

META-VBO: UTILIZING PRIOR TASKS IN OPTIMIZING RISK MEASURES WITH GAUSSIAN PROCESSES

Quoc Phong Nguyen¹, Bryan Kian Hsiang Low² & Patrick Jaillet¹

¹LIDS and EECS, Massachusetts Institute of Technology, USA

²School of Computing, National University of Singapore, Singapore

qphongmp@gmail.com, lowkh@comp.nus.edu.sg, jaillet@mit.edu

ABSTRACT

Research on optimizing the risk measure of a blackbox function using Gaussian processes, especially *Bayesian optimization* (BO) of risk measures, has become increasingly important due to the inevitable presence of uncontrollable variables in real-world applications. Nevertheless, existing works on BO of risk measures start the optimization from scratch for every new task without considering the results of prior tasks. In contrast, its vanilla BO counterpart has received a thorough investigation on utilizing prior tasks to speed up the current task through the body of works on *meta-BO* which, however, have not considered risk measures. To bridge this gap, this paper presents the first algorithm for *meta-BO of risk measures* (i.e., *value-at-risk* (VaR) and the *conditional VaR*), namely meta-VBO, by introducing a novel adjustment to the upper confidence bound acquisition function. Our proposed algorithm exhibits two desirable properties: (i) invariance to scaling and vertical shifting of the blackbox function and (ii) robustness to prior harmful tasks. We provide a theoretical performance guarantee for our algorithm and empirically demonstrate its performance using several synthetic function benchmarks and real-world objective functions.

1 INTRODUCTION

Real world optimization problems, such as maximizing the performance of a manufacturing line, the crop yield of a farm, and the performance of *machine learning* (ML) models, often face with the challenge of dealing with blackbox objective functions (i.e., without closed-form expressions, convexity, or derivatives) and time-consuming or expensive noisy zeroth-order evaluations (Snoek et al., 2012; Nguyen et al., 2021b). These issues have attracted significant attention, particularly on the *Bayesian optimization* (BO) framework (Srinivas et al., 2010; Snoek et al., 2012; Shahriari et al., 2015; Hernández-Lobato et al., 2014; Wang and Jegelka, 2017; Garnett, 2022). Recently, there has been an increasing interest in *BO of risk measures*, driven by the need to address practical aspects related to uncontrollable variables referred to as *environmental random variables* (Cakmak et al., 2020), e.g., random errors in sensors of the manufacturing line, weather conditions in farms (Nguyen et al., 2021b;a), and random perturbation in training data of ML models (Bogunovic et al., 2018). BO of risk measure algorithms specifically search for the controllable *decision variables* (e.g., applied torque, input force, nutrient amounts, and ML model hyperparameters) that optimize a *risk measure* of the blackbox function where the risk is induced by the environmental random variables. Risk measures investigated in existing works include the worst-case value (Bogunovic et al., 2018), *value-at-risk* (VaR) and *conditional VaR* (CVaR) (Cakmak et al., 2020; Nguyen et al., 2021b;a; Picheny et al., 2022), probability threshold robustness (Iwazaki et al., 2021a), and the variance (Iwazaki et al., 2021b; Makarova et al., 2021).

While there is an extensive research on BO of risk measures, using learning experiences from *prior tasks* to accelerate the optimization of the *current task* has not been explored. In contrast, this approach has been shown to effectively reduce the number of costly observations required by the *vanilla BO* counterpart through numerous works on *meta-BO*. These works encompass a wide spectrum of methodologies, ranging from those that involve the transfer of the probabilistic models from prior tasks to the current one (Swersky et al., 2013), to more nuanced approaches centered on transferring the optimization result (Feurer et al., 2018; Wistuba et al., 2016; 2018; Perrone et al.,

2019; Ramachandran et al., 2019; Volpp et al., 2020; Dai et al., 2022). Although the former category, such as the multi-task *Gaussian processes* (GP) (Swersky et al., 2013), is well-defined, it often suffers from computational challenges as discussed by Dai et al. (2022) and is not specifically tailored to the goal of optimization - the blackbox function’s maximizer. In contrast, the latter category, which attracts the majority of research attention, presents greater challenges and lacks a clearly defined framework. For example, Feurer et al. (2018) transfers the ranking of function evaluations. Perrone et al. (2019) narrows down the search space. Ramachandran et al. (2019) considers the distribution of the optima. Dai et al. (2022) considers a weighted average of the acquisition functions.

In this paper, we focus on the challenge of transferring the optimization outcome. Specifically, we pioneer the first solution to the *meta-BO of risk measure* problem, namely *meta-VBO*. Our work accounts for the risk introduced by the environmental random variables - a crucial aspect not addressed in any existing meta-BO works. We focus on two well-known risk measures: *value-at-risk* (VaR) and *conditional VaR* (CVaR) (Cakmak et al., 2020; Torossian et al., 2020). Moreover, our solution also extends to the adversarially robust BO (Bogunovic et al., 2018) and applies to the meta-BO problem. We assume that the distribution of the environmental random variables is given, e.g., by estimating from historical data or domain knowledge. Additionally, we assume that during the optimization, noisy evaluations of the blackbox function can be acquired (queried) at any realizations of the environmental random variables. In the above examples, this can be achieved running simulations of manufacturing lines with controlled random errors or growing indoor plants in a laboratory environment where weather conditions are regulated. After the optimization is completed, we recommend implementing the best decision variable in the real-world environment where the environmental random variables are beyond our control.¹ These assumptions are commonly made in several BO works with risk measures (Cakmak et al., 2020; Nguyen et al., 2021b;a).

Under the above settings, there are several major challenges in addressing meta-BO of risk measures. **First**, the choice of the information to be transferred from *prior tasks* to the *current task* should be tailored to the optimization goal of determining the optima’s locations. Since they are invariant under scaling and vertical-shifting (i.e., increasing or decreasing the function by a constant) of the blackbox function, it is crucial to ensure that the proposed meta-BO of risk measure solution maintains this invariance. Hence, it rules out the use of the function gap (i.e., the maximum difference between function evaluations) in the theoretical study of Dai et al. (2022) as a similarity measure between tasks. **Second**, as the number of prior tasks grows, the possibility of encountering a task that hampers the current optimization task arises (Feurer et al., 2018; Dai et al., 2022). Thus, it is essential to be resilient against such prior harmful tasks.

Moreover, one cannot easily extend existing meta-BO solutions to address the BO problem with risk measures. While the posterior belief of the blackbox function follows a GP in the vanilla BO, many risk measures including VaR and CVaR do not follow any well-known stochastic process. Hence, it is challenging to adopt approaches based on the Gaussian posterior distribution in the works of Ramachandran et al. (2019) and Dai et al. (2022). Furthermore, risk measures are not directly observed. They are only estimated from noisy evaluations of the blackbox function. As a result, although Feurer et al. (2018) is able to discover and assign low weights to prior harmful tasks by proposing a ranking loss function, extending this loss function to tackle risk measures is not immediately obvious. Last, there are not any existing meta-BO works that possess both theoretical performance guarantees and robustness against scaling and vertical-shifting of the blackbox function and prior harmful tasks.

