s)$ reduces to a scalar kernel, and the problem recovers standard scalar-valued Bayesian optimization. This regime corresponds to minimal information acquisition, where each query reveals only a one-dimensional projection of $f(x)$.

These cases show that our framework unifies both full vector-valued observation and scalar-valued observation as special instances of a general formulation. Intermediate choices of M interpolate between these two cases, providing a principled way to model the trade-off between observation richness and algorithm efficiency.

4 Algorithm and regret analysis

We construct upper and lower bounds for $F(x)$ at any iteration $t \geq 1$. In particular, given the dataset X_{t-1}, Y_{t-1}, $F(x) \in [\underline{F}_{t-1}(x), \overline{F}_{t-1}(x)]$, where

$$\overline{F}_{t-1}(x) = \langle m, M\mu_{t-1}(x)\rangle_{\mathcal{M}} + \beta_{t-1}\|m\|_{\mathcal{M}}\|K_{t-1}^M(x, x)\|_{\text{op}}^{1/2},$$
$$\underline{F}_{t-1}(x) = \langle m, M\mu_{t-1}(x)\rangle_{\mathcal{M}} - \beta_{t-1}\|m\|_{\mathcal{M}}\|K_{t-1}^M(x, x)\|_{\text{op}}^{1/2}.$$

To solve equation 1, we adopt the principle of optimism in the face of uncertainty, leading to a UCB-based algorithm as described in Algorithm 1.

Algorithm. We are now ready to present the **vector-valued Bayesian optimization (vvBO)** algorithm with structured measurements. We define the confidence parameter as $\beta_{t-1} = \Gamma + \frac{\sigma\sqrt{\text{Tr}(R_{\mathcal{V}})}}{\sqrt{\lambda}}\sqrt{2\log(1/\zeta) + \log\det\big(I + \lambda^{-1}K^M_{\text{tr},X_{t-1}X_{t-1}}\big)}$. At each iteration t, the learner constructs a KRR estimator in the measurement space using the dataset X_{t-1}, Y_{t-1}, and forms an upper confidence bound based on Theorem 1. The next query point is selected as $x_t = \arg\max_{x\in\mathcal{X}} \langle m, M\mu_{t-1}(x)\rangle_{\mathcal{M}} + \beta_{t-1}\|m\|_{\mathcal{M}}\|K_{t-1}^M(x, x)\|_{\text{op}}^{1/2}$, which balances exploration and exploitation through the confidence parameter β_{t-1}. The observation is then obtained as $y_t = Mf(x_t) + \delta_t$. Notably, the entire procedure operates directly in the measurement space, leveraging structured observations through the induced kernel K^M.

4.1 Regret analysis

The following theorem establishes a general regret bound for Algorithm 1.

Theorem 2. *Let Assumptions 1–2 and the assumptions in Theorem 1 hold. For any $\lambda > 0$ and $\zeta \in (0, 1)$, with probability at least $1 - \zeta$, the regret of Algorithm 1 satisfies*

$$\mathcal{R}_T \leq 2\|m\|_{\mathcal{M}}\left(\Gamma + \frac{\sigma\sqrt{\text{Tr}(R_{\mathcal{V}})}}{\sqrt{\lambda}}\sqrt{2\log(\frac{1}{\zeta}) + \log\det\big(I + \lambda^{-1}K^M_{\text{tr},X_TX_T}\big)}\right)\sqrt{T\sum_{t=1}^{T}\|K_{t-1}^M(x_t, x_t)\|_{\text{op}}}.$$

The proof of Theorem 2 is deferred to Appendix E. The resulting bound has a similar structure to classical results in scalar-valued Bayesian optimization Chowdhury & Gopalan (2017); Srinivas et al. (2012), as well as extensions to finite-dimensional vector-valued settings Chowdhury & Gopalan (2021). Despite this similarity, establishing sublinear regret in our setting is nontrivial. The term $\log\det(I + \lambda^{-1}K^M_{\mathrm{tr},X_TX_T})$ involves an operator-valued kernel and is well-defined only when K^M_{tr,X_TX_T} is trace-class. However, this condition alone does not characterize how the quantity grows with T. Likewise, boundedness of $\|K^M_{t-1}(x_t, x_t)\|_{\mathrm{op}}$ does not control the cumulative growth of the uncertainty terms. These challenges motivate restricting attention to structured kernel classes, under which explicit bounds can be derived.

Separable kernel. We restrict attention to a widely used class of separable kernels Carmeli et al. (2010). Specifically, we consider kernels of the form $K = G \otimes B$, where $G : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ is a scalar-valued kernel and $B \in \mathcal{L}(\mathcal{Y})$ is a positive self-adjoint operator. Under a measurement operator M, this induces $K^M(x, s) = G(x, s)\,(MBM^*) \in \mathcal{L}(\mathcal{M})$. We write $B^M = MBM^* \in \mathcal{L}(\mathcal{M})$ and

$$K^M_{X_tX_t} = G_{X_tX_t} \otimes B^M, \qquad K^M_{xX_t} = G_{xX_t} \otimes B^M, \qquad K^M_{X_tx} = G_{X_tx} \otimes B^M,$$

where $G_{X_tX_t} \in \mathbb{R}^{t\times t}$ is the kernel Gram matrix and $G_{xX_t} \in \mathbb{R}^{1\times t}$, $G_{X_tx} \in \mathbb{R}^{t\times 1}$. We assume that $\sup_x G(x, x) \le 1$ and $\|B^M\|_{\mathrm{op}} \le C^M$, w.l.o.g. We further assume that B^M has the decomposition $B^M = \iota_{\mathcal{M}}(B^M_0 \otimes I_{\mathcal{V}})\iota^{-1}_{\mathcal{M}}$ where $B^M_0 \in \mathcal{L}(\tilde{\mathcal{M}})$ is a self-adjoint positive operator with finite rank. Note that the finite-rank assumption on B^M_0 does not imply that B^M is finite rank. Accordingly, for any $s, t \in \mathcal{X}$, we can write $K^M_{\mathrm{tr}}(s, t) = G(s, t)B^M_0$.

Theorem 3 (Kernel-specific bound)**.** *Let Assumptions 1-2 and the assumptions in Theorem 1 hold. For a separable kernel $K^M = G \otimes B^M$ for which B^M has a UBD composition $B^M = \iota_{\mathcal{M}}(B^M_0 \otimes I_{\mathcal{V}})\iota^{-1}_{\mathcal{M}}$, if $\mathcal{X} \subset \mathbb{R}^d$ is compact, with probability at least $1 - \zeta$, the regret for Algorithm 1 satisfies:*

1. *Linear kernel: $\mathcal{R}_T \le \mathcal{O}\big(\log T\sqrt{T}\big)$,*
2. *Gaussian kernel: $\mathcal{R}_T \le \mathcal{O}\big((\log T)^{d+1}\sqrt{T}\big)$,*
3. *Matérn kernel: $\mathcal{R}_T \le \mathcal{O}\big(T^{d(d+1)/(2\nu+d(d+1))}\log T\sqrt{T}\big)$, for $\nu > \frac{d(d+1)}{2}$.*

The proof of Theorem 3 is deferred to Appendix F. The above results show that, with separable kernels, the proposed method achieves sublinear regret rates comparable to those in classical scalar-valued Bayesian optimization. In particular, the dependence on T matches known results despite the system outputs being vector-valued elements in Hilbert spaces. These guarantees follow from combining the high-probability confidence bounds in Theorem 1 with the information-gain bounds used in the proof of Theorem 3, demonstrating that efficient sequential optimization is achievable under general measurement operators.

Extension to time-varying and nonlinear functions. As discussed in Section 2, the proposed algorithm readily extends to time-varying objectives with time-dependent linear functionals m_t. Details are provided in Appendix H. Furthermore, the results in this work can also be naturally generalized to any nonlinear function $F' : \mathcal{M} \to \mathbb{R}$ satisfying a Lipschitz condition. A detailed discussion is deferred to Appendix H.

5 Numerical experiments

In this section, we demonstrate that leveraging structured measurements improves efficiency in optimizing black-box systems with vector-valued outputs via several numerical experiments, particularly with changing objectives. We compare the proposed vector-valued Bayesian optimization (vvBO) algorithm against multiple baselines, including vanilla BO Chowdhury & Gopalan (2017), multi-task BO (MTBO) Chowdhury & Gopalan (2021), function-on-function BO (FFBO) Huang et al. (2026), and contextual BO (CTBO), where the context is defined as M^*m in the augmented input space (see details in Appendix J.1). In all experiments, the underlying system produces structured outputs (e.g., trajectories or functions), and the objective is defined as a linear functional of these outputs. We consider both synthetic benchmarks and a real-world controller tuning task. All results are averaged over 10 independent runs with different starting points without pretraining. Simulation details, including baseline implementation, hyperparameters, are provided in Appendix J.1.

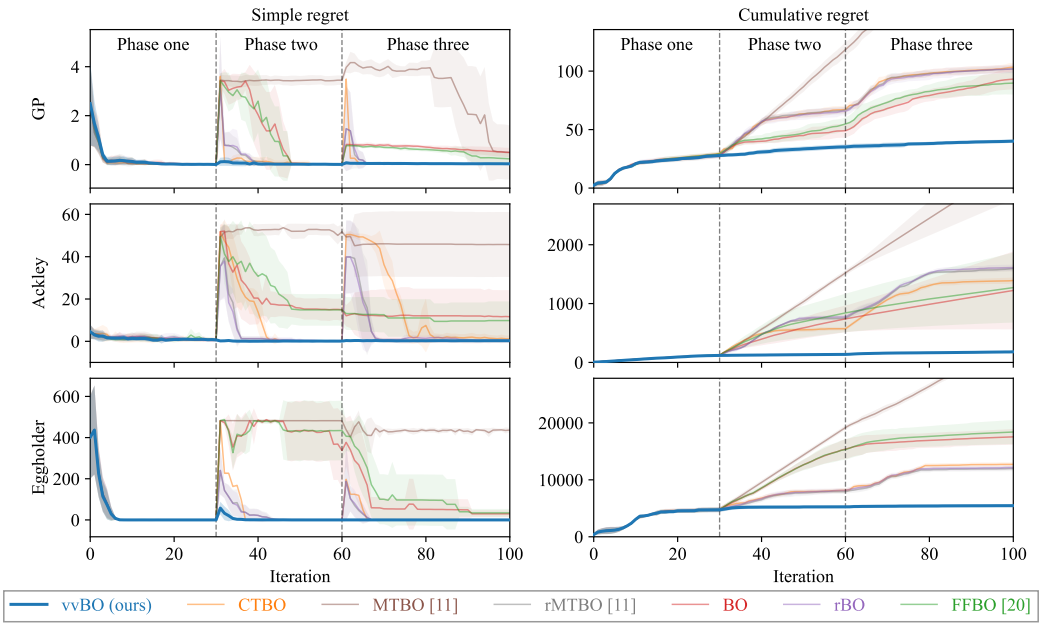


Figure 2: Comparison of simple regret and cumulative regret for three test operators with different baseline methods. vvBO with full trajectory measurements has the lowest regret in all scenarios, which indicates that it efficiently accommodates different objectives.

Synthetic problems. Here, we explore how structured measurements help to adapt to different objectives. We consider the setting where $M = I$ for vvBO, i.e., measuring the full vector-valued outputs. Additional results for partial observations ($M \neq I$) are provided in Appendix J.2. The optimization process is divided into three phases, each corresponding to a different objective defined by $m_i, i \in \{1, 2, 3\}$ (see detailed formulation in Appendix J.2). All the methods retain and reuse all previously collected samples across phases. Additionally, we compare against vanilla BO (rBO) and multi-task BO (rMTBO) re-initialized when shifting to the next phase. Performance is evaluated using both simple and cumulative regret. The results for three test functions are shown in Fig. 2. Across all experiments, vvBO with full vector-valued measurements consistently outperforms the baseline methods for different objectives. Since most baseline methods operate on scalar observations or fixed finite-dimensional representations, they cannot fully exploit the structure of the output or adjust effectively to changes in the objective. In contrast, our approach models the underlying vector-valued structure, facilitating effective reuse of information and rapid tracking. As for CTBO, while capable of handling varying objectives, it incurs larger cumulative regret since for each new context, it requires a large amount of exploration before establishing an accurate model. A detailed comparison between vvBO and CTBO is provided in Appendix J.2.

Time-varying objectives in real-world controller tuning. We next evaluate vvBO in a more realistic setting with time-varying objectives. We consider a real-world problem of tuning a Model Predictive Control (MPC) controller, which regulates indoor temperature in a commercial building testbed Yang et al. (2020). The simulated building has a surface area of $8500\,\mathrm{m}^2$ and serves approximately 1350 occupants, with heating provided by a district heating network. To control the temperature, the MPC controller regulates the valve position of a radiator and is parameterized by three variables. The goal is to choose the optimal parameters that minimize a weighted sum of energy consumption, CO_2 emissions, and thermal discomfort. This economic objective is commonly found in building control and captures the trade-off between energy efficiency and thermal comfort Chen et al. (2015). To create realistic conditions, we introduce a time-varying heating price following a sinusoidal pattern, resulting in a sequence of time-varying objectives (see detailed formulation in Appendix J.3). Controller parameters are updated daily, and temperature and heating trajectories are continuously measured each day. Experiments are conducted over 200 iterations with varying initial conditions.

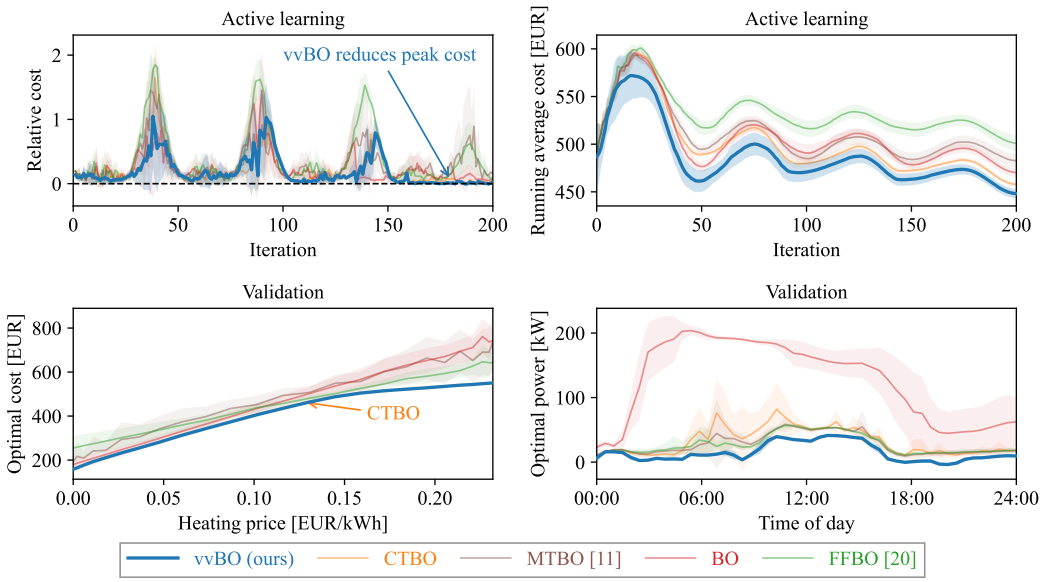


Figure 3: Real-world tuning of an MPC controller for building climate control.