To address the above problems, we introduce a new perspective in constructing our solution, namely *meta-VBO*. Specifically, while the classic *Gaussian process - upper confidence bound* (GP-UCB) algorithm (Srinivas et al., 2010) offers only one option of the input query at each BO iteration, we pursue a different approach of constructing a *versatile query set* (V-set) in Sec. 3.1. Interestingly, we can obtain a no-regret algorithm by selecting any input in this set as the input query (Theorem 3.2). To the best of our knowledge, this characteristic has not been introduced in any BO algorithm. On one hand, it implies that by restricting the choice of the input query to V-set, the algorithm is no-regret regardless of prior harmful tasks. On the other hand, the flexibility in choosing the input query from a set (as opposed to one option in GP-UCB) allows us to transfer information from prior tasks to the current task by assigning different preferences/priorities to inputs in V-set (Sec. 3.2). Essentially, inputs that are likely to be local (or global) maximizers of many prior tasks are assigned high priorities.

¹The notion of the best decision variable is discussed in Sec. 3.3.

The size of V-set can be tuned via two hyperparameters, which allows a trade-off between the amount of information to be incorporated from prior tasks and the cumulative regret bound of the current task. This reflects the intuition that without knowing the similarity between prior tasks and the current task, utilizing more information from prior tasks may not benefit the current task. We justify the selected hyperparameters in the experiments based on the incurred regret. Furthermore, V-set is adaptively updated as observations are accumulated from one iteration to the next. This naturally adjusts the effect of prior tasks on the current task as discussed in Remark 3.4, ensuring meta-VBO’s robustness to prior harmful tasks. In Sec. 3.3, we propose a strategy to recommend an input that approximates the optimal solution at each iteration such that its instantaneous regret diminishes towards zero as the iteration progresses (Theorem 3.7). Last, we empirically evaluate the performance of meta-VBO in optimizing VaR and CVaR of several synthetic and real-world objective functions in Sec. 4.

2 PRELIMINARY

2.1 BO OF RISK MEASURES

We are interested in maximizing a blackbox objective function f (e.g., the performance of a system) that depends on both the *decision variable* $\mathbf{x} \in \mathcal{X}$ and the *environmental variable* $\mathbf{z} \in \mathcal{Z}$, where \mathcal{X} and \mathcal{Z} are bounded subsets of \mathbb{R}^d and $\mathbb{R}^{d'}$, respectively. We assume that \mathcal{X} and \mathcal{Z} are discrete for simplicity, though the results can be extended to continuous domains (Srinivas et al., 2010; Chowdhury and Gopalan, 2017). Furthermore, in the real-world environment, \mathbf{x} is controllable while \mathbf{z} (e.g., weather conditions and random errors) is not. Hence, the environmental variable is represented as a random variable, denoted as \mathbf{Z} . We assume that the distribution of \mathbf{Z} is known, e.g., by estimating from historical data or domain knowledge. These settings are adopted from the existing works of Cakmak et al. (2020); Nguyen et al. (2021b;a).

Due to the randomness in \mathbf{Z} , evaluating the blackbox function at a particular value of \mathbf{x} can result in a high value under one realization of \mathbf{Z} but leading to a low value under another realization of \mathbf{Z} . In other words, the evaluation of f at \mathbf{x} is a random variable, denoted as $f(\mathbf{x}, \mathbf{Z})$, whose randomness is induced by \mathbf{Z} . Therefore, to determine the “optimal” design variable \mathbf{x}_* , one cannot directly compare the function evaluations at different values of \mathbf{x} . An alternative approach is to consider a risk measure of the distribution of $f(\mathbf{x}, \mathbf{Z})$ which accounts for the probability (or the risk) that $f(\mathbf{x}, \mathbf{Z})$ suffers from poor realizations. In this paper, we examine two risk measures for $f(\mathbf{x}, \mathbf{Z})$: *value-at-risk* (VaR), denoted as $v_f(\mathbf{x}; \alpha)$, and *conditional VaR* (CVaR), denoted as $c_f(\mathbf{x}; \alpha)$. Here, $\alpha \in (0, 1)$ represents a known *risk level*. VaR and CVaR are defined as follows:

$$v_f(\mathbf{x}; \alpha) \triangleq \inf\{\gamma : P(f(\mathbf{x}, \mathbf{Z}) \leq \gamma) \geq \alpha\} \quad c_f(\mathbf{x}; \alpha) \triangleq \frac{1}{\alpha} \int_0^\alpha v_f(\mathbf{x}; \alpha') d\alpha'. \quad (1)$$

While $v_f(\mathbf{x}; \alpha)$ can be viewed as 100 α %-percentile of $f(\mathbf{x}, \mathbf{Z})$ (i.e., the event that $f(\mathbf{x}, \mathbf{Z})$ suffers from realizations lower than $v_f(\mathbf{x}; \alpha)$ happens with a probability of at most α), $c_f(\mathbf{x}; \alpha)$ is the expected value of $f(\mathbf{x}, \mathbf{Z})$ in the worst 100 α % of cases. To avoid unnecessary repetition, we unify the notations of both VaR and CVaR as $\rho_f(\mathbf{x}; \alpha)$. The notations $v_f(\mathbf{x}; \alpha)$ and $c_f(\mathbf{x}; \alpha)$ are only used when VaR and CVaR require distinct treatments.

Given the risk measure $\rho_f(\mathbf{x}; \alpha)$ and the distribution of \mathbf{Z} , the *BO of risk measure* problem aims to search for the maximizer $\mathbf{x}_* \triangleq \arg \max_{\mathbf{x} \in \mathcal{X}} \rho_f(\mathbf{x}; \alpha)$ (Cakmak et al., 2020). Since f is unknown, determining this maximizer typically involves iteratively observing noisy evaluations $y(\mathbf{x}_t, \mathbf{z}_t) \triangleq f(\mathbf{x}_t, \mathbf{z}_t) + \epsilon(\mathbf{x}_t, \mathbf{z}_t)$ for iteration $t \geq 1$ where $\epsilon(\mathbf{x}_t, \mathbf{z}_t) \sim \mathcal{N}(0, \sigma_n^2)$ are i.i.d. Gaussian noises.² Furthermore, it is often expensive to obtain these evaluations (e.g., requiring running expensive simulations). Thus, the crux of the problem lies in devising a selection strategy of the *input query* $(\mathbf{x}_t, \mathbf{z}_t)$ at iteration t , considering observations in the previous $t - 1$ iterations, denoted as $\mathbf{y}_{\mathbf{D}_t} \triangleq (y(\mathbf{x}_{t'}, \mathbf{z}_{t'}))_{t'=1}^{t-1}$ where $\mathbf{D}_t \triangleq ((\mathbf{x}_{t'}, \mathbf{z}_{t'}))_{t'=1}^{t-1}$, to quickly discover \mathbf{x}_* .