The first row of Fig. 3 presents the cost during the active learning stage. After 200 iterations, vvBO achieves the lowest cost among all baseline methods. The impact of the time-varying heating price is reflected in the left plot. In the early phase, vvBO behaves similarly to baseline methods, focusing on exploration. As more data are collected, it adapts effectively to the varying prices by reducing the temperature when prices are high to save energy (peaks) and increasing the temperature when prices are low to improve comfort (valleys). Consequently, after around 150 iterations, vvBO significantly reduces the peak cost. Details on how the heating price affects the total cost can be found in Appendix J.3. This reveals the ability of vvBO to respond to time-varying objectives through rich measurements.

In the second row of Fig. 3, the validation of the learned model is illustrated. In the left plot, we vary the heating price and evaluate the economic cost associated with the parameter recommended by each method. vvBO achieves the lowest cost across all price levels, while other baseline methods incur high costs. In the right plot, we consider a new objective that is commonly found in building control Gelazanskas & Gamage (2014); Kosek et al. (2013), minimizing the peak heating power at specific times of the day. Since this objective is never encountered during the active learning stage, this experiment evaluates the ability to transfer information across different objectives. vvBO continues to recommend effective parameters across all time instances, whereas CTBO, which achieves comparable performance for the economic objective, performs poorly due to large uncertainty in unseen contexts. Indeed, it requires substantially more data to learn the new context. These results highlight that, by leveraging structured measurements, vvBO can generalize across objectives without modeling changing components, allowing adaptation to time-varying and previously unseen settings.

6 Conclusion

In this work, we study Bayesian optimization for systems that produce structured, vector-valued outputs observed through linear measurement operators. By modeling directly in the measurement space, we derive general concentration bounds for vector-valued KRR in the measurement space. Building on these results, we propose a UCB-based algorithm and establish sublinear regret guarantees for common kernels. This shows that effective learning is also possible when measurements lie in a general Hilbert space. Furthermore, our work highlights the importance of the observation model in Bayesian optimization: when structured measurements are available, each measurement can reveal substantially richer information about the underlying system than scalar observations. Leveraging this structure can facilitate more efficient learning and information sharing across multiple objectives. Numerical experiments validate these benefits in both synthetic and real-world scenarios.

AI USE STATEMENT

In this work, we used generative AI tools to assist with proofreading, including typo and grammatical error detection, and with polishing the wording of the manuscript. We also used generative AI tools to assist with the visualization of data obtained from our simulation results. All AI-assisted text was reviewed and revised by the authors, and all AI-assisted visualizations were checked against the original simulation outputs for correctness. We take responsibility for the final content of this work, including all text, claims, and artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

The code used to generate all plots in this work is provided in the supplementary materials. Complete proofs of all theoretical results, together with detailed derivations, are provided in the appendix. The appendix also contains the complete simulation details, including the ground-truth functions, hyperparameter settings, and the hardware used to conduct the experiments.

REFERENCES

Naum Ilich Akhiezer. *Theory of linear operators in Hilbert space*. Pitman Publishing, 1981.

Raul Astudillo and Peter Frazier. Bayesian optimization of composite functions. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 354–363. PMLR, 09–15 Jun 2019.

Raul Astudillo and Peter I. Frazier. Thinking inside the box: A tutorial on grey-box Bayesian optimization, 2022. URL https://arxiv.org/abs/2201.00272.

David Blum, Javier Arroyo, Sen Huang, Ján Drgoňa, Filip Jorissen, Harald Taxt Walnum, Yan Chen, Kyle Benne, Draguna Vrabie, Michael Wetter, et al. Building optimization testing framework (BOPTEST) for simulation-based benchmarking of control strategies in buildings. *Journal of Building Performance Simulation*, 14(5):586–610, 2021.

Edwin V Bonilla, Kian Chai, and Christopher Williams. Multi-task Gaussian process prediction. *Advances in Neural Information Processing Systems*, 20, 2007.

Romain Brault, Alex Lambert, Zoltán Szabó, Maxime Sangnier, and Florence d'Alché Buc. Infinite task learning in RKHSs. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1294–1302. PMLR, 2019.

Paul Brunzema, Alexander Von Rohr, and Sebastian Trimpe. On controller tuning with time-varying Bayesian optimization. In *2022 IEEE 61st Conference on Decision and Control*, pp. 4046–4052. IEEE, 2022.

Andrea Caponnetto, Charles A Micchelli, Massimiliano Pontil, and Yiming Ying. Universal multi-task kernels. *The Journal of Machine Learning Research*, 9:1615–1646, 2008.

Claudio Carmeli, Ernesto De Vito, and Alessandro Toigo. Vector valued reproducing kernel Hilbert spaces of integrable functions and Mercer theorem. *Analysis and Applications*, 4(04):377–408, 2006.

Claudio Carmeli, Ernesto De Vito, Alessandro Toigo, and Veronica Umanitá. Vector valued reproducing kernel Hilbert spaces and universality. *Analysis and Applications*, 8(01):19–61, 2010.

Xiao Chen, Qian Wang, and Jelena Srebric. Model predictive control for indoor thermal comfort and energy optimization using occupant feedback. *Energy and Buildings*, 102:357–369, 2015.

Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *International Conference on Machine Learning*, pp. 844–853. PMLR, 2017.

Sayak Ray Chowdhury and Aditya Gopalan. No-regret algorithms for multi-task Bayesian optimization. In *International Conference on Artificial Intelligence and Statistics*, pp. 1873–1881. PMLR, 2021.

João PL Coutinho, You Peng, Ricardo Rendall, Kaiwen Ma, Swee-Teng Chin, Ivan Castillo, and Marco S Reis. Efficient human-in-the-loop MPC tuning with multi-task preferential Bayesian optimization. *Control Engineering Practice*, 173:106972, 2026.

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.

Theodoros Evgeniou and Massimiliano Pontil. Regularized multi–task learning. In *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 109–117, 2004.

Peter I Frazier. A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.

Roman Garnett. *Bayesian optimization*. Cambridge University Press, 2023.

Linas Gelazanskas and Kelum A.A. Gamage. Demand side management in smart grid: A review and proposals for future direction. *Sustainable Cities and Society*, 11:22–30, 2014. ISSN 2210-6707.

Jan Gertheiss, David Rügamer, Bernard XW Liew, and Sonja Greven. Functional data analysis: An introduction and recent developments. *Biometrical Journal*, 66(7):e202300363, 2024.

Henrik Hose, Paul Brunzema, Alexander Von Rohr, Alexander Gräfe, Angela P Schoellig, and Sebastian Trimpe. Fine-tuning of neural network approximate MPC without retraining via Bayesian optimization. *arXiv preprint arXiv:2512.14350*, 2025.

Jingru Huang, Haijie Xu, Manrui Jiang, and Chen Zhang. Function-on-function Bayesian optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 21994–22002, 2026.

Hachem Kadri, Emmanuel Duflos, Manuel Davy, Philippe Preux, and Stephane Canu. *General framework for nonlinear functional regression with reproducing kernel Hilbert spaces*. PhD thesis, INRIA, 2009.

Hachem Kadri, Emmanuel Duflos, Philippe Preux, Stéphane Canu, and Manuel Davy. Nonlinear functional regression: a functional RKHS approach. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pp. 374–380. JMLR Workshop and Conference Proceedings, 2010.

Hachem Kadri, Emmanuel Duflos, Philippe Preux, Stéphane Canu, Alain Rakotomamonjy, and Julien Audiffren. Operator-valued kernels for learning from functional response data. *Journal of Machine Learning Research*, 17(20):1–54, 2016.

Nazan Khan, David E Goldberg, and Martin Pelikan. Multi-objective Bayesian optimization algorithm. In *Proceedings of the 4th Annual Conference on Genetic and Evolutionary Computation*, pp. 684–684, 2002.

Anna Magdalena Kosek, Giuseppe Tommaso Costanzo, Henrik W Bindner, and Oliver Gehrke. An overview of demand side management control schemes for buildings in smart grids. In *2013 IEEE International Conference on Smart Energy Grid Engineering*, pp. 1–9. IEEE, 2013.

Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to PDEs. *Journal of Machine Learning Research*, 24(89):1–97, 2023.

Andreas Krause and Cheng Ong. Contextual Gaussian process bandit optimization. *Advances in Neural Information Processing Systems*, 24, 2011.

Akshay Kudva and Joel A. Paulson. Bonsai: Structure-exploiting robust Bayesian optimization for networked black-box systems under uncertainty. *Computers & Chemical Engineering*, 204:109393, Jan 2026. ISSN 0098-1354.

Amon Lahr, Anna Scampicchio, Johannes Köhler, and Melanie N Zeilinger. Optimal uncertainty bounds for multivariate kernel regression under bounded noise: A Gaussian process-based dual function. *arXiv preprint arXiv:2603.16481*, 2026.

Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Graph kernel network for partial differential equations. *arXiv preprint arXiv:2003.03485*, 2020.

Pierre-François Massiani, Christian Fiedler, Lukas Haverbeck, Friedrich Solowjow, and Sebastian Trimpe. Kernel conditional tests from learning-theoretic bounds. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

Johanna Menn, David Stenger, and Sebastian Trimpe. Preferential Bayesian optimization with crash feedback. *IEEE Robotics and Automation Letters*, 2026.

Charles A Micchelli and Massimiliano Pontil. On learning vector-valued functions. *Neural computation*, 17(1):177–204, 2005.

Charles A Micchelli, Yuesheng Xu, and Haizhang Zhang. Universal kernels. *Journal of Machine Learning Research*, 7(12), 2006.

Ha Quang Minh. Infinite-dimensional log-determinant divergences between positive definite trace class operators. *Linear Algebra and Its Applications*, 528:331–383, 2017.

James O Ramsay and Bernard W Silverman. *Functional data analysis*. Springer, 1997.

Carlos Ruiz, Carlos M Alaíz, and José R Dorronsoro. A survey on kernel-based multi-task learning. *Neurocomputing*, 577:127255, 2024.

Pier Giuseppe Sessa, Pierre Laforgue, Nicolò Cesa-Bianchi, and Andreas Krause. Multitask learning with no regret: from improved confidence bounds to active learning. *Advances in Neural Information Processing Systems*, 36:6770–6781, 2023.

Jicheng Shi, Christophe Salzmann, and Colin N Jones. Disturbance-adaptive data-driven predictive control: Trading comfort violation for savings in building climate control. *IEEE Transactions on Control Systems Technology*, 2026.

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.

Kevin Swersky, Jasper Snoek, and Ryan P Adams. Multi-task Bayesian optimization. *Advances in Neural Information Processing Systems*, 26, 2013.

Ngo Anh Vien, Heiko Zimmermann, and Marc Toussaint. Bayesian functional optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

Wenbin Wang, Jicheng Shi, and Colin N Jones. Personalized building climate control with contextual preferential Bayesian optimization. *arXiv preprint arXiv:2512.09481*, 2025.

Wenjie Xu, Yuning Jiang, Bratislav Svetozarevic, and Colin N Jones. Bayesian optimization of expensive nested grey-box functions. *arXiv preprint arXiv:2306.05150*, 2023.

Wenjie Xu, Bratislav Svetozarevic, Loris Di Natale, Philipp Heer, and Colin N. Jones. Data-driven adaptive building thermal controller tuning with constraints: A primal–dual contextual Bayesian optimization approach. *Applied Energy*, 358:122493, 2024. ISSN 0306-2619. doi: https://doi.org/10.1016/j.apenergy.2023.122493. URL https://www.sciencedirect.com/science/article/pii/S0306261923018573.

Tao Yang, Konstantin Filonenko, Krzysztof Arendt, and Christain Veje. Implementation and performance analysis of a multi-energy building emulator. In *2020 6th IEEE International Energy Conference*, pp. 451–456. IEEE, 2020.

Jiarui Yu, Jicheng Shi, Wenjie Xu, and Colin N Jones. Which price to pay? Auto-tuning building MPC controller for optimal economic cost. *arXiv preprint arXiv:2501.10859*, 2025.

Ming Yuan and T. Tony Cai. A reproducing kernel Hilbert space approach to functional linear regression. *The Annals of Statistics*, 38(6), December 2010. ISSN 0090-5364. doi: 10.1214/09-aos772. URL http://dx.doi.org/10.1214/09-AOS772.

A Related work

Learning nonlinear vector-valued operators. Learning vector-valued or function-valued operators has been widely studied in the functional data analysis (FDA) literature Gertheiss et al. (2024); Ramsay & Silverman (1997). Classical approaches typically rely on discretization, representing the outputs in finite-dimensional Euclidean spaces Kadri et al. (2010; 2009); Li et al. (2020). Some recent work models functions directly in Hilbert spaces Brault et al. (2019); Kadri et al. (2016), using tools such as multi-output Gaussian processes Bonilla et al. (2007); Ruiz et al. (2024) and neural operators Kovachki et al. (2023). While these methods enable flexible modeling of vector-valued operators, they primarily focus on supervised learning and do not address sequential optimization problems. **BO with vector-valued outputs.** Several works extend Bayesian optimization to structured or multi-output settings. Multi-task Bayesian optimization Chowdhury & Gopalan (2021); Sessa et al. (2023); Swersky et al. (2013) leverages correlations across multiple outputs to improve sample efficiency, while multi-objective Bayesian optimization Khan et al. (2002); Daulton et al. (2022) focuses on learning Pareto-optimal trade-offs. However, these approaches typically assume a finite number of scalar outputs and do not consider measurements such as functions or trajectories. **BO with augmented inputs.** A related line of work assumes that observations are obtained through linear functionals of the vector-valued output Huang et al. (2026); Lahr et al. (2026). They view the linear functional as the augmented input of a scalar-valued function with an augmented input space, which reduces the problem to a scalar-valued setting. Consequently, this formulation cannot quantify the uncertainty of the vector-valued operator. A comprehensive analysis of uncertainty quantification under general linear measurement operators is still missing.

B Common linear functionals

In this section, we provide examples of linear functionals in common Hilbert spaces. In all cases, continuous linear functionals admit a representation via the inner product, as characterized by the Riesz representation theorem Akhiezer (1981).

Finite-dimensional spaces. In $\mathbb{R}^n$ with inner product

$$\langle x, y\rangle = \sum_{i=1}^{n} x_i y_i,$$

every linear functional φ can be written as

$$\varphi(x) = \sum_{i=1}^{n} a_i x_i = \langle x, y\rangle,$$

where $y = (a_1, \ldots, a_n) \in \mathbb{R}^n$.

Sequence space. The Hilbert space ℓ^2 consists of square-summable sequences with inner product

$$\langle x, y\rangle = \sum_{k=1}^{\infty} x_k y_k.$$

Every continuous linear functional φ admits the representation

$$\varphi(x) = \sum_{k=1}^{\infty} c_k x_k = \langle x, y\rangle,$$

for some $y = \{c_k\}_{k=1}^{\infty} \in \ell^2$.

Square-integrable function spaces. Let (Ω, μ) be a measure space. The space $L^2(\Omega, \mu)$ has inner product

$$\langle f, g\rangle = \int_{\Omega} f(x) g(x)\, d\mu(x).$$

By the Riesz representation theorem, every continuous linear functional φ can be expressed as

$$\varphi(f) = \int_{\Omega} f(x) g(x)\, d\mu(x) = \langle f, g\rangle,$$

for a unique $g \in L^2(\Omega, \mu)$.

Sobolev spaces. Sobolev spaces are Hilbert spaces equipped with inner products involving derivatives. For example, in $H^1(\Omega)$,

$$\langle u, v \rangle = \int_\Omega \nabla u \cdot \nabla v + uv \, dx.$$

Every continuous linear functional φ admits a representation

$$\varphi(u) = \langle u, h \rangle_{H^1},$$

for a unique $h \in H^1(\Omega)$.

Reproducing kernel Hilbert spaces. Let $\mathcal{H}$ be a scalar-valued RKHS with kernel $G : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$. By the reproducing property,

$$f(x) = \langle f, G(x, \cdot) \rangle_{\mathcal{H}}.$$

The evaluation functional $L(f) = f(x_0)$ is represented by $G(x_0, \cdot)$. More generally, for

$$L(f) = \sum_{i=1}^{n} a_i f(x_i),$$

the representer is

$$y = \sum_{i=1}^{n} a_i G(x_i, \cdot).$$

Integral functionals can also be expressed in this form. For example, let

$$L(f) = \int_{\mathcal{X}} g(x) f(x) \, dx,$$

where the integral is well-defined, and the resulting functional is continuous. Using the reproducing property,

$$L(f) = \int_{\mathcal{X}} g(x) \langle f, G(x, \cdot) \rangle_{\mathcal{H}} dx = \left\langle f, \int_{\mathcal{X}} g(x) G(x, \cdot) \, dx \right\rangle_{\mathcal{H}}.$$

Thus, the representer is

$$y = \int_{\mathcal{X}} g(x) G(x, \cdot) \, dx \in \mathcal{H}.$$

For two such functionals with weights g_1 and g_2, their inner product is

$$\langle y_1, y_2 \rangle_{\mathcal{H}} = \int_{\mathcal{X}} \int_{\mathcal{X}} g_1(x) g_2(y) G(x, y) \, dx \, dy.$$

More generally, for

$$L(f) = \sum_{i=1}^{n} a_i \int_{\mathcal{X}} g_i(x) f(x) \, dx,$$

the corresponding representer is

$$y = \sum_{i=1}^{n} a_i \int_{\mathcal{X}} g_i(x) G(x, \cdot) \, dx.$$

C Analytical formulation for the vector-valued estimator

In the following lemma, we provide the explicit formulation for equation 2.

Lemma 2. *Let Assumptions 1-2 hold. For any $x \in \mathcal{X}$, $\mu_t(x)$ is given by*

$$\mu_t(x) = \sum_{i=1}^{t} K(x, x_i) M^* \alpha_i,$$

where $\alpha_i \in \mathcal{M}$ solve the following linear system

$$\sum_{j=1}^{t}(MK(x_i,x_j)M^* + \lambda I_{i=j})\alpha_j = y_i,$$

where $I_{i=j}$ is the indicator operator, i.e., $I_{i=j} = I$ if $i = j$ and $I_{i=j} = 0$ otherwise.

Proof. For any $f \in \mathcal{H}$, we can express it as $f = \mu_t + g$, where $g \in \mathcal{H}$ represents the residual. Note that

$$\begin{aligned}
&\sum_{i=1}^{t} \|Mf(x_i) - y_i\|_{\mathcal{M}}^2 + \lambda\|f\|_{\mathcal{H}}^2 \\
&\qquad = \sum_{i=1}^{t} \|M\mu_t(x_i) + Mg(x_i) - y_i\|_{\mathcal{M}}^2 + \lambda\|\mu_t + g\|_{\mathcal{H}}^2 \\
&\qquad = \sum_{i=1}^{t} \|Mg(x_i)\|_{\mathcal{M}}^2 + \sum_{i=1}^{t} \|M\mu_t(x_i) - y_i\|_{\mathcal{M}}^2 - 2\sum_{i=1}^{t}\langle Mg(x_i), y_i - M\mu_t(x_i)\rangle_{\mathcal{M}} \\
&\qquad\quad + \lambda\|\mu_t\|_{\mathcal{H}}^2 + \lambda\|g\|_{\mathcal{H}}^2 + 2\lambda\langle\mu_t, g\rangle_{\mathcal{H}} \\
&\qquad = \sum_{i=1}^{t} \|M\mu_t(x_i) - y_i\|_{\mathcal{M}}^2 + \lambda\|\mu_t\|_{\mathcal{H}}^2 \\
&\qquad\quad + \sum_{i=1}^{t} \|Mg(x_i)\|_{\mathcal{M}}^2 + \lambda\|g\|_{\mathcal{H}}^2 \\
&\qquad\quad + 2\lambda\langle\mu_t, g\rangle_{\mathcal{H}} - 2\sum_{i=1}^{t}\langle Mg(x_i), y_i - M\mu_t(x_i)\rangle_{\mathcal{M}}.
\end{aligned}$$

For the last two terms, we have

$$\begin{aligned}
\lambda\langle\mu_t, g\rangle_{\mathcal{H}} &= \lambda\Big\langle \sum_{i=1}^{t} \Phi(x_i)M^*\alpha_i, g\Big\rangle_{\mathcal{H}} \\
&= \lambda\sum_{i=1}^{t}\langle\Phi(x_i)M^*\alpha_i, g\rangle_{\mathcal{H}} \\
&= \lambda\sum_{i=1}^{t}\langle M^*\alpha_i, g(x_i)\rangle_{\mathcal{Y}} \\
&= \lambda\sum_{i=1}^{t}\langle\alpha_i, Mg(x_i)\rangle_{\mathcal{M}},
\end{aligned}$$

$$\begin{aligned}
\sum_{i=1}^{t}\langle Mg(x_i), y_i - M\mu_t(x_i)\rangle_{\mathcal{M}} &= \sum_{i=1}^{t}\Big\langle Mg(x_i), \sum_{j=1}^{t}(MK(x_i,x_j)M^* + \lambda I_{i=j})\alpha_j - M\mu_t(x_i)\Big\rangle_{\mathcal{M}} \\
&= \lambda\sum_{i=1}^{t}\Big\langle Mg(x_i), \sum_{j=1}^{t} I_{i=j}\alpha_j\Big\rangle_{\mathcal{M}} \\
&= \lambda\sum_{i=1}^{t}\langle Mg(x_i), \alpha_i\rangle_{\mathcal{M}}.
\end{aligned}$$

Hence, the minimum is attained at $g = 0$. □

D PROOF OF THEOREM 1

This proof follows the same structure as Chowdhury & Gopalan (2021); Massiani et al. (2025), adapted to a general vector-valued setting. We highlight the necessary modifications when they arise. We employ the feature representation of an operator-valued kernel Micchelli & Pontil (2005), writing

$$K^M(x,s) = \Phi^*(x)\Phi(s),$$

where $\Phi(x) \in \mathcal{L}(\mathcal{M}, \mathcal{W})$ for some Hilbert space $\mathcal{W}$. Then there exists $w \in \mathcal{W}$ such that

$$(Mf)(\cdot) = \Phi^*(\cdot)w.$$

Define the operator $\Phi_{X_t} \in \mathcal{L}(\mathcal{W}, \mathcal{M}^t)$ by

$$(\Phi_{X_t} w)_i = \Phi^*(x_i)w, \quad \forall w \in \mathcal{W},$$

and denote its adjoint by $\Phi^*_{X_t}$. The vector-valued KRR estimator $\mu_t(x)$ can be expressed as

$$\begin{aligned}
M\mu_t(x) &= \Phi^*(x)\Phi^*_{X_t}(\Phi_{X_t}\Phi^*_{X_t} + \lambda I)^{-1} Y_t \\
&= \Phi^*(x)\Phi^*_{X_t}(\Phi_{X_t}\Phi^*_{X_t} + \lambda I)^{-1}(\Phi_{X_t} w + \boldsymbol{\delta}) \\
&= \Phi^*(x)(\Phi^*_{X_t}\Phi_{X_t} + \lambda I)^{-1}\Phi^*_{X_t}(\Phi_{X_t} w + \boldsymbol{\delta}) \\
&= \Phi^*(x)(\Phi^*_{X_t}\Phi_{X_t} + \lambda I)^{-1}(\Phi^*_{X_t}\Phi_{X_t} w + \Phi^*_{X_t}\boldsymbol{\delta}) \\
&= Mf(x) - \Phi^*(x)(\Phi^*_{X_t}\Phi_{X_t} + \lambda I)^{-1}(\lambda w - \Phi^*_{X_t}\boldsymbol{\delta}),
\end{aligned}$$

where $\boldsymbol{\delta} = (\delta_1, \cdots, \delta_t) \in \mathcal{M}^t$.

Consequently,

$$\begin{aligned}
\|M\mu_t(x) - Mf(x)\|_{\mathcal{M}} &= \left\|\Phi^*(x)(\Phi^*_{X_t}\Phi_{X_t} + \lambda I)^{-1}(\lambda w - \Phi^*_{X_t}\boldsymbol{\delta})\right\|_{\mathcal{M}} \\
&\leq \|(\Phi^*_{X_t}\Phi_{X_t} + \lambda I)^{-1/2}\Phi(x)\|_{\mathrm{op}} \\
&\quad \cdot \|(\Phi^*_{X_t}\Phi_{X_t} + \lambda I)^{-1/2}(\lambda w - \Phi^*_{X_t}\boldsymbol{\delta})\|_{\mathcal{W}} \\
&\leq \|\Phi^*(x)(\Phi^*_{X_t}\Phi_{X_t} + \lambda I)^{-1}\Phi(x)\|^{1/2}_{\mathrm{op}} \\
&\quad \cdot \Big(\|\Phi^*_{X_t}\boldsymbol{\delta}\|_{(\Phi^*_{X_t}\Phi_{X_t}+\lambda I)^{-1}} + \lambda^{1/2}\|w\|_{\mathcal{W}}\Big),
\end{aligned}$$

where we used $\|A\|_{\mathrm{op}} \leq \|A^*A\|^{1/2}_{\mathrm{op}}$ and the bound

$$\|w\|_{(\Phi^*_{X_t}\Phi_{X_t}+\lambda I)^{-1}} \leq \lambda^{-1/2}\|w\|_{\mathcal{W}}.$$

Next, observe that

$$\begin{aligned}
\lambda\Phi^*(x)(\Phi^*_{X_t}\Phi_{X_t} + \lambda I)^{-1}\Phi(x) &= \Phi^*(x)\Phi(x) - \Phi^*(x)\Phi^*_{X_t}(\Phi_{X_t}\Phi^*_{X_t} + \lambda I)^{-1}\Phi_{X_t}\Phi(x) \\
&= K^M(x,x) - K^M_{xX_t}(K^M_{X_tX_t} + \lambda I)^{-1}K^M_{X_tx} \\
&= K^M_t(x,x).
\end{aligned}$$

Hence, we have

$$\|\Phi^*(x)(\Phi^*_{X_t}\Phi_{X_t} + \lambda I)^{-1}\Phi(x)\|^{1/2}_{\mathrm{op}} = \lambda^{-1/2}\|K^M_t(x,x)\|^{1/2}_{\mathrm{op}}.$$

For the noise term, applying Corollary E.7 of Massiani et al. (2025) (with V being the identity operator) yields that, with probability at least $1-\zeta$, for all $t \geq 1$,

$$\|\Phi^*_{X_t}\boldsymbol{\delta}\|_{(\Phi^*_{X_t}\Phi_{X_t}+\lambda I)^{-1}} \leq \sigma\sqrt{\mathrm{Tr}(R_{\mathcal{V}})}\sqrt{2\log(1/\zeta) + \log\det\big(I + \lambda^{-1}K^M_{\mathrm{tr},X_tX_t}\big)}.$$

Combining the above bounds and observing that $\|w\|_{\mathcal{W}} = \|(Mf)\|_{\mathcal{H}_{K^M}} \leq \Gamma$, we obtain

$$\begin{aligned}
&\|Mf(x) - M\mu_t(x)\|_{\mathcal{M}} \\
&\quad\leq \left(\Gamma + \frac{\sigma\sqrt{\mathrm{Tr}(R_{\mathcal{V}})}}{\sqrt{\lambda}}\sqrt{2\log(1/\zeta) + \log\det\big(I + \lambda^{-1}K^M_{\mathrm{tr},X_tX_t}\big)}\right)\|K^M_t(x,x)\|^{1/2}_{\mathrm{op}}.
\end{aligned}$$

This concludes the proof for Theorem 1.

E Proof of Theorem 2

We first present a useful lemma from Chowdhury & Gopalan (2021).

Lemma 3. *For any $\lambda > 0$ and $t \geq 1$, we have*

$$\log\det\big(I + \lambda^{-1}K^M_{\mathrm{tr},X_tX_t}\big) = \sum_{i=1}^{t} \log\det\big(I + \lambda^{-1}K^M_{\mathrm{tr},i-1}(x_i,x_i)\big).$$

From Lemma 3, we know that $\log\det\big(I + \lambda^{-1}K^M_{\mathrm{tr},X_tX_t}\big)$ is an increasing function of t. Hence, we conclude that $\beta_t \leq \beta_{t+1}$ for all $t \geq 0$.