We model f with a *Gaussian process* (GP) (Rasmussen and Williams, 2006). Specifically, we use a constant GP prior mean function μ_0 and a squared exponential kernel $k_{(\mathbf{x}, \mathbf{z}), (\mathbf{x}', \mathbf{z}')} \triangleq \text{cov}[f(\mathbf{x}, \mathbf{z}), f(\mathbf{x}', \mathbf{z}')] = \sigma_s^2 \exp(-0.5(\mathbf{x} - \mathbf{x}')^\top \Lambda_x^{-2}(\mathbf{x} - \mathbf{x}') - 0.5(\mathbf{z} - \mathbf{z}')^\top \Lambda_z^{-2}(\mathbf{z} - \mathbf{z}'))$ where $\Lambda_x \triangleq$

²We adopt the settings outlined in previous works of Cakmak et al. (2020); Nguyen et al. (2021a) such that \mathbf{z} is controllable during the optimization (discussed in Sec. 1).

$\text{diag}(l_1, \dots, l_d)$, $\Lambda_z \triangleq \text{diag}(l_{d+1}, \dots, l_{d+d'})$, and σ_s^2 are the lengthscales and the signal variance, respectively. Given $\mathbf{y}_{\mathbf{D}_t}$, the posterior belief of $f(\mathbf{x}, \mathbf{z})$ is a Gaussian distribution with the following mean and variance

$$\mu_t(\mathbf{x}, \mathbf{z}) \triangleq \mu_0 + \mathbf{K}_{(\mathbf{x}, \mathbf{z}), \mathbf{D}_t} \Lambda_{\mathbf{D}_t} (\mathbf{y}_{\mathbf{D}_t} - \mu_0) \quad \sigma_t^2(\mathbf{x}, \mathbf{z}) \triangleq k_{(\mathbf{x}, \mathbf{z})} - \mathbf{K}_{(\mathbf{x}, \mathbf{z}), \mathbf{D}_t} \Lambda_{\mathbf{D}_t} \mathbf{K}_{\mathbf{D}_t, (\mathbf{x}, \mathbf{z})}$$

where $\Lambda_{\mathbf{D}_t} \triangleq (\mathbf{K}_{\mathbf{D}_t, \mathbf{D}_t} + \sigma_n^2 \mathbf{I})^{-1}$, \mathbf{K} is the covariance matrix, and \mathbf{I} is the identity matrix.

The performance of a BO algorithm is often analyzed through the *cumulative regret*, denoted as $R_T \triangleq \sum_{t=1}^T r(\mathbf{x}_t)$ where $r(\mathbf{x}_t) \triangleq \rho_f(\mathbf{x}_*; \alpha) - \rho_f(\mathbf{x}_t; \alpha)$ is the *instantaneous regret* at the input query \mathbf{x}_t . The algorithm is *no-regret* if its cumulative regret is sublinear, i.e., $\lim_{T \rightarrow \infty} R_T/T \rightarrow 0$. In such a case, the *simple regret* $S_T \triangleq \min_{t=1, \dots, T} r(\mathbf{x}_t) \leq R_T/T$ approaches 0. This suggests that one should recommend the best observed input so far as the outcome of the optimization, e.g., in the vanilla BO problem with a small noise. Unfortunately, in BO of risk measures, $\rho_f(\mathbf{x}_t; \alpha)$ is unobserved due to the unobserved $f(\mathbf{x}_t, \mathbf{z})$ for $(\mathbf{x}_t, \mathbf{z}) \notin \mathbf{D}_t$. This prevents us from determining the best observed input. Therefore, we will design a recommendation strategy such that the instantaneous regret $r(\mathbf{x}_t^*)$ at the recommended input \mathbf{x}_t^* approaches zero as $t \rightarrow \infty$ in Sec. 3.3.

2.2 META-BO OF RISK MEASURES

Unlike existing works on BO of risk measures, our approach in this paper considers a set \mathcal{T} of prior tasks. Let $f : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$ denote the blackbox function of the *current task* and $f_\tau : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$ denote the blackbox function in a prior task $\tau \in \mathcal{T}$. Furthermore, to facilitate later theoretical analysis, we follow the works of Chowdhury and Gopalan (2017) to assume that all blackbox functions belong to the *reproducing kernel Hilbert space* (RKHS) associated with the kernel k and a bounded RKHS norm, i.e., $\|f\|_k \leq B$ and $\forall \tau \in \mathcal{T}$, $\|f_\tau\|_k \leq B$. The optimization result of a prior task $\tau \in \mathcal{T}$ includes the GP posterior mean and variance of f_τ , denoted as $\hat{\mu}_\tau : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$ and $\hat{\sigma}_\tau^2 : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}^+$, respectively. Meta-BO of the risk measure ρ aims to utilize $(\hat{\mu}_\tau, \hat{\sigma}_\tau^2)_{\tau \in \mathcal{T}}$ to improve the optimization of ρ_f . However, the similarity between f and f_τ (hence, between ρ_f and ρ_{f_τ}) is unknown for all $\tau \in \mathcal{T}$. In particular, if ρ_{f_τ} differs from ρ_f , naively utilizing $\hat{\mu}_\tau, \hat{\sigma}_\tau^2$ to optimize ρ_f can deteriorate the optimization of ρ_f .

Under the assumption that blackbox functions belong to the RKHS space with a norm bounded by B , we first revisit the well-known confidence bound of $f(\mathbf{x}, \mathbf{z})$ from the work of Chowdhury and Gopalan (2017). Let γ_t denote the maximum information gain about f that can be obtained from any set of $t - 1$ observations. Choosing $\delta \in (0, 1)$, $\beta_t = (B + \sigma_n \sqrt{2(\gamma_t + 1 + \log 1/\delta)})^2$, then the event that $\forall t \geq 1, \forall \mathbf{x} \in \mathcal{X}, \forall \mathbf{z} \in \mathcal{Z}$, $l_t(\mathbf{x}, \mathbf{z}) \leq f(\mathbf{x}, \mathbf{z}) \leq u_t(\mathbf{x}, \mathbf{z})$ holds with probability $\geq 1 - \delta$ where the lower and upper confidence bounds are defined as

$$l_t(\mathbf{x}, \mathbf{z}) \triangleq \mu_t(\mathbf{x}, \mathbf{z}) - \beta_t^{1/2} \sigma_t(\mathbf{x}, \mathbf{z}) \quad \text{and} \quad u_t(\mathbf{x}, \mathbf{z}) \triangleq \mu_t(\mathbf{x}, \mathbf{z}) + \beta_t^{1/2} \sigma_t(\mathbf{x}, \mathbf{z}), \text{ resp.} \quad (2)$$