We can write the regret as

$$\begin{aligned}
\mathcal{R}_T &= \sum_{t=1}^{T} F(x^*) - F(x_t) \\
&= \sum_{t=1}^{T} F(x^*) - \overline{F}_{t-1}(x_t) + \overline{F}_{t-1}(x_t) - F(x_t) \\
&= \sum_{t=1}^{T} F(x^*) - \overline{F}_{t-1}(x^*) + \overline{F}_{t-1}(x^*) - \overline{F}_{t-1}(x_t) + \overline{F}_{t-1}(x_t) - F(x_t) \\
&\leq 2\sum_{t=1}^{T} \beta_{t-1}\|m\|_{\mathcal{M}}\|K^M_{t-1}(x_t,x_t)\|^{1/2}_{\mathrm{op}} \\
&\leq 2\beta_T\|m\|_{\mathcal{M}}\sum_{t=1}^{T}\|K^M_{t-1}(x_t,x_t)\|^{1/2}_{\mathrm{op}} \\
&\leq 2\beta_T\|m\|_{\mathcal{M}}\sqrt{T\sum_{t=1}^{T}\|K^M_{t-1}(x_t,x_t)\|_{\mathrm{op}}}.
\end{aligned}$$

Substituting the definition of β_T yields the result.

F Proof of Theorem 3

The following lemma shows that the operator-valued information quantities can be reduced to scalar information gain with separable kernels. This reduction is key to obtaining explicit regret bounds.

Lemma 4. *Let $K^M_{\mathrm{tr}} = GB^M_0$. The operator-valued quantities in Theorem 2 are upper bounded by the scalar information gain associated with G. In particular,*

$$\log\det\big(I + \lambda^{-1}G_{X_TX_T} \otimes B^M_0\big) \leq \sum_{i=1}^{n} \gamma_T\left(G, \frac{\lambda}{\lambda_i^{B^M_0}}\right),$$

$$\sum_{t=1}^{T}\|K^M_{t-1}(x_t,x_t)\|_{\mathrm{op}} \leq 2C_1\,\gamma_T(G, \frac{\lambda}{C_{K^M}}),$$

where $\{\lambda_i^{B^M_0}\}_i$ denotes the point spectrum of B^M_0, and $\gamma_T(G,\lambda) = \max_{X_T}\log\det\big(I+\lambda^{-1}G_{X_TX_T}\big)$ is the information gain associated with the scalar kernel G, $C_1 = \frac{C^2_{K^M}}{\log(1+\frac{C^2_{K^M}}{\lambda})}$.

The proof of Lemma 4 is provided in Appendix G. For many commonly used scalar kernels G, the information gain $\gamma_T(G,\lambda)$ admits sublinear growth in T. Applying the result from Lemma 4 in Theorem 2, we can bound the cumulative regret by

$$\mathcal{R}_T \leq 2\|m\|_{\mathcal{M}}\left(\Gamma + \frac{\sigma\sqrt{\mathrm{Tr}(R_{\mathcal{V}})}}{\sqrt{\lambda}}\sqrt{2\log(\frac{1}{\zeta}) + \sum_{i=1}^{n}\gamma_T\left(G, \frac{\lambda}{\lambda_i^{B^M_0}}\right)}\right)\sqrt{2C_1T\gamma_T(G, \frac{\lambda}{C_{K^M}})},$$

where $C_1 = \frac{C_{K^M}^2}{\log(1+\frac{C_{K^M}^2}{\lambda})}$. For the linear kernel,

$$\gamma_T = \mathcal{O}(\log T),$$

for the square-exponential kernel,

$$\gamma_T = \mathcal{O}((\log T)^{d+1}),$$

for the Matérn kernel with $\nu > \frac{d(d+1)}{2}$,

$$\gamma_T = \mathcal{O}(T^{d(d+1)/(2\nu+d(d+1))} \log T).$$

Consequently, we have

1. Linear kernel: $\mathcal{R}_T \leq \mathcal{O}\big(\log T\sqrt{T}\big)$,
2. Gaussian kernel: $\mathcal{R}_T \leq \mathcal{O}\big((\log T)^{d+1}\sqrt{T}\big)$,
3. Matérn kernel: $\mathcal{R}_T \leq \mathcal{O}\big(T^{d(d+1)/(2\nu+d(d+1))} \log T\sqrt{T}\big)$.

G Proof of Lemma 4

For the first part, according to the Fredholm determinant for positive definite trace-class operators Minh (2017), we observe that

$$\log\det\big(I + \lambda^{-1} G_{X_T X_T} \otimes B_0^M\big) = \sum_{t=1}^{T}\sum_{i=1}^{n} \log(1 + \frac{1}{\lambda}\lambda_t^G \lambda_i^{B_0^M}) \leq \frac{1}{\lambda}\sum_{t=1}^{T}\sum_{i=1}^{n} \lambda_t^G \lambda_i^{B_0^M} < \infty.$$

Hence, we obtain

$$\begin{aligned}
\log\det\big(I + \lambda^{-1} G_{X_T X_T} \otimes B_0^M\big) &= \sum_{t=1}^{T}\sum_{i=1}^{n} \log(1 + \frac{1}{\lambda}\lambda_t^G \lambda_i^{B_0^M}) \\
&= \sum_{i=1}^{n}\sum_{t=1}^{T} \log(1 + \frac{1}{\lambda}\lambda_t^G \lambda_i^{B_0^M}) \\
&= \sum_{i=1}^{n} \log\det\big(I + \frac{\lambda_i^{B_0^M}}{\lambda} G_{X_T X_T}\big) \\
&\leq \sum_{i=1}^{n} \gamma_T(G, \frac{\lambda}{\lambda_i^{B_0^M}}).
\end{aligned}$$

For the second part, we first present a useful lemma.

Lemma 5. *The operator norm of $K_{t-1}^M(x_t, x_t)$ is given by*

$$\|K_{t-1}^M(x_t, x_t)\|_{\mathrm{op}} = \max_j \lambda_j^{B_0^M} \Big(G(x_t, x_t) - G_{x_t X_{t-1}}\Big(G_{X_{t-1}X_{t-1}} + \frac{\lambda}{\lambda_j^{B_0^M}}\Big)^{-1} G_{X_{t-1}x_t}\Big),$$

where $\lambda_j^{B_0^M}$ is the eigenvalue of B_0^M.

Proof. Since B^M is UBD, i.e., $B^M = \iota_{\mathcal{M}}(B_0^M \otimes I_{\mathcal{V}})\iota_{\mathcal{M}}^{-1}$, we denote the spectral decomposition of B_0^M by $B_0^M = \sum_{j=1}^{n} \lambda_j^{B_0^M} \phi_j \phi_j^*$, where $\phi_j \in \tilde{\mathcal{M}}$ is the eigenfunction. Similarly, $G_{X_{t-1}X_{t-1}} = \sum_{i=1}^{t-1} \lambda_i^G \theta_i \theta_i^*$. We have

$$\begin{aligned}
&K_{t-1}^M(x_t, x_t) \\
&= G(x_t, x_t)\iota_{\mathcal{M}}(B_0^M \otimes I_{\mathcal{V}})\iota_{\mathcal{M}}^{-1} \\
&\quad - (G_{x_t X_{t-1}} \otimes \iota_{\mathcal{M}}(B_0^M \otimes I_{\mathcal{V}})\iota_{\mathcal{M}}^{-1}) \cdot (G_{X_{t-1}X_{t-1}} \otimes \iota_{\mathcal{M}}(B_0^M \otimes I_{\mathcal{V}})\iota_{\mathcal{M}}^{-1} + \lambda I)^{-1} \\
&\qquad \cdot (G_{X_{t-1}x_t} \otimes \iota_{\mathcal{M}}(B_0^M \otimes I_{\mathcal{V}})\iota_{\mathcal{M}}^{-1}) \\
&\overset{(s.1)}{=} \iota_{\mathcal{M}} \left(\big(G(x_t, x_t)B_0^M - (G_{x_t X_{t-1}} \otimes B_0^M)(G_{X_{t-1}X_{t-1}} \otimes B_0^M + \lambda I)^{-1}(G_{X_{t-1}x_t} \otimes B_0^M)\big) \otimes I_{\mathcal{V}}\right) \iota_{\mathcal{M}}^{-1}.
\end{aligned}$$

where in (s.1), we use the fact that $I = I \otimes I_{\mathcal{V}}$. Consequently, we have

$$\|K_{t-1}^M(x_t, x_t)\|_{\text{op}} = \Big\| G(x_t, x_t)B_0^M - (G_{x_t X_{t-1}} \otimes B_0^M)(G_{X_{t-1}X_{t-1}} \otimes B_0^M + \lambda I)^{-1}(G_{X_{t-1}x_t} \otimes B_0^M) \Big\|_{\text{op}},$$

given that $\iota_{\mathcal{M}}$ is a unitary isomorphism. Then, we have

$$\begin{aligned}
(G_{X_{t-1}X_{t-1}} \otimes B_0^M + \lambda I)^{-1} &= \sum_{i=1}^{t-1}\sum_{j=1}^{n}(\lambda_i^G \lambda_j^{B_0^M} + \lambda)^{-1}(\theta_i \otimes \phi_j)(\theta_i \otimes \phi_j)^* \\
&= \sum_{i=1}^{t-1}\sum_{j=1}^{n}(\lambda_i^G \lambda_j^{B_0^M} + \lambda)^{-1}(\theta_i \otimes \phi_j)(\theta_i^* \otimes \phi_j^*) \\
&= \sum_{j=1}^{n}\sum_{i=1}^{t-1}(\lambda_i^G \lambda_j^{B_0^M} + \lambda)^{-1}(\theta_i \theta_i^*) \otimes (\phi_j \phi_j^*) \\
&= \sum_{j=1}^{n}\Big(\lambda_j^{B_0^M} G_{X_{t-1}X_{t-1}} + \lambda I\Big)^{-1} \otimes (\phi_j \phi_j^*).
\end{aligned}$$

For $\|K_{t-1}^M(x_t, x_t)\|_{\text{op}}$, we have

$$\begin{aligned}
&\|K_{t-1}^M(x_t, x_t)\|_{\text{op}} \\
&= \Big\| G(x_t, x_t)B_0^M - (G_{x_t X_{t-1}} \otimes B_0^M)(G_{X_{t-1}X_{t-1}} \otimes B_0^M + \lambda I)^{-1}(G_{X_{t-1}x_t} \otimes B_0^M) \Big\|_{\text{op}} \\
&= \Big\| \sum_{j=1}^{n} G(x_t, x_t)\lambda_j^{B_0^M}\phi_j\phi_j^* - \sum_{j=1}^{n} G_{x_t X_{t-1}}\Big(\lambda_j^{B_0^M} G_{X_{t-1}X_{t-1}} + \lambda I\Big)^{-1} G_{X_{t-1}x_t}(B_0^M \phi_j \phi_j^* B_0^M) \Big\|_{\text{op}} \\
&= \Big\| \sum_{j=1}^{n} \Big(G(x_t, x_t)\lambda_j^{B_0^M} - (\lambda_j^{B_0^M})^2 G_{x_t X_{t-1}}\Big(\lambda_j^{B_0^M} G_{X_{t-1}X_{t-1}} + \lambda I\Big)^{-1} G_{X_{t-1}x_t}\Big)\phi_j\phi_j^* \Big\|_{\text{op}}.
\end{aligned}$$

Following this expression, we conclude the proof by choosing the largest eigenvalue. □

Now, according to Mercer's theorem Micchelli et al. (2006), there exists a mapping $\varphi : \mathcal{X} \to \ell^2$ such that $G(x, x) = \varphi^\top(x)\varphi(x)$. We also define an operator $\varphi_{X_{t-1}} : \ell^2 \to \mathbb{R}^{t-1}$ via $\varphi_{X_{t-1}}a = [\varphi(x_1)^\top a, \ldots, \varphi(x_{t-1})^\top a]^\top, \forall a \in \ell^2$. Then $G_{X_{t-1}X_{t-1}}$ can be expressed as $\varphi_{X_{t-1}}\varphi_{X_{t-1}}^\top$. Then, applying the Woodbury identity, we have

$$\begin{aligned}
&G(x, x) - G_{xX_{t-1}}(G_{X_{t-1}X_{t-1}} + \lambda I)^{-1} G_{X_{t-1}x} \\
&\qquad = \varphi(x)^\top\varphi(x) - \varphi(x)^\top\varphi_{X_{t-1}}^\top(\varphi_{X_{t-1}}\varphi_{X_{t-1}}^\top + \lambda I)^{-1}\varphi_{X_{t-1}}\varphi(x) \\
&\qquad = \lambda\varphi(x)^\top(\varphi_{X_{t-1}}^\top\varphi_{X_{t-1}} + \lambda I)^{-1}\varphi(x).
\end{aligned}$$

Applying Lemma 5, we have

$$\begin{aligned}
\|K_{t-1}^M(x_t, x_t)\|_{\text{op}} &= \max_j \lambda_j^{B_0^M}\Big(G(x_t, x_t) - G_{x_t X_{t-1}}\Big(G_{X_{t-1}X_{t-1}} + \frac{\lambda}{\lambda_j^{B_0^M}} I\Big)^{-1} G_{X_{t-1}x_t}\Big) \\
&= \max_j \lambda\varphi(x_t)^\top(\varphi_{X_{t-1}}^\top\varphi_{X_{t-1}} + \frac{\lambda}{\lambda_j^{B_0^M}} I)^{-1}\varphi(x_t) \\
&\leq \lambda\varphi(x_t)^\top(\varphi_{X_{t-1}}^\top\varphi_{X_{t-1}} + \frac{\lambda}{C_{K^M}} I)^{-1}\varphi(x_t) \\
&= C_{K^M}\Big(G(x_t, x_t) - G_{x_t X_{t-1}}(G_{X_{t-1}X_{t-1}} + \frac{\lambda}{C_{K^M}} I)^{-1} G_{X_{t-1}x_t}\Big).
\end{aligned}$$

Summing over t and writing $G(x_t, x_t) - G_{x_t X_{t-1}}(G_{X_{t-1}X_{t-1}} + \frac{\lambda}{C_{K^M}} I)^{-1} G_{X_{t-1}x_t}$ as $s_{t-1}(x_t)$, we have

$$\begin{aligned}
\sum_{t=1}^{T} \|K_{t-1}^M(x_t, x_t)\|_{\mathrm{op}} &\le C_{K^M} \sum_{t=1}^{T} G(x_t, x_t) - G_{x_t X_{t-1}}(G_{X_{t-1}X_{t-1}} + \frac{\lambda}{C_{K^M}} I)^{-1} G_{X_{t-1}x_t} \\
&= C_{K^M} \frac{\lambda}{C_{K^M}} \sum_{t=1}^{T} (\frac{C_{K^M}}{\lambda} s_{t-1}(x_t)) \\
&\overset{\text{(s.1)}}{\le} \lambda \sum_{t=1}^{T} \frac{\frac{C_{K^M}^2}{\lambda}}{\log(1 + \frac{C_{K^M}^2}{\lambda})} \log(1 + \frac{C_{K^M}}{\lambda} s_{t-1}(x_t)) \\
&= \frac{C_{K^M}^2}{\log(1 + \frac{C_{K^M}^2}{\lambda})} \log\det(I + \frac{C_{K^M}}{\lambda} G_{X_T X_T}) \\
&\le 2C_1 \gamma_T(G, \frac{\lambda}{C_{K^M}}),
\end{aligned}$$

where we denote $\frac{C_{K^M}^2}{\log(1+\frac{C_{K^M}^2}{\lambda})}$ compactly by C_1. In (s.1), we use the monotonicity of $\frac{x}{\log(1+x)}$ for $x \in [0, \frac{C_{K^M}^2}{\lambda}]$ and the fact that $\frac{C_{K^M}}{\lambda} s_{t-1}(x_t) \in [0, \frac{C_{K^M}^2}{\lambda}]$.

H Extension to time-varying and nonlinear objectives

H.1 Time-varying objective

When the functional m_t varies over time, the objective becomes $F_t(x) = \langle m_t, Mf(x) \rangle_{\mathcal{M}}$. At each iteration, the objective admits uncertainty bounds $F_t(x) \in [\underline{F}_t(x), \overline{F}_t(x)]$, where

$$\begin{aligned}
\overline{F}_t(x) &= \langle m_t, M\mu_{t-1}(x) \rangle_{\mathcal{M}} + \beta_{t-1} \|m_t\|_{\mathcal{M}} \|K_{t-1}^M(x, x)\|_{\mathrm{op}}^{1/2}, \\
\underline{F}_t(x) &= \langle m_t, M\mu_{t-1}(x) \rangle_{\mathcal{M}} - \beta_{t-1} \|m_t\|_{\mathcal{M}} \|K_{t-1}^M(x, x)\|_{\mathrm{op}}^{1/2}.
\end{aligned}$$

Since only the functional m_t varies while the estimator of Mf remains unchanged, the uncertainty quantification in the measurement space can be reused across time. This allows the framework to handle time-varying objectives without modifying the underlying estimator.

Based on this formulation, we present Algorithm 2 for sequential optimization under time-varying objectives. The algorithm follows the same principle as its static counterpart, with the key difference that the confidence width $\beta_{t-1} \|m_t\|_{\mathcal{M}} \|K_{t-1}^M(x, x)\|_{\mathrm{op}}^{1/2}$ now depends on the current objective m_t. At iteration t, the next point is therefore selected according to

$$x_t \in \arg\max_{x \in \mathcal{X}} \overline{F}_t(x).$$

The objective is evaluated from the vector-valued model rather than being modeled directly as an input-to-scalar map. Therefore, once confidence bounds for the structured output have been established, the same framework naturally accommodates different known objectives defined on that output.

We next show that the regret analysis extends directly to this setting. Let

$$x_t^\star \in \arg\max_{x \in \mathcal{X}} F_t(x)$$

denote an optimizer of the objective at iteration t. The cumulative regret for the time-varying objective is

$$\begin{aligned}
\mathcal{R}_T^{\mathrm{tv}} &= \sum_{t=1}^{T} (F_t(x_t^\star) - F_t(x_t)) \\
&= \sum_{t=1}^{T} \langle m_t, Mf(x_t^\star) - Mf(x_t) \rangle_{\mathcal{M}}. \qquad (4)
\end{aligned}$$

Algorithm 2 Time-varying vector-valued Bayesian optimization

Require: Total budget T, regularization parameter $\lambda > 0$, confidence level $\zeta \in (0, 1)$
1: Initialize dataset $X_0 = \emptyset$, $Y_0 = \emptyset$ with $M\mu_0(x) = 0$, $K_0^M(x, x) = K^M(x, x)$
2: **for** $t = 1, \ldots, T$ **do**
3: Compute $M\mu_{t-1}(x)$ and $K_{t-1}^M(x, x)$
4: Select

$$x_t \in \arg\max_{x \in \mathcal{X}} \langle m_t, M\mu_{t-1}(x) \rangle_{\mathcal{M}} + \beta_{t-1} \|m_t\|_{\mathcal{M}} \|K_{t-1}^M(x, x)\|_{\mathrm{op}}^{1/2}$$

5: Query f at x_t and observe $y_t = Mf(x_t) + \delta_t$
6: Update X_t, Y_t
7: **end for**

On the high-probability event of Theorem 1, the same decomposition as in the proof of Theorem 2 gives

$$\begin{aligned} F_t(x_t^\star) - F_t(x_t) &\le \overline{F}_t(x_t) - F_t(x_t) \\ &\le 2\beta_{t-1} \|m_t\|_{\mathcal{M}} \|K_{t-1}^M(x_t, x_t)\|_{\mathrm{op}}^{1/2}. \end{aligned}$$

Consequently,

$$\mathcal{R}_T^{\mathrm{tv}} \le 2\beta_T \sum_{t=1}^{T} \|m_t\|_{\mathcal{M}} \|K_{t-1}^M(x_t, x_t)\|_{\mathrm{op}}^{1/2}. \tag{5}$$

If the time-varying functionals are uniformly bounded such that $\|m_t\|_{\mathcal{M}} \le C_m$ for all t, then, by the Cauchy–Schwarz inequality,

$$\mathcal{R}_T^{\mathrm{tv}} \le 2C_m\beta_T \sqrt{T \sum_{t=1}^{T} \|K_{t-1}^M(x_t, x_t)\|_{\mathrm{op}}}. \tag{6}$$

The remaining sum is exactly the same uncertainty term as in Theorem 2 and can therefore be bounded using the same information-gain argument.

H.2 Nonlinear Lipschitz objective

Let $F' : \mathcal{M} \to \mathbb{R}$ be $L_{F'}$-Lipschitz, i.e., $|F'(m_1) - F'(m_2)| \le L_{F'} \|m_1 - m_2\|_{\mathcal{M}}$. Applying this with $m_1 = Mf(x)$ and $m_2 = M\mu_t(x)$ gives

$$\begin{aligned} F'(Mf(x)) &\le F'(M\mu_t(x)) + L_{F'} \|Mf(x) - M\mu_t(x)\|_{\mathcal{M}}, \\ F'(Mf(x)) &\ge F'(M\mu_t(x)) - L_{F'} \|Mf(x) - M\mu_t(x)\|_{\mathcal{M}}. \end{aligned}$$

Thus, the same uncertainty bound in the measurement space can be propagated to nonlinear objectives via Lipschitz continuity. The remainder of the argument follows from the linear case.

I Efficient computation

When the measurement space $\mathcal{M}$ is an infinite-dimensional separable Hilbert space. For any element in $\mathcal{M}$, it can be represented via an orthonormal basis $\{\phi\}_i^\infty$. Those basis vectors can be taken as the eigenfunctions of B^M if B^M is a positive self-adjoint trace-class operator given

$$B^M = \sum_{i=1}^{\infty} \lambda_i^{B^M} \phi_i \phi_i^*.$$

Finite-rank approximation. To obtain a numerically tractable representation, we truncate the basis $\{\phi\}_i^\infty$ to the first n basis functions $\{\phi\}_i^n$. Under this finite-dimensional approximation, any observation y_i can be approximated in the basis $\{\phi_i\}_{i=1}^n$ as

$$y_i = \sum_{j=1}^{n} \bar{Y}_{ij} \phi_j,$$

where $\bar{Y}_t \in \mathbb{R}^{t \times n}$ collects the corresponding coefficients.

Representation of the estimator. We denote the finite matrix representation of B^M in the basis $\{\phi_i\}_{i=1}^n$ by B_n^M and denote its eigenvalues by $\lambda_i^{B_n^M}$. Using this basis representation, the coefficients of $M\mu_t(x)$ can be computed explicitly. In particular, the coefficient corresponding to ϕ_i is given by

$$\lambda_i^{B_n^M} G_{xX_t}\big(\lambda_i^{B_n^M} G_{X_tX_t} + \lambda I\big)^{-1}\bar{Y}_{:,i},$$

where $\bar{Y}_{:,i} \in \mathbb{R}^t$ denotes the i-th column of $\bar{Y}_t$.

Efficient evaluation. The above representation allows efficient computation of key quantities. For example,

$$\begin{aligned}
&\log\det\big(\lambda^{-1}G_{X_tX_t}\otimes B_n^M + I\big)\\
&\qquad = \sum_{j=1}^{n} \log\det\big(\lambda_j^{B_n^M}\lambda^{-1}G_{X_tX_t} + I\big),\\
\langle m, M\mu_t(x)\rangle_{\mathcal{M}} &= \sum_{i=1}^{n} \bar{m}_i \lambda_i^{B_n^M} G_{xX_t}\big(\lambda_i^{B_n^M} G_{X_tX_t} + \lambda I\big)^{-1}\bar{Y}_{:,i},\\
K_t^M(x,x) &= \sum_{j=1}^{n}\Big(G(x,x)\lambda_j^{B_n^M} - (\lambda_j^{B_n^M})^2 G_{xX_t}\big(\lambda_j^{B_n^M}G_{X_tX_t} + \lambda I\big)^{-1}G_{X_tx}\Big)\phi_j\phi_j^*,
\end{aligned}$$

where $\bar{m}_i$ is the coordinate of m in the basis $\{\phi_i\}_{i=1}^n$, i.e., $m = \sum_{i=1}^n \bar{m}_i\phi_i$.

Optimization. To select the next query point, we solve

$$x_t \in \arg\max_{x\in\mathcal{X}} \langle m, M\mu_{t-1}(x)\rangle_{\mathcal{M}} + \beta_{t-1}\|m\|_{\mathcal{M}}\|K_{t-1}^M(x,x)\|_{\mathrm{op}}^{1/2}.$$

This objective is generally non-convex in x. For low-dimensional problems, one can discretize the input space and perform grid search. For higher-dimensional settings, gradient-based optimization with multiple random initializations can be used.

Finite-dimensional approximation. Here, the infinite-dimensional measurement space is approximated by a finite-dimensional subspace obtained by retaining only a finite number of basis functions. This approximation captures the dominant components of the measurements while discarding the remaining higher-order directions. The resulting truncation error is therefore determined by the amount of information contained in these neglected components. As the number of retained basis functions increases, the finite-dimensional representation becomes increasingly accurate and converges to the original Hilbert-space formulation. Thus, the truncation level provides a natural trade-off between computational efficiency and approximation accuracy.

J Numerical experiments

J.1 Experiment setup

Objective design. To enable a fair comparison across methods with different observation models, we construct the objective independently of the measurement operator. Let $f(x) \in \mathcal{Y}$ denote the underlying vector-valued output. We define the objective as $F(x) = \langle m, f(x)\rangle_{\mathcal{Y}}$, $m = \sum_{j=1}^q w_j\xi_j$, where $\Xi = [\xi_1, \ldots, \xi_q] : \mathbb{R}^q \to \mathcal{Y}$ is a collection of linear functionals represented in $\mathcal{Y}$, and $w \in \mathbb{R}^q$ are the associated weights. We compactly write $m = \Xi w$. This formulation is convenient for evaluating vvBO under multiple observation models. In the experiments, we consider: (i) full observation, where $M = I$, $\mathcal{M} = \mathcal{Y}$, and the full function $f(x)$ is observed; (ii) finite-dimensional observation, where $M = \Xi^*$, $\mathcal{M} = \mathbb{R}^q$, and only the projections $\Xi^* f(x) = [\langle\xi_1, f(x)\rangle_{\mathcal{Y}}, \ldots, \langle\xi_q, f(x)\rangle_{\mathcal{Y}}]^\top \in \mathbb{R}^q$ are available. Under the finite-dimensional observation model, $F(x) = \langle m, f(x)\rangle_{\mathcal{Y}} = \langle w, \Xi^* f(x)\rangle_{\mathbb{R}^q}$, which provides a unified formulation for benchmark methods with different measurement assumptions. In particular, methods operating on scalar, finite-dimensional, or function-valued measurements can all be compared while optimizing an equivalent objective derived from the same black-box operator f.

Implementation of baseline methods. Here are the implementation details for each baseline method.

- *Scalar-valued BO (BO).* For vanilla Bayesian optimization, only the scalar objective value $F(x)$ is measured at each query.
- *Retrained scalar-valued BO (rBO).* A variant of BO that is re-initialized whenever the objective changes, discarding previously collected data.
- *Multi-task BO (MTBO).* Multi-task BO assumes access to vector-valued observations in a finite-dimensional space Chowdhury & Gopalan (2021). In our experiments, for each $x \in \mathcal{X}$, MTBO measures
$$\Xi^* f(x) \in \mathbb{R}^q,$$
and constructs a surrogate model mapping from $x \in \mathcal{X}$ to $\Xi^* f(x) \in \mathbb{R}^q$. The objective is then evaluated as
$$F(x) = \langle w, \Xi^* f(x) \rangle_{\mathbb{R}^q}.$$
For vvBO, when $M = \Xi^*$, vvBO and MTBO operate on the same finite-dimensional observations. In contrast, under full observations ($M = I$), vvBO directly models the full output $f(x)$, while MTBO remains restricted to the projected measurements $\Xi^* f(x) \in \mathbb{R}^q$.
- *Retrained multi-task BO (rMTBO).* A variant of MTBO that is re-initialized whenever the objective changes, discarding previously collected data. Since the objective changes through the weights w, rMTBO updates the scalarized objective accordingly while reconstructing the model from scratch using the newly collected observations.
- *Contextual BO (CTBO).* The objective $F(x) = \langle m, f(x) \rangle_{\mathcal{Y}}$ can equivalently be written as a scalar-valued function $\tilde{F} : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ defined by
$$\tilde{F}(x, m) := \langle m, f(x) \rangle_{\mathcal{Y}}.$$
Thus, $F(x) = \tilde{F}(x, m)$ where m acts as a context variable when the objective changes.
If f belongs to a vector-valued RKHS with kernel K, then $\tilde{F}$ lies in a scalar-valued RKHS with kernel Carmeli et al. (2006); Micchelli & Pontil (2005)
$$\tilde{K}\big((x_1, m_1), (x_2, m_2)\big) = \langle m_1, K(x_1, x_2) m_2 \rangle_{\mathcal{Y}}.$$
Given observations
$$\tilde{X}_t = \{(x_i, m_i)\}_{i=1}^t, \qquad \tilde{Y}_t = \{\tilde{F}(x_i, m_i)\}_{i=1}^t,$$
the standard scalar-valued GP regression can be applied with posterior mean and variance given by
$$\tilde{\mu}_t(x, m) = \tilde{K}_{(x,m)\tilde{X}_t} \big(\tilde{K}_{\tilde{X}_t \tilde{X}_t} + \lambda I\big)^{-1} \tilde{Y}_t,$$
$$\tilde{\sigma}_t^2(x, m) = \tilde{K}((x, m), (x, m)) - \tilde{K}_{(x,m)\tilde{X}_t} \big(\tilde{K}_{\tilde{X}_t \tilde{X}_t} + \lambda I\big)^{-1} \tilde{K}_{\tilde{X}_t (x,m)}.$$
Common acquisition functions such as UCB, expected improvement, and Thompson sampling can then be applied Garnett (2023). In this work, we use the UCB acquisition function.
- *Function-on-function BO (FFBO).* FFBO is implemented following Huang et al. (2026), corresponding to CTBO with a fixed context. The observations are scalar objective values $F(x)$.