Leveraging on the above confidence bounds l_t and u_t , Nguyen et al. (2021b;a) propose the following confidence interval of the risk measure $\rho_f(\mathbf{x}; \alpha)$ (applicable to both VaR and CVaR): with probability $\geq 1 - \delta$,

$$\forall t \geq 1, \forall \mathbf{x} \in \mathcal{X}, \rho_f(\mathbf{x}; \alpha) \in [\rho_{l_t}(\mathbf{x}; \alpha), \rho_{u_t}(\mathbf{x}; \alpha)]. \quad (3)$$

where l_t and u_t are defined in equation (2) and $\rho_f, \rho_{l_t}, \rho_{u_t}$ can be replaced by v_f, v_{l_t} , and v_{u_t} for VaR or by c_f, c_{l_t}, c_{u_t} for CVaR in equation (1). For a prior task $\tau \in \mathcal{T}$, we denote the confidence interval of $f_\tau(\mathbf{x}, \mathbf{z})$ as $[\hat{l}_\tau(\mathbf{x}, \mathbf{z}), \hat{u}_\tau(\mathbf{x}, \mathbf{z})]$ which is defined in equation (2) using the GP posterior belief $(\hat{\mu}_\tau, \hat{\sigma}_\tau^2)$ of f_τ and the index of the last optimization iteration of task τ . Then, the confidence interval of the risk measure $\rho_{f_\tau}(\mathbf{x}; \alpha)$ of prior task τ is $[\rho_{\hat{l}_\tau}(\mathbf{x}; \alpha), \rho_{\hat{u}_\tau}(\mathbf{x}; \alpha)]$.

3 OUR PROPOSED SOLUTION TO META-BO OF RISK MEASURES

Typically, at each iteration of the GP-UCB algorithm (Srinivas et al., 2010), there is only one option to select the input query from (assuming a unique maximizer of the GP-UCB acquisition function). Hence, in order to incorporate the information from prior tasks, the existing RM-GP-UCB (Dai et al., 2022) takes a weighted average of the GP-UCB acquisition functions of the current task and prior tasks. However, these weights depend on the maximum difference between f and f_τ evaluations,




Figure 1: (a) Plot of $\mathcal{V}_t = \{\mathbf{x}_0, \mathbf{x}_t^+, \mathbf{x}_2\}$ with $\lambda = 0$ and $\eta = 1$. The dashed line shows the confidence interval of $\rho_f(\mathbf{x}_*; \alpha)$. While \mathbf{x}_t^- and \mathbf{x}_1 violate UC, \mathbf{x}_3 violates OC. (b) Plot of an informative prior task τ where $\mathcal{V}_t \setminus \mathcal{P}_{\tau|\mathcal{V}_t} \neq \emptyset$. (c) Plot of an uninformative prior task τ where $\mathcal{P}_{\tau|\mathcal{V}_t} = \mathcal{V}_t$.

known as the *function gap*, at observed inputs. Defining such a gap notion for risk measures is not obvious. Moreover, the function gap is susceptible to function transformations such as scaling and vertical-shifting. Last, in order to achieve theoretical performance guarantee, RM-GP-UCB requires the weights of prior tasks to decay to zero.

In contrast, our meta-VBO approach tackles the problem from a different perspective that not only exhibits robustness against scaling and vertical-shifting transformations but also eliminates the explicit need for diminishing the effect of prior tasks via weight decay. Specifically, in the next section, we offer multiple choices of the input query in the form of a *versatile query set* (V-set), denoted as \mathcal{V}_t (depending on iteration t), such that any input query sequence in the Cartesian product $\prod_{t=1}^T \mathcal{V}_t$ incurs a sublinear cumulative regret,

$$\forall (\mathbf{x}_t)_{t=1}^T \in \prod_{t=1}^T \mathcal{V}_t, \quad \lim_{T \rightarrow \infty} \frac{1}{T} R_T = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T r(\mathbf{x}_t) = 0. \quad (4)$$

The flexibility of choosing any input query $\mathbf{x}_t \in \mathcal{V}_t$ grants us the opportunity to incorporate information from prior tasks into the current optimization task as elaborated in Sec. 3.2.

3.1 VERSATILE QUERY SET

Intuitively, the *versatile query set* (V-set) \mathcal{V}_t should (C1) contain those inputs that are likely to be the maximizer of ρ_f . Furthermore, we would like to avoid over-exploitation at inputs with well-estimated risk measures where gathering observations leads to a minimal progress towards discovering \mathbf{x}_* . This is done by introducing another criterion (C2) to exclude inputs with well-estimated risk measures from \mathcal{V}_t . These two key criteria (C1) and (C2) are formulated as follows.

(C1) To determine whether an input is likely to be the maximizer of ρ_f , we obtain the following confidence interval of $\rho_f(\mathbf{x}_*; \alpha)$ in App. A. With probability $\geq 1 - \delta$,

$$\forall t \geq 1, \rho_f(\mathbf{x}_*; \alpha) \in [\rho_{l_t}(\mathbf{x}_t^-; \alpha), \rho_{u_t}(\mathbf{x}_t^+; \alpha)], \quad (5)$$

where $\mathbf{x}_t^- \triangleq \arg \max_{\mathbf{x} \in \mathcal{X}} \rho_{l_t}(\mathbf{x}; \alpha)$ and $\mathbf{x}_t^+ \triangleq \arg \max_{\mathbf{x} \in \mathcal{X}} \rho_{u_t}(\mathbf{x}; \alpha)$. It is noted that \mathbf{x}_t^+ is the input query at iteration t in V-UCB (Nguyen et al., 2021b) and CV-UCB (Nguyen et al., 2021a). For an input \mathbf{x}' to be the maximizer \mathbf{x}_* of ρ_f (i.e., the above condition (C1)), the confidence interval of its risk measure $\rho_f(\mathbf{x}'; \alpha)$ should overlap with that of the maximum risk measure $\rho_f(\mathbf{x}_*; \alpha)$, i.e., $\rho_{u_t}(\mathbf{x}'; \alpha) \geq \rho_{l_t}(\mathbf{x}_t^-; \alpha)$ from equation (3) and equation (5). Hence, we propose to formulate (C1) as the following *overlapping condition* (OC): For an input \mathbf{x}' to be in \mathcal{V}_t ,

$$\rho_{u_t}(\mathbf{x}'; \alpha) \geq \rho_{l_t}(\mathbf{x}_t^-; \alpha) + \lambda (\rho_{u_t}(\mathbf{x}_t^+; \alpha) - \rho_{l_t}(\mathbf{x}_t^-; \alpha)) \quad \text{for } \lambda \in [0, 1]. \quad (6)$$

As $\lambda \rightarrow 1$, $\rho_{u_t}(\mathbf{x}'; \alpha) \rightarrow \rho_{u_t}(\mathbf{x}_t^+; \alpha)$, i.e., the overlap between the confidence interval of $\rho_f(\mathbf{x}'; \alpha)$ and that of $\rho_f(\mathbf{x}_*; \alpha)$ at iteration t increases. When $\lambda = 1$, the set of inputs satisfying OC reduce to only the input query \mathbf{x}_t^+ of V-UCB (Nguyen et al., 2021b) or CV-UCB (Nguyen et al., 2021a).