RKHS representation of vector-valued outputs. In our experiments, the system output $f(x)$ is a function or a trajectory. For numerical implementation, we construct a finite-dimensional representation. Specifically, for each input x, the function $f(x)$ is evaluated on a dense grid $\{t_j\}_{j=1}^{50}$, yielding samples $y_j(x) = f(x)(t_j)$. We then approximate $f(x)$ using $\widehat{f(x)}$, which is in a scalar-valued RKHS and expressed as $\widehat{f(x)}(t) = \sum_{j=1}^{50} \alpha_{f(x),j}\, k(t, t_j)$, where k is a scalar-valued kernel defined on the output domain. The coefficient vector $\alpha_{f(x)} = [\alpha_{f(x),1}, \ldots, \alpha_{f(x),50}]^\top \in \mathbb{R}^{50}$ is obtained by solving a scalar-valued KRR problem
$$\min_{\widehat{f(x)} \in \mathcal{H}_k} \sum_{j=1}^{50} \big(\widehat{f(x)}(t_j) - y_j(x)\big)^2 + 0.01 \|\widehat{f(x)}\|_{\mathcal{H}_k}^2.$$

Table 1: Hyperparameters for different test operators.

Test Operator	Kernel	λ	Noise level	Input length scale(s)	Output length scale(s)
GP	RBF	10^{-2}	10^{-2}	0.1	0.1
GP (3D)	RBF	10^{-2}	10^{-2}	0.1	0.1
Ackley	RBF	10^{-2}	10^{-2}	3	3
Eggholder	RBF	10^{-2}	10^{0}	50	50
Bukin	RBF	10^{-2}	10^{0}	0.6	1
Holder Table	RBF	10^{-2}	10^{0}	1	1
Shubert	RBF	10^{-2}	10^{-3}	0.5	0.5
Langermann	RBF	10^{-2}	10^{-3}	0.5	0.5
Controller tuning	RBF	10^{-2}	10^{-1}	$[0.6, 0.8, 0.8]$	$[5000, 10000]$

Table 2: Computation time (ms) of vvBO across different iterations (mean $\pm$ std) for different test operators.

Test Operator	Iteration 30	Iteration 60	Iteration 100
GP	8.16 ± 0.04	8.59 ± 0.10	9.22 ± 0.17
Ackley	10.08 ± 0.22	10.62 ± 0.23	11.42 ± 0.19
Eggholder	8.13 ± 0.19	8.53 ± 0.18	9.20 ± 0.20
Bukin	8.24 ± 0.15	8.56 ± 0.18	9.39 ± 0.18
Holder Table	10.13 ± 0.25	10.60 ± 0.18	11.39 ± 0.27
Shubert	7.98 ± 0.17	8.39 ± 0.15	9.22 ± 0.26
Langermann	8.06 ± 0.16	8.42 ± 0.09	9.11 ± 0.19

Similarly, each linear functional ξ is represented in the same basis as $\xi(\cdot) = \sum_{j=1}^{50} \alpha_{\xi,j}\, k(\cdot, t_j)$, with coefficients $\alpha_{\xi,j} \in \mathbb{R}$. Linear functionals of $f(x)$ are then evaluated as RKHS inner products in this finite-dimensional representation. All experiments use this finite representation of $f(x)$ and the functionals ξ, ensuring a consistent and fair comparison. This construction is introduced solely for numerical implementation and does not alter the theoretical results derived in this work.

Hyperparameters. All simulations are conducted on an HP OMEN 45L Desktop equipped with an Intel Core Ultra 9 285K processor (24 cores, up to 6.5 GHz) and 62 GB RAM, running Ubuntu 24.04.3 LTS (Linux 6.14). The computation times of vvBO for the different test operators at different iterations are reported in Table 2.

The hyperparameters for the different test operators are summarized in Table 1 and are kept identical across the corresponding baseline methods. For vvBO, the operator-valued kernel is constructed as the product of an RBF kernel and the identity operator. The kernel hyperparameters, in particular the RBF length scales, were selected manually based on prior knowledge of the characteristic scales of each test problem and were fixed before running the reported experiments. The regularization parameter λ was likewise selected manually to provide a reasonable balance between data fitting and numerical stability and was kept fixed across independent runs.

The exploration parameter β used in the numerical experiments was also selected manually to provide a reasonable balance between exploration and exploitation and was fixed for each experimental setting. This practical choice is distinguished from the theoretical confidence parameter β_t introduced in Section 4, which is obtained from the high-probability confidence bound and depends on λ, ζ, σ, Γ, and the kernel matrix. Directly using the theoretical expression can be overly conservative in practice, as it is designed to provide a uniform high-probability guarantee. We therefore treat β as a tunable exploration parameter in the experiments. More systematic data-driven procedures, such as marginal-likelihood optimization under a Gaussian-process interpretation, could be used to tune the kernel and regularization/noise parameters. Developing a systematic hyperparameter-selection procedure is not the main objective of this work; therefore, we select these parameters manually based on prior knowledge of the corresponding test problems.

J.2 SYNTHETIC PROBLEM

J.2.1 TEST OPERATORS

We evaluate the proposed method on synthetic problems constructed from standard optimization benchmarks. For each benchmark, the input is decomposed as $u = (x, t)$, where x denotes the input variable and t indexes the structured output. The latent vector-valued function is generated as

$$f(x)(t) = h(x, t),$$

where h is one of the following ground-truth functions. The scalar objective is then obtained by applying a linear functional to the trajectory $f(x)$.

Gaussian process. For the GP benchmark, we generate a smooth function using an RBF kernel expansion. Let $X_{N_x} = \{x_i\}_{i=1}^{N_x}$ and $T_{N_t} = \{t_j\}_{j=1}^{10}$ denote fixed grids on $[0, 1]$. N_x changes with grid density. For one-dimensional inputs, $N_x = 10$. For three-dimensional inputs, $N_x = 125$. Let $\alpha \in \mathbb{R}^{N_x \times N_t}$ be a matrix with randomly sampled elements, i.e., α_{ij} is sampled uniformly from $[-3.5, 3.5]$. The ground truth is $h(x, t) = G_{xX_{N_x}} \alpha G_{T_{N_t} t}$, where G is an RBF kernel with length scale $\ell = 0.1$. We implement both a one-dimensional version, denoted GP, and a higher-dimensional version, denoted GP (3D), where $x \in \mathbb{R}^3$.

Ackley. For $u \in \mathbb{R}^d$, the Ackley benchmark is

$$h(u) = -\left(-20\exp\left(-0.2\sqrt{\frac{1}{d}\sum_{i=1}^{d} u_i^2}\right) - \exp\left(\frac{1}{d}\sum_{i=1}^{d}\cos(0.2\pi u_i)\right) + 20 + e - 20\right).$$

The domain is $u_i \in [-32.768, 32.768]$.

Eggholder. For $u = (u_1, u_2)$,

$$h(u) = -\left(\frac{u_2}{2} + 47\right)\sin\left(\sqrt{\frac{1}{2}\left|\frac{u_2}{2} + \frac{u_1}{4} + 47\right|}\right) - \frac{u_1}{2}\sin\left(\sqrt{\frac{1}{2}\left|\frac{u_1}{2} - \left(\frac{u_2}{2} + 47\right)\right|}\right),$$

with domain $u_i \in [-512, 512]$.

Bukin. For $u = (u_1, u_2)$,

$$h(u) = -100\sqrt{|u_2 - 0.01u_1^2|} + 0.01|u_1 + 10| + 180,$$

with $u_1 \in [-15, -5]$ and $u_2 \in [-3, 3]$.

Holder Table. For $u = (u_1, u_2)$,

$$h(u) = \left|\sin(u_1)\cos(u_2)\exp\left(\left|1 - \frac{\sqrt{u_1^2 + u_2^2}}{\pi}\right|\right)\right|,$$

with $u_i \in [-10, 10]$.

Shubert. For $u = (u_1, u_2)$,

$$h(u) = \frac{1}{100}\left[\sum_{i=1}^{5} i\cos\Big((i+1)\frac{u_1}{2} + i\Big)\right]\left[\sum_{i=1}^{5} i\cos\Big((i+1)\frac{u_2}{2} + i\Big)\right],$$

with $u_i \in [-10, 10]$.

Langermann. For $u = (u_1, u_2)$,

$$h(u) = \sum_{i=1}^{5} c_i \exp\left(-\frac{\|u/2 - A_i\|_2^2}{\pi}\right) \cos\left(\pi \|u/2 - A_i\|_2^2\right),$$

where

$$c = (1, 2, 5, 2, 3), \qquad A = \begin{pmatrix} 3 & 5 \\ 5 & 2 \\ 2 & 1 \\ 1 & 4 \\ 7 & 9 \end{pmatrix},$$

with $u_i \in [0, 10]$.

J.2.2 Implementation details

We consider three phases with changing objectives. Each objective is defined as $F_i(x) = \langle m_i, f(x) \rangle_{\mathcal{Y}}$, $m_i = \Xi_i w_i$, where $\Xi_i = [\xi_{i,1}, \ldots, \xi_{i,5}] : \mathbb{R}^5 \to \mathcal{Y}$ collects a set of base functionals and $w_i \in \mathbb{R}^5$ are the associated weights. From phase one to phase two, only the weights w_i change while the functional basis Ξ_i remains fixed. From phase two to phase three, both Ξ_i and w_i are modified. For vvBO, we consider two observation regimes. Under full observations ($M = I$), the full output $f(x)$ is observed and all three phases are evaluated. Under partial observations ($M = \Xi^*$), only projected measurements $\Xi^* f(x) \in \mathbb{R}^5$ are available. In this case, $F_i(x) = \langle m_i, f(x) \rangle_{\mathcal{Y}} = \langle w_i, \Xi_i^* f(x) \rangle_{\mathbb{R}^5}$. Since the measurement operator is fixed to Ξ^*, only the transition from phase one to phase two is considered under partial observations, while phase three is omitted because changing Ξ_i would also modify the measurement operator.

The corresponding functional bases Ξ_i and weights w_i for each test operator are summarized in Table 3. Here, ξ_{int} denotes a randomly generated integral functional, while $\xi(t)$ denotes a point-evaluation functional at t. Table 4 summarizes the inputs and measurements available to each method across phases. In particular, vvBO with full observations ($M = I$) directly models the full output $f(x)$, whereas vvBO with partial observations ($M = \Xi^*$) only accesses the projected measurements $\Xi^* f(x) \in \mathbb{R}^5$. MTBO and rMTBO operate on finite-dimensional observations constructed from $\Xi_i^* f(x) \in \mathbb{R}^5$, while BO and rBO only observe the scalar objective values. CTBO treats m_i as a context variable and models the scalar-valued mapping $(x, m_i) \mapsto \tilde{F}(x, m_i)$. FFBO corresponds to CTBO with fixed context and only observes scalar objective values.

J.2.3 Additional experiments

Full observation. We report additional experiments under full observation ($M = I$), where the complete output $f(x)$ is available. The results are shown in Fig. 4. These results confirm that vvBO adjusts effectively across varying objectives and that the benefit of structured measurements persists across different benchmarks.

Partial observation. We further evaluate the proposed method under partial observation settings, where the measurement operator maps the output to a finite-dimensional space ($M = \Xi^*$). The results are shown in Fig. 5 and Fig. 6. Across all test operators, vvBO consistently outperforms the baseline methods in terms of both simple regret and cumulative regret, effectively adapting to different objectives. In contrast, scalar-valued BO methods suffer from information loss due to observing only the objective value, while MTBO has limited performance and does not adapt to changes in the objective function. These results demonstrate that vvBO retains its advantage under different observation models.

Comparison between vvBO and CTBO. The performance of CTBO in tracking changing objectives depends critically on the variation of the context m. When m shifts to a distant region, CTBO lacks sufficient samples in that region to construct an accurate model, resulting in more exploration. In contrast, when m changes only slightly, CTBO can effectively reuse existing samples from the previous context to reduce uncertainty. This behavior is illustrated in Fig. 7. From phase one to phase two, we first change the weight from $[1, 0, 0, 0, 0]$ to $[0, 0, 1, 0, 0]$, corresponding to optimizing only the third component ξ_3. This induces a large shift in the context, leading to a significant increase in

Table 3: Functional basis Ξ_i and weights w_i defining the objectives $F_i(x) = \langle \Xi_i w_i, f(x) \rangle_{\mathcal{Y}}$ across phases. From phase one to phase two, only the weights w_i change, while from phase two to phase three both Ξ_i and w_i are modified.

Test Operator	Phase	Functional basis Ξ_i	$w_i \in \mathbb{R}^5$	β
GP, GP(3D)	one	$[\xi(0), \xi(0.1), \xi(0.2), \xi(0.3), \xi(0.4)]$	$\frac{1}{5}[1,1,1,1,1]$	6
	two	same as phase one	$[0,1,0,0,0]$	6
	three	$[\xi(0.5), \xi(0.6), \xi(0.7), \xi(0.8), \xi(0.9)]$	$\frac{1}{5}[1,1,1,1,1]$	6
Ackley	one	$[\xi_{\text{int}}^{(1)}, \ldots, \xi_{\text{int}}^{(5)}]$	$\frac{1}{5}[1,1,1,1,1]$	10
	two	same as phase one	$\frac{1}{4}[1,0,1,1,1]$	40
	three	$[\xi_{\text{int}}^{(1)}, \ldots, \xi_{\text{int}}^{(5)}]$	$\frac{1}{5}[1,1,1,1,1]$	70
Bukin	one	$[\xi(0), \xi(0.5), \xi(1), \xi(1.5), \xi(2)]$	$\frac{1}{5}[1,1,1,1,1]$	100
	two	same as phase one	$[0,0,0,0,1]$	115
	three	$[\xi(0), \xi(-0.5), \xi(-1), \xi(-1.5), \xi(-2)]$	$\frac{1}{5}[1,1,1,1,1]$	80
Eggholder	one	$[\xi(500), \xi(400), \xi(300), \xi(200), \xi(100)]$	$[1,0,0,0,0]$	400
	two	same as phase one	$[0,0,1,0,0]$	250
	three	$[\xi(0), \xi(-100), \xi(-200), \xi(-300), \xi(-400)]$	$[0,0,0,0,1]$	300
Holder Table	one	$[\xi_{\text{int}}^{(1)}, \ldots, \xi_{\text{int}}^{(5)}]$	$\frac{1}{5}[1,1,1,1,1]$	30
	two	same as phase one	$[1,0,0,0,0]$	30
	three	$[\xi_{\text{int}}^{(1)}, \ldots, \xi_{\text{int}}^{(5)}]$	$\frac{1}{5}[1,1,1,1,1]$	5
Shubert	one	$[\xi(0), \xi(1), \xi(2), \xi(3), \xi(4)]$	$\frac{1}{5}[1,1,1,1,1]$	0.5
	two	same as phase one	$[0,0,0,1,0]$	0.5
	three	$[\xi(0), \xi(-1), \xi(-2), \xi(-3), \xi(-4)]$	$[0,0,0,0,1]$	1
Langermann	one	$[\xi(5), \xi(6), \xi(7), \xi(8), \xi(9)]$	$\frac{1}{5}[1,1,1,1,1]$	3
	two	same as phase one	$[1,0,0,0,0]$	3
	three	$[\xi(0), \xi(1), \xi(2), \xi(3), \xi(4)]$	$[1,0,0,0,0]$	3

cumulative regret due to the need for extensive exploration in a previously unseen region. In contrast, when the weight is changed from $[1,0,0,0,0]$ to $[0.8,0.2,0,0,0]$, the induced change is small. In this case, CTBO can leverage previously collected data more effectively, resulting in reduced exploration cost and lower cumulative regret. The corresponding uncertainty estimates for CTBO are shown in Fig. 8. When transitioning to $[0,0,1,0,0]$, the model exhibits high uncertainty due to the lack of nearby samples. For the smaller change, the uncertainty remains well-controlled, enabling more efficient exploration. These results highlight the key limitation of CTBO. The performance is highly sensitive to the variation of the context. In contrast, vvBO models the underlying vector-valued operator directly and therefore maintains consistent performance regardless of changes in the objective functional.

J.3 Real-world controller tuning

We evaluate the proposed method on a real-world building control task using the Building Optimization Testing Framework (BOPTEST) Blum et al. (2021), a high-fidelity simulation framework for benchmarking advanced HVAC control strategies. BOPTEST provides physics-based building emulators implemented in Modelica and exposed through standardized APIs, enabling realistic closed-loop evaluation under dynamic weather, occupancy, and pricing conditions Blum et al. (2021).

We consider the '`singlezone_commercial_hydronic`' test case, which models a commercial building zone served by a hydronic radiator connected to a district heating network. The building dynamics include thermal interactions between the indoor air, building envelope, external weather conditions, solar gains, and internal heat loads. Heating is delivered through a water-based radiator system, resulting in slow thermal dynamics and delayed control responses that make long-horizon optimization challenging. The controller regulates the radiator valve opening $u_1 \in [0,1]$, which

Table 4: Inputs, measurements, and objective evaluation for different baseline methods in the synthetic experiments.

Method	Phase	Input	Measurement	Surrogate objective
vvBO ($M = I$)	one	x	$f(x)$	$\langle \Xi_1 w_1, f(x) \rangle_{\mathcal{Y}}$
	two	x	$f(x)$	$\langle \Xi_2 w_2, f(x) \rangle_{\mathcal{Y}}$
	three	x	$f(x)$	$\langle \Xi_3 w_3, f(x) \rangle_{\mathcal{Y}}$
vvBO ($M = \Xi^*$)	one	x	$\Xi_1^* f(x)$	$\langle w_1, \Xi_1^* f(x) \rangle_{\mathbb{R}^5}$
	two	x	$\Xi_2^* f(x)$	$\langle w_2, \Xi_2^* f(x) \rangle_{\mathbb{R}^5}$
	three	–	–	–
MTBO	one	x	$\Xi_1^* f(x)$	$\langle w_1, \Xi_1^* f(x) \rangle_{\mathbb{R}^5}$
	two	x	$\Xi_1^* f(x)$	$\langle w_1, \Xi_1^* f(x) \rangle_{\mathbb{R}^5}$
	three	x	$\Xi_1^* f(x)$	$\langle w_1, \Xi_1^* f(x) \rangle_{\mathbb{R}^5}$
rMTBO	one	x	$\Xi_1^* f(x)$	$\langle w_1, \Xi_1^* f(x) \rangle_{\mathbb{R}^5}$
	two	x	$\Xi_2^* f(x)$ (re-initialized)	$\langle w_2, \Xi_2^* f(x) \rangle_{\mathbb{R}^5}$
	three	x	$\Xi_3^* f(x)$ (re-initialized)	$\langle w_3, \Xi_3^* f(x) \rangle_{\mathbb{R}^5}$
BO	one	x	$F_1(x)$	$F_1(x)$
	two	x	$F_2(x)$	$F_2(x)$
	three	x	$F_3(x)$	$F_3(x)$
rBO	one	x	$F_1(x)$	$F_1(x)$
	two	x	$F_2(x)$ (re-initialized)	$F_2(x)$
	three	x	$F_3(x)$ (re-initialized)	$F_3(x)$
CTBO	one	$(x, \Xi_1 w_1)$	$F_1(x)$	$\tilde{F}(x, \Xi_1 w_1)$
	two	$(x, \Xi_2 w_2)$	$F_2(x)$	$\tilde{F}(x, \Xi_2 w_2)$
	three	$(x, \Xi_3 w_3)$	$F_3(x)$	$\tilde{F}(x, \Xi_3 w_3)$
FFBO	one	$(x, \Xi_1 w_1)$	$F_1(x)$	$\tilde{F}(x, \Xi_1 w_1)$
	two	$(x, \Xi_1 w_1)$	$F_2(x)$	$\tilde{F}(x, \Xi_1 w_1)$
	three	$(x, \Xi_1 w_1)$	$F_3(x)$	$\tilde{F}(x, \Xi_1 w_1)$

determines the heat supplied to the zone through the district heating loop. Further details on this test case can be found in Yang et al. (2020).

Controller design. We employ a model predictive control (MPC) scheme to regulate the indoor temperature. The system dynamics are modeled using an autoregressive with exogenous input (ARX) model:

$$y_{k+1} = \boldsymbol{a}^\top \boldsymbol{y}_{k-10:k} + \sum_{i=1}^{3} \boldsymbol{b}_i^\top \boldsymbol{u}_{i,k-10:k},$$

where y is the indoor temperature, and u_1, u_2, u_3 correspond to the valve position, outdoor temperature, and solar irradiation. In this formulation, only the valve position is controlled. The outdoor temperature and solar irradiation are read directly from the BOPTEST API. The system model coefficients are obtained by solving a least-squares problem using data collected from a randomized bang-bang controller, i.e., the radiator is turned on when the temperature falls below a lower bound, and turned off when it exceeds an upper bound, and randomly actuated otherwise.

The optimization problem for the MPC controller is then formulated as

$$\begin{aligned}
\min_{\boldsymbol{u}_1, \boldsymbol{\epsilon}} \quad & \sum_{k=0}^{N-1} \left(u_{1,k} + \theta_2 \epsilon_k + \theta_3 (u_{1,k} - u_{1,k-1})^2 \right) \\
\text{s.t.} \quad & y_{k+1} = \boldsymbol{a}^\top \boldsymbol{y}_{k-10:k} + \sum_{i=1}^{3} \boldsymbol{b}_i^\top \boldsymbol{u}_{i,k-10:k}, \\
& \theta_1 - \epsilon_k \le y_{k+1} \le 26 + \epsilon_k, \qquad k = 0, \dots, N-1, \\
& u_{1,k} \in [0, 1], \qquad \epsilon_k \ge 0, \qquad k = 0, \dots, N-1.
\end{aligned}$$

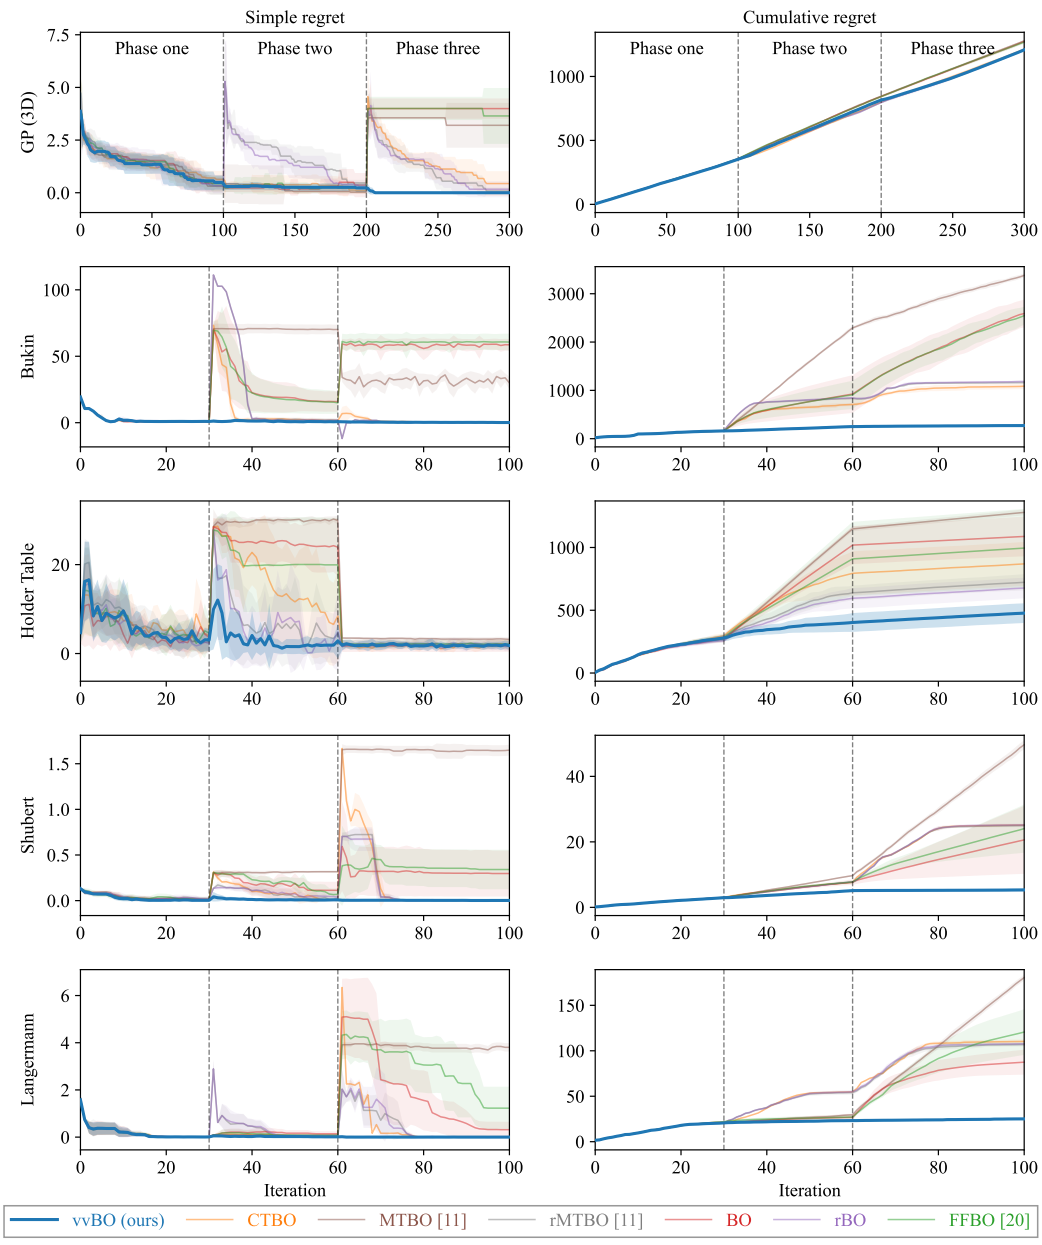


Figure 4: Comparison of simple regret and cumulative regret for additional test operators under full observations ($M = I$).

where $N = 64$ is the horizon and ϵ_k is the slack variable to ensure feasibility. The MPC controller is parameterized by $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3)$, corresponding to the comfort bound, slack penalty, and smoothness penalty. The controller operates with a sampling time of 900 s, meaning that the control input is updated every 15 minutes. We evaluate the controller on a daily basis. For a fixed $\boldsymbol{\theta}$, the controller regulates the indoor temperature over a 24-hour window, producing the heating trajectory $Q(t)$ and the indoor temperature trajectory $y(t)$, where t denotes the time of day in seconds.