(C2) We would like to exclude inputs with well-estimated risk measures by ensuring that the uncertainty of their risk measures is relatively high compared to that of the maximum risk measure $\rho_f(\mathbf{x}_*; \alpha)$. Specifically, we propose to formulate (C2) as the following *uncertainty condition*

(UC): For an input \mathbf{x}' to be in \mathcal{V}_t ,

$$\underbrace{\rho_{u_t}(\mathbf{x}'; \alpha) - \rho_{l_t}(\mathbf{x}'; \alpha)}_{\text{uncertainty of } \rho_f(\mathbf{x}'; \alpha)} \geq \frac{1}{\eta} \left(\underbrace{\rho_{u_t}(\mathbf{x}_t^+; \alpha) - \rho_{l_t}(\mathbf{x}_t^-; \alpha)}_{\text{uncertainty of } \rho_f(\mathbf{x}_*; \alpha)} \right) \quad \text{for } \eta \in [1, 1/\lambda]. \quad (7)$$

If $\eta > 1/\lambda$, OC implies UC since $\rho_{l_t}(\mathbf{x}_t^-; \alpha) \geq \rho_{l_t}(\mathbf{x}; \alpha) \forall \mathbf{x} \in \mathcal{X}$ (from the definition of \mathbf{x}_t^-). Hence, we are only interested in UC when $\eta \leq 1/\lambda$ as specified in equation (7).

To sum up, we define V-set \mathcal{V}_t as the set of inputs $\mathbf{x} \in \mathcal{X}$ that satisfy both OC in equation (6) and UC in equation (7) (illustrated in Fig. 1a). The conditions of $\lambda \leq 1$ and $\eta \geq 1$ ensure \mathcal{V}_t is always non-empty which follows from the following Lemma 3.1 (proved in App. B).

Lemma 3.1. *For all $t \geq 1$, \mathcal{V}_t contains \mathbf{x}_t^+ .*

We note that the size of \mathcal{V}_t is non-decreasing when λ decreases or η increases. Specifically, $\lambda = 1$ implies $\eta = 1$ (since $\eta \in [1, 1/\lambda]$) and OC implies UC. In this case, $\mathcal{V}_t = \{\mathbf{x}_t^+\}$ (assuming ρ_{u_t} has a unique maximizer), i.e., \mathcal{V}_t reduces to the UCB-typed input query such as those in V-UCB (Nguyen et al., 2021b) and CV-UCB (Nguyen et al., 2021a). There is not any flexibility in choosing \mathbf{x}_t from $\mathcal{V}_t = \{\mathbf{x}_t^+\}$, which prevents us from incorporating the information of prior tasks into the current task. However, when $\lambda < 1$ and $\eta \in [1, 1/\lambda]$, \mathcal{V}_t contains other inputs than the popular input query \mathbf{x}_t^+ of UCB-typed algorithms. This is a trade-off between prior tasks (existing knowledge) and the current task (new knowledge). When $\lambda = 1$ ($\mathcal{V}_t = \{\mathbf{x}_t^+\}$), we utilize only observations from the current task. Conversely, decreasing λ (increasing $|\mathcal{V}_t|$) promotes the exploitation/utilization of prior tasks to select an input query in \mathcal{V}_t . This *task-specific* trade-off is distinct from the conventional (*function-evaluation-specific*) exploration-exploitation trade-off associated with estimating the maximizer in V-UCB (Nguyen et al., 2021b) and CV-UCB (Nguyen et al., 2021a) (pertaining to a particular task).

Given the selected $\mathbf{x}_t \in \mathcal{V}_t$, we adopt the choice of \mathbf{z}_t in the existing works of Nguyen et al. (2021b;a). In particular, if the risk measure is VaR, then \mathbf{z}_t is selected as a *lacing value* (LV) w.r.t. \mathbf{x}_t, l_t, u_t , and α .³ If the risk measure is CVaR, we first find the risk level $\alpha_t \triangleq \arg \max_{\alpha' \in (0, \alpha]} v_{u_t}(\mathbf{x}_t; \alpha') - v_{l_t}(\mathbf{x}_t; \alpha')$. Then, \mathbf{z}_t is selected as an LV w.r.t. \mathbf{x}_t, l_t, u_t , and α_t .

Interestingly, given the above choice of \mathbf{z}_t , any choice of \mathbf{x}_t in our proposed \mathcal{V}_t results in a sublinear cumulative regret as shown in the following Theorem 3.2 (proved in App. D). In other words, we achieve the goal of constructing \mathcal{V}_t that satisfies equation (4).

Theorem 3.2. *Given the selection of \mathbf{x}_t as any input in \mathcal{V}_t and the selection of \mathbf{z}_t from the works of Nguyen et al. (2021b;a), the algorithm achieves a sublinear cumulative regret:*

$$\forall T \geq 1, R_T \leq (\eta(1 - \lambda) + 1) \sqrt{C_1 T \beta_T \gamma_T} \quad (8)$$

holds with probability $\geq 1 - \delta$ where $C_1 \triangleq 8 / \log(1 + \sigma_n^{-2})$; $\gamma_T, \beta_T, \delta$ are defined in equation (2); λ and η are defined in equation (6) and equation (7), respectively.

Given the flexibility of choosing \mathbf{x}_t in \mathcal{V}_t , we are now able to transfer the information from prior tasks to the current task by assigning different priorities to inputs in \mathcal{V}_t .

3.2 ASSIGNING PRIORITY TO INPUTS IN VERSATILE QUERY SET \mathcal{V}_t

Since selecting any input query $\mathbf{x}_t \in \mathcal{V}_t$ results in a no-regret algorithm (Theorem 3.2), we would like to exploit the results of prior tasks to assign varying degrees of importance (i.e., *priorities*) to inputs in \mathcal{V}_t . To be robust to scaling and vertical-shifting of the blackbox function, the priority is based on the location of the maximizer rather than the estimated risk measure value. Besides, to provide more opportunities for transferring information between tasks, we consider the case that the global maximizer of the current task is the local maximizer of a prior task. Hence, we assign higher priorities to inputs that have a higher chance of being a local maximizer (restricted to \mathcal{V}_t) of prior tasks. For this reason, we propose the notion of *the probable local maximizer set*, denoted as $\mathcal{P}_{\tau|\mathcal{V}_t}$, which contains inputs that are probably a local maximizer restricted to \mathcal{V}_t of ρ_{f_τ} :

$$\mathcal{P}_{\tau|\mathcal{V}_t} \triangleq \{\mathbf{x} \in \mathcal{V}_t \mid \rho_{\bar{u}_\tau}(\mathbf{x}; \alpha) \geq \max_{\mathbf{x}' \in \mathcal{V}_t} \rho_{\bar{l}_\tau}(\mathbf{x}'; \alpha)\}. \quad (9)$$

³Lacing values \mathbf{z}_{LV} w.r.t. \mathbf{x}, l_t, u_t , and α satisfy $l_t(\mathbf{x}, \mathbf{z}_{LV}) \leq v_{l_t}(\mathbf{x}; \alpha) \leq v_{u_t}(\mathbf{x}; \alpha) \leq u_t(\mathbf{x}, \mathbf{z}_{LV})$ (Nguyen et al., 2021b).