Objective. The goal is to optimize the controller parameters by minimizing the daily cost

$$\mathcal{J}_\tau(\boldsymbol{\theta}) = \int Q(t) p_{\text{heating},\tau}(t)\, dt + 0.1 \int Q(t) p_{\text{CO}_2}(t)\, dt + 100 \left(\frac{1}{|\mathcal{T}|} \int y(t)\, dt - 24 \right)^2,$$

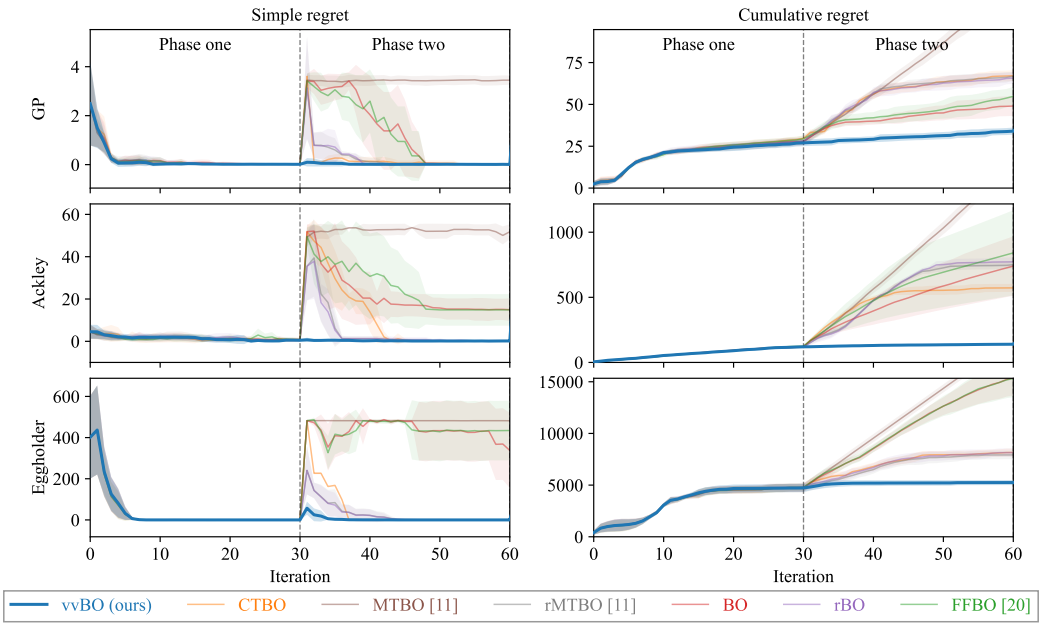


Figure 5: Comparison of simple regret and cumulative regret for three test operators under partial observations ($M = \Xi^*$).

where τ denotes the day index. These three components, which are nonlinear black-box functions with respect to $\boldsymbol{\theta}$, capture energy cost, environmental impact, and thermal discomfort, respectively. Let $f(\boldsymbol{\theta}) = (Q(t), y(t))$ represent the vector-valued system output. The first two terms correspond to linear functionals of $f(\boldsymbol{\theta})$, while the third term is a nonlinear functional of the temperature trajectory. The heating price $p_{\text{heating},\tau}(t)$ varies across days following a sinusoidal pattern between 0 and 0.232 EUR/kWh, making the objective $\mathcal{J}_\tau(\boldsymbol{\theta})$ inherently time-varying.

Learning setup. The mapping from $\boldsymbol{\theta}$ to trajectories $Q(t)$ and $y(t)$ is generally unknown. In vvBO, we model this mapping using a vector-valued RKHS model with a separable kernel, i.e., $K = G \otimes B$. Specifically, we choose the scalar kernel G as a radial basis function kernel and set B to be the identity operator on $\mathcal{Y}$. We assume the full trajectories $f(\boldsymbol{\theta})$ are available for measurement. For numerical implementation, trajectories are discretized over time and represented using the scalar-valued KRR estimator as discussed in Appendix J.1. The optimization proceeds sequentially. At the end of each day, the measurements are added to the dataset, and the model is updated. Given the forecasted prices for the next day, the algorithm selects parameters $\boldsymbol{\theta}$ that minimize $\mathcal{J}_\tau(\boldsymbol{\theta})$.

Measurements for baseline methods. All methods optimize the same time-varying objective $\mathcal{J}_\tau(\boldsymbol{\theta})$, but differ in the measurements available during learning. vvBO models the mapping from $\boldsymbol{\theta}$ to the full trajectories $f(\boldsymbol{\theta}) = (Q(t), y(t))$, and evaluates the objective through linear and nonlinear functionals of these trajectories. MTBO assumes access only to finite-dimensional task measurements, which correspond to energy consumption, environmental impact, and thermal discomfort. The surrogate objective is the weighted combination of these three measurements. BO directly observes the scalar objective value $\mathcal{J}_\tau(\boldsymbol{\theta})$. CTBO and FFBO treat the price signals as contextual variables and learn scalar surrogate models for the objective. CTBO allows the context to vary across iterations, while FFBO assumes a fixed context during learning. Table 5 summarizes the measurements available to each method.

Additional results. Fig. 9 shows the daily cost during both the learning and validation phases. In the first plot, we report the daily cost during the training phase. The black dashed line indicates the optimal cost for each day. After an initial exploration period, vvBO adjusts effectively to the time-varying objective (around iteration 150). In particular, when the heating price is high, vvBO selects parameters that reduce energy consumption by lowering the indoor temperature. Conversely,

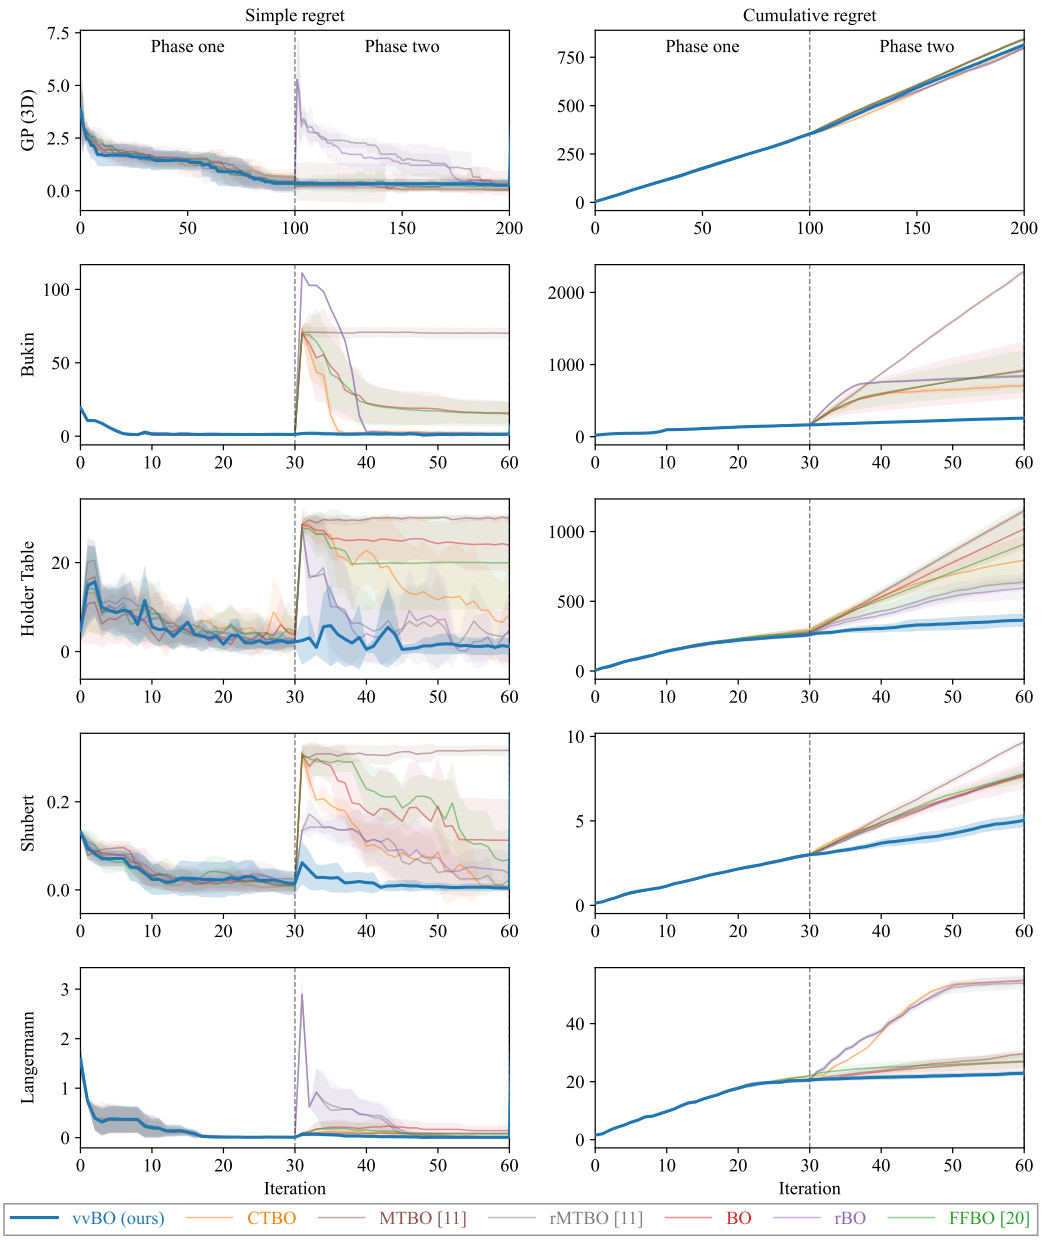


Figure 6: Comparison of simple regret and cumulative regret for additional test operators under partial observations ($M = \Xi^*$).

when the price is low, it increases the temperature to improve thermal comfort, thereby reducing the overall cost. In contrast, most baseline methods fail to track this optimal policy and exhibit significantly higher costs. The second plot shows the validation performance of the learned models. In this phase, no additional data are collected. Each method recommends parameters based on the trained model. vvBO consistently achieves the lowest cost across all days, demonstrating its ability to generalize and adjust to time-varying conditions. This highlights the advantage of leveraging structured measurements for efficient learning in real-world control problems.

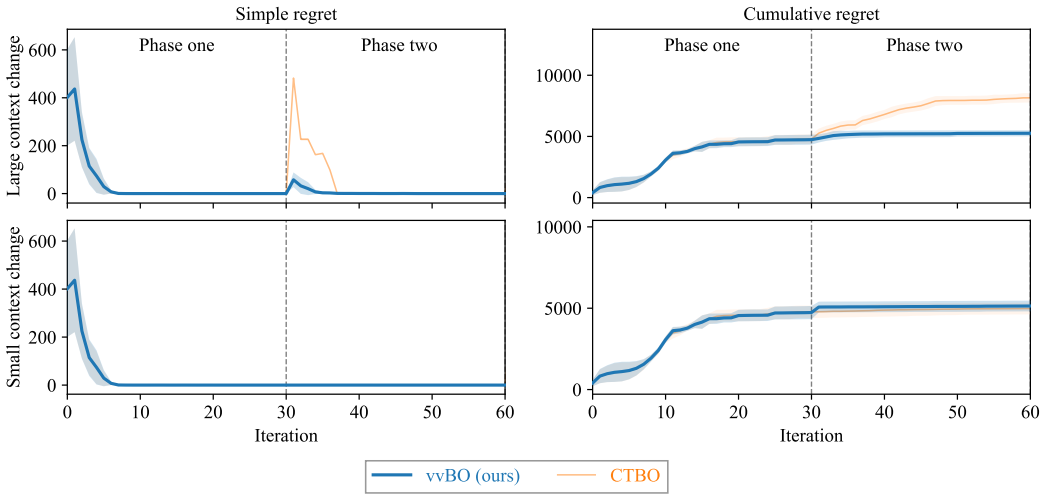


Figure 7: Comparison between vvBO and CTBO for the Eggholder test operator under different contexts. When the change in the context is small, CTBO has a lower exploration cost.

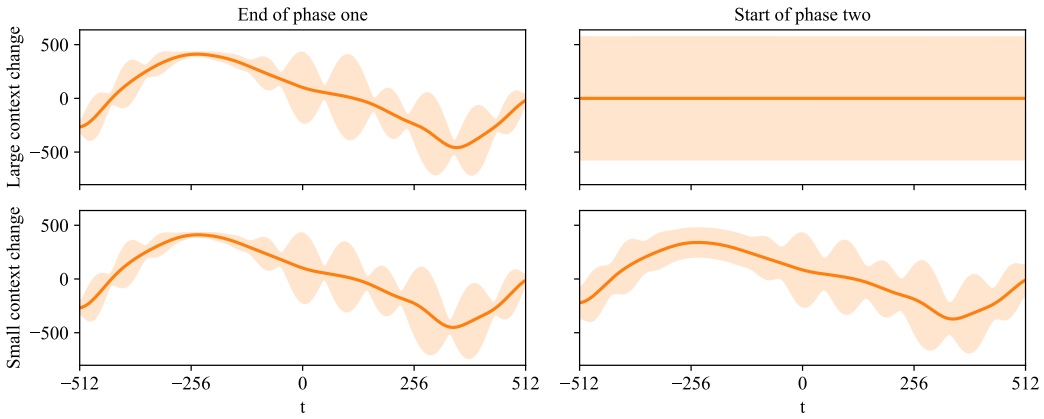


Figure 8: Comparison of the confidence bounds of CTBO for the Eggholder test operator with different objectives. When the change in the context is small, CTBO can formulate a more informative confidence bound.

Table 5: Measurements available to each method.

Method	Input	Measurements
vvBO	$\boldsymbol{\theta}$	$Q(t),\ y(t)$
MTBO	$\boldsymbol{\theta}$	energy cost, CO_2 cost, and thermal discomfort
BO	$\boldsymbol{\theta}$	$\mathcal{J}_\tau(\boldsymbol{\theta})$
CTBO	$(\boldsymbol{\theta},\ p_{\text{heating},\tau}(t),\ p_{CO_2}(t))$	$\mathcal{J}_\tau(\boldsymbol{\theta})$
FFBO	$(\boldsymbol{\theta},\ p_{\text{heating},0}(t),\ p_{CO_2}(t))$	$\mathcal{J}_\tau(\boldsymbol{\theta})$

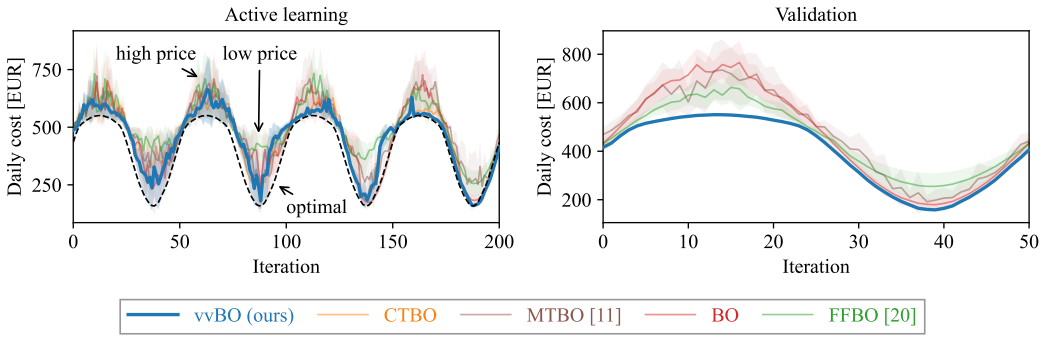


Figure 9: Daily cost during learning and validation phases.