An alternative assumption is that the global maximizer of the current task is close to that of a prior task, then we can consider *the probable global maximizer set*, defined as $\mathcal{P}_{\tau|\mathcal{X}} \triangleq \{\mathbf{x} \in \mathcal{X} \mid \rho_{\hat{u}_\tau}(\mathbf{x}; \alpha) \geq \max_{\mathbf{x}' \in \mathcal{X}} \rho_{\hat{l}_\tau}(\mathbf{x}'; \alpha)\}$. Since using $\mathcal{P}_{\tau|\mathcal{X}}$ does not substantially change the implementation of meta-VBO in comparison to using $\mathcal{P}_{\tau|\mathcal{V}_t}$, we only focus on the probable local maximizer set $\mathcal{P}_{\tau|\mathcal{V}_t}$ in this paper. Figs. 1b and 1c show the set $\mathcal{P}_{\tau|\mathcal{V}_t}$ of inputs $\mathbf{x} \in \mathcal{V}_t$ such that $\rho_{\hat{u}_\tau}(\mathbf{x}; \alpha)$ is above the dash-dotted line representing $\max_{\mathbf{x}' \in \mathcal{V}_t} \rho_{\hat{l}_\tau}(\mathbf{x}'; \alpha)$. Essentially, we transfer the information about the *local maximizer* restricted to \mathcal{V}_t of prior tasks to the current task. This is an unexplored direction in the current meta-BO literature.

Given the probable local maximizer set $\mathcal{P}_{\tau|\mathcal{V}_t}$, the priority of an input $\mathbf{x} \in \mathcal{V}_t$, denoted as $\pi_t(\mathbf{x})$, is measured by the number of the probable local maximizer sets that \mathbf{x} belongs to, i.e.,

$$\pi_t(\mathbf{x}) \triangleq \sum_{\tau \in \mathcal{T}} \mathbb{1}_{\mathbf{x} \in \mathcal{P}_{\tau|\mathcal{V}_t}}. \quad (10)$$

The higher $\pi_t(\mathbf{x})$ is, the more likely that \mathbf{x} is in the probable local maximizer set of a prior task. Hence, we select the input query \mathbf{x}_t as

$$\mathbf{x}_t \triangleq \max_{\mathbf{x} \in \mathcal{V}_{t,\max}} \rho_{u_t}(\mathbf{x}; \alpha) \text{ where } \mathcal{V}_{t,\max} \triangleq \{\mathbf{x} \in \mathcal{V}_t \mid \pi_t(\mathbf{x}) = \max_{\mathbf{x}' \in \mathcal{V}_t} \pi_t(\mathbf{x}')\}. \quad (11)$$

Remark 3.3 (Prior harmful tasks). The local maximizers of ρ_{f_τ} in a prior harmful task restricted to \mathcal{V}_t are different from that of ρ_f of the current task. Contrary to the intuition suggesting that prior harmful tasks may bias meta-VBO away from the desirable maximizer \mathbf{x}_* , Theorem 3.2 asserts that our algorithm maintains a sublinear cumulative regret by constraining the choice of \mathbf{x}_t to the \mathcal{V}_t . However, one could reasonably anticipate a potential reduction in the convergence rate, as empirically shown in Sec. 4, due to the influence of prior harmful tasks.

Remark 3.4 (Informativeness of a prior task). If $\pi_t(\mathbf{x})$ is constant for all $\mathbf{x} \in \mathcal{V}_t$, it results in $\mathcal{V}_{t,\max} = \mathcal{V}_t$ and $\mathbf{x}_t = \mathbf{x}_t^+$. This happens when there are not any prior tasks or when the BO results of prior tasks are *uninformative* in the local region \mathcal{V}_t , i.e., $\mathcal{P}_{\tau|\mathcal{V}_t} = \mathcal{V}_t \forall \tau \in \mathcal{T}$. As $\mathbf{x}_t = \mathbf{x}_t^+$ in these cases, our meta-VBO algorithm reduces to V-UCB (Nguyen et al., 2021b) or CV-UCB (Nguyen et al., 2021a). Fig. 1c illustrates an iteration t where prior task τ is uninformative as $\mathcal{V}_t = \mathcal{P}_{\tau|\mathcal{V}_t}$, while Fig. 1b shows the case that τ is an informative task. It is noted that the informativeness of a prior task w.r.t. the current task depends on \mathcal{V}_t , hence, the iteration t . A prior task may be informative when we have not gathered many observations on the current task (\mathcal{V}_t is large such that $\mathcal{P}_{\tau|\mathcal{V}_t}$ is a strict subset of \mathcal{V}_t), but it may become uninformative when many observations are acquired (\mathcal{V}_t is small such that $\mathcal{P}_{\tau|\mathcal{V}_t} = \mathcal{V}_t$). When a prior task is uninformative, it does not influence the optimization of the current task in our meta-VBO algorithm (from the definition of $\pi_t(\mathbf{x})$ in equation (40)). Hence, the effect of a prior task on the current task is adaptively controlled by \mathcal{V}_t as illustrated in Appendix I.1. This approach contrasts with that of Dai et al. (2022) that explicitly requires a schedule of the weight decay based on the function gap.

Remark 3.5 (Scaling and vertical-shifting of the blackbox functions). Our proposed meta-VBO depends on \mathcal{V}_t and $\mathcal{P}_{\tau|\mathcal{V}_t}$. These sets rely on comparing risk measures at different inputs which remains unchanged under scaling or vertical-shifting of the blackbox function. Hence, our algorithm is robust to these transformations.

Remark 3.6 (Meta-VBO for vanilla BO and adversarially robust BO). As our proposed meta-VBO depends on the upper and lower confidence bounds of function evaluations, it is sufficiently general to address the vanilla meta-BO problem (see Appendix F). Its potential applicability to the meta-learning of other UCB-typed algorithms remains an area for future exploration. Furthermore, since adversarially robust BO is a special case of BO of VaR $\alpha \rightarrow 0^+$ (Nguyen et al., 2021b), our proposed meta-VBO directly applies to adversarially robust BO (see App. G).

3.3 DECISION VARIABLE RECOMMENDATION STRATEGY

At iteration t , we propose to recommend $\mathbf{x}_t^* \triangleq \mathbf{x}_{\zeta(t)}^-$ where $\zeta(t) \triangleq \arg \max_{t' \in \{1, \dots, t\}} \rho_{l_{t'}}(\mathbf{x}_{t'}^-; \alpha)$. We prove that its instantaneous regret approaches zero in App. E.

Theorem 3.7. *Let $\zeta(t) \triangleq \arg \max_{t' \in \{1, \dots, t\}} \rho_{l_{t'}}(\mathbf{x}_{t'}^-; \alpha)$, then, by recommending $\mathbf{x}_t^* = \mathbf{x}_{\zeta(t)}^-$, the instantaneous regret at \mathbf{x}_t^* is bounded: $r(\mathbf{x}_t^*) \leq \eta \sqrt{C_1 \beta_t \gamma_t / t}$ with probability $\geq 1 - \delta$ where η , C_1 , β_t , and γ_t are defined in Theorem 3.2.*

As sublinear bounds of γ_t are established for common kernels such as the SE kernel in the work of Srinivas et al. (2010), the instantaneous regret at $\mathbf{x}_t^* \triangleq \mathbf{x}_{\zeta(t)}^-$ approaches zero as $t \rightarrow \infty$. Furthermore, we should choose $\eta = 1$ to minimize the upper confidence bound of $r(\mathbf{x}_t^*)$ in Theorem 3.7. Given $\eta = 1$, we choose $\lambda \in [0, 1]$ to trade off between the flexibility in choosing \mathbf{x}_t (i.e., the size of V-set) and the upper confidence bound of the cumulative regret R_T at the input queries in Theorem 3.2. In the experiments, as we want to exploit the solutions of prior tasks, we maximize the size of V-set by setting $\lambda = 0$. In this case, the upper confidence bound of R_T is two times that of V-UCB (Nguyen et al., 2021b) and CV-UCB (Nguyen et al., 2021a): $R_T \leq 2\sqrt{C_1 T \beta_T \gamma_T}$.

4 EXPERIMENTS

Given the absence of prior research on meta-BO of risk measures, we compare the performance of our proposed algorithm under different sets of prior tasks to that of V-UCB (Nguyen et al., 2021b) which is equivalent to our proposed algorithm when there are not any prior tasks, i.e., $\mathcal{T} = \emptyset$. Recall that our emphasis in this paper is on transferring the optimization results, rather than transferring the surrogate model, as discussed in Sec. 1. Therefore, we exclusively focus on independent GPs. They are also more practical due to their lower computational overhead, while still enabling us to provide fair comparisons to empirically demonstrate our advantage over V-UCB.⁴ Our goal is to empirically verify whether our proposed algorithm is (i) able to exploit the information from prior informative tasks to improve its converge rate in comparison to V-UCB even if harmful tasks exist in \mathcal{T} and (ii) robust to prior harmful tasks even if there are not any useful tasks in \mathcal{T} . To achieve this, we design different sets of prior tasks: $\mathcal{T}_{\text{useful-pos-scale}}$ (i.e., positive scaling of the blackbox function of the current task: $f_{\tau}(\mathbf{x}, \mathbf{z}) = af(\mathbf{x}, \mathbf{z})$ for $a > 0$), $\mathcal{T}_{\text{useful-vshift}}$ (i.e., vertical-shifting $f_{\tau}(\mathbf{x}, \mathbf{z}) = f(\mathbf{x}, \mathbf{z}) + b$ for $b \in \mathbb{R}$), $\mathcal{T}_{\text{harmful-neg-scale}}$ (i.e., negative scaling $f_{\tau}(\mathbf{x}, \mathbf{z}) = -f(\mathbf{x}, \mathbf{z})$), and $\mathcal{T}_{\text{harmful-hshift}}$ (i.e., horizontal-shifting $f_{\tau}(\mathbf{x}, \mathbf{z}) = f(\mathbf{x} + \xi, \mathbf{z})$ for $\xi \in \mathbb{R}^d$). Additionally, \mathcal{T}_{all} is the union of the above 4 sets. We note that $\mathcal{T}_{\text{harmful-neg-scale}}$ contains the most harmful task as the function evaluations are negated. On the other hand, horizontal-shifting along the \mathbf{x} dimensions in $\mathcal{T}_{\text{harmful-hshift}}$ moves the location of the maximizers away from that of the current task. As explained in Sec. 3.3, we set $\lambda = 0$ and $\eta = 1$. We present additional experimental results of meta-BO for CVaR in App. I.

The experiments are performed on 6 synthetic functions: a Gaussian curve, the Branin-Hoo, the Goldstein-Price, the six-hump camel, the Hartmann-3D, and the Hartmann-6D (obtained from <https://www.sfu.ca/~ssurjano>); and using 2 real-world datasets: the yacht hydrodynamics dataset (Dua and Graff, 2017) and a portfolio optimization dataset (obtained from the existing works of Cakmak et al. (2020); Nguyen et al. (2021a)). In the first 4 experiments, both the dimensions of \mathbf{x} and \mathbf{z} are set to $d = d' = 1$. In the Hartmann-3D experiment, $d = 2$ and $d' = 1$, while in the Hartmann-6D experiment, $d = 5$ and $d' = 1$. In the yacht hydrodynamics experiment, the optimization problem is to minimize the VaR of the residuary resistance per unit weight of displacement of a yacht by controlling the its 5-dimensional hull geometry coefficients. The environmental random variable is the 1-dimensional Froude number. The blackbox function is generated by fitting a GP to the yacht hydrodynamics dataset (Dua and Graff, 2017). In the portfolio optimization experiment, we follow the work of (Cakmak et al., 2020) to optimize a portfolio by controlling $d = 3$ decision variables subject to the randomness in the bid-ask spread and the borrow cost (i.e., $d' = 2$). The blackbox function is generated by fitting a GP to the simulated portfolio dataset. The risk level is set at $\alpha = 0.1$ for all experiments.

The average and the standard error of $r(\mathbf{x}_t^*)$ (in Sec. 3.3) over 30 random repetitions are shown in Fig. 2. It is noted that the curves of `useful-pos-scale`, `useful-vshift`, and `all` overlap each other in Figs. 2a, 2e, and 2g. In these figures, we observe that when \mathcal{T} contains only prior useful tasks in $\mathcal{T}_{\text{useful-pos-scale}}$ and $\mathcal{T}_{\text{useful-vshift}}$, our proposed meta-VBO outperforms V-UCB which does not utilize any prior tasks. Notably, even when \mathcal{T} contains both harmful and useful tasks in \mathcal{T}_{all} , meta-VBO still outperforms V-UCB. In the extreme case when \mathcal{T} contains only harmful tasks in $\mathcal{T}_{\text{harmful-neg-scale}}$ and $\mathcal{T}_{\text{harmful-hshift}}$, meta-VBO does not outperform V-UCB, but it still converges to a low-regret solution. In Fig. 2f, meta-VBO given the set $\mathcal{T}_{\text{harmful-neg-scale}}$ performs poorly probably due to the large input space and the function surface of the Hartmann-6D. We can

⁴When computation cost are affordable, one may combine a method that transfers the surrogate model with our algorithm by simply replacing independent GPs with a multitask GP because our algorithm only relies on upper and lower confidence bounds that are readily obtainable from a multitask GP model.

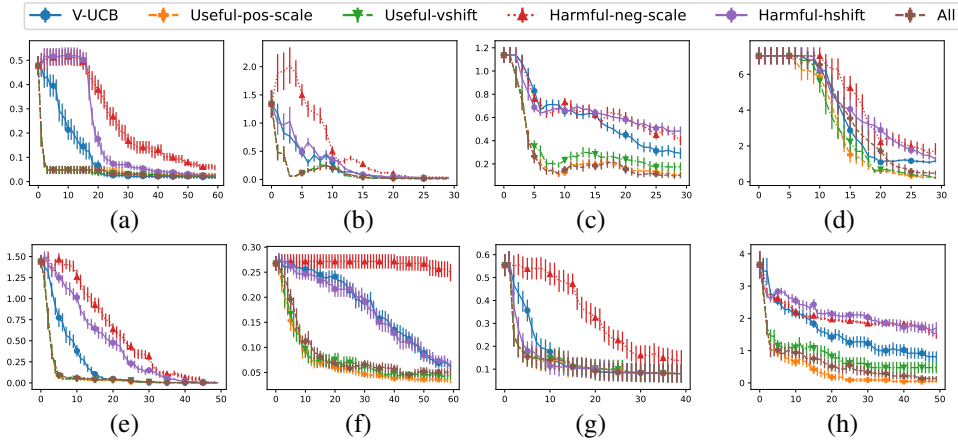


Figure 2: Plots of the regret at the recommended input against the BO iteration in experiments with (a) a Gaussian curve, (b) the Branin-Hoo function, (c) the Goldstein-Price function, (d) the six-hump camel function, (e) the Hartmann-3D function, (f) the Hartmann-6D function, (g) the yacht hydrodynamics dataset, and (h) the portfolio optimization dataset.




Figure 3: Percentage of prior tasks whose probable local maximizer sets contain the input query (grouped by the transformation) at different BO iterations in the experiment with the Gaussian curve and the set \mathcal{T}_{all} of prior tasks. The percentage is computed over 30 repeated experiments.

also observe the percentage of prior tasks whose probable maximizer sets $\mathcal{P}_{\tau|\mathcal{V}_t}$ containing \mathbf{x}_t in Fig. 3. It shows that as we gather more observations, the percentage of harmful tasks whose $\mathcal{P}_{\tau|\mathcal{V}_t}$ contains \mathbf{x}_t reduces.

5 CONCLUSION

In this paper, we propose the first algorithm for meta-BO of risk measures by proposing a new concept of V-set and leveraging prior tasks to prioritize inputs in this set. Our algorithm offers a theoretical performance guarantee even when prior tasks are harmful. Furthermore, we propose a recommendation strategy that suggests an input as the result of the optimization, with a proof of its instantaneous regret approaching zero. The empirical performance of the proposed solution is demonstrated through several synthetic function benchmarks and real-world objective functions. In the future, we would like to extend our meta-VBO approach to perform meta-learning for other BO variants such as preferential BO (González et al., 2017; Nguyen et al., 2021c), BO for finding Nash equilibria (Tay et al., 2023), and BO with unknown constraints (Nguyen et al., 2023).

REPRODUCIBILITY STATEMENT

To ensure the reproducibility of both the theoretical and experimental results in our paper, we have clearly elaborated on our assumptions, the proofs of the theoretical results, and provided the implementation of the experiments. In particular, our main assumptions are stated in the third paragraph in Sec. 1 and the first paragraph in Sec. 3. The details proofs of our theoretical results are shown in Appendices A, B, D, E, and G.

ACKNOWLEDGEMENT

This research is supported by AI Singapore, under grant AISG2-RP-2020-018.

This research is supported by the National Research Foundation (NRF), Prime Minister’s Office, Singapore under its Campus for Research Excellence and Technological Enterprise (CREATE) programme. The Mens, Manus, and Machina (M3S) is an interdisciplinary research group (IRG) of the Singapore MIT Alliance for Research and Technology (SMART) centre.

REFERENCES

- I. Bogunovic, J. Scarlett, S. Jegelka, and V. Cevher. Adversarially robust optimization with Gaussian processes. In *Proc. NeurIPS*, pages 5760–5770, 2018.
- S. Cakmak, R. Astudillo, P. I. Frazier, and E. Zhou. Bayesian optimization of risk measures. In *Proc. NeurIPS*, 2020.
- S. R. Chowdhury and A. Gopalan. On kernelized multi-armed bandits. In *Proc. ICML*, pages 844–853, 2017.
- Z. Dai, Y. Chen, H. Yu, B. K. H. Low, and P. Jaillet. On provably robust meta-Bayesian optimization. In *Proc. UAI*, pages 475–485, 2022.
- Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- M. Feurer, B. Letham, and E. Bakshy. Scalable meta-learning for Bayesian optimization using ranking-weighted Gaussian process ensembles. In *AutoML Workshop at ICML*, volume 7, 2018.
- Roman Garnett. *Bayesian Optimization*. Cambridge University Press, 2022. in preparation.
- J. González, Z. Dai, A. Damianou, and N. D. Lawrence. Preferential Bayesian optimization. In *Proc. ICML*, pages 1282–1291, 2017.
- J. M. Hernández-Lobato, M. W. Hoffman, and Z. Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In *Proc. NIPS*, pages 918–926, 2014.
- S. Iwazaki, Y. Inatsu, and I. Takeuchi. Bayesian quadrature optimization for probability threshold robustness measure. *Neural Computation*, 33(12):3413–3466, 2021a.
- S. Iwazaki, Y. Inatsu, and I. Takeuchi. Mean-variance analysis in Bayesian optimization under uncertainty. In *Proc. AISTAT*, pages 973–981, 2021b.
- A. Makarova, I. Usmanova, I. Bogunovic, and A. Krause. Risk-averse heteroscedastic Bayesian optimization. In *Proc. NeurIPS*, pages 17235–17245, 2021.
- Q. P. Nguyen, Z. Dai, B. K. H. Low, and P. Jaillet. Optimizing conditional value-at-risk of black-box functions. In *Proc. NeurIPS*, pages 4170–4180, 2021a.
- Q. P. Nguyen, Z. Dai, B. K. H. Low, and P. Jaillet. Value-at-risk optimization with Gaussian processes. In *Proc. ICML*, pages 8063–8072, 2021b.
- Q. P. Nguyen, S. Tay, B. K. H. Low, and P. Jaillet. Top-k ranking Bayesian optimization. In *Proc. AAAI*, volume 35, pages 9135–9143, 2021c.
- Q. P. Nguyen, W. T. R. Chew, L. Song, B. K. H. Low, and P. Jaillet. Optimistic Bayesian optimization with unknown constraints. In *Proc. ICLR*, 2023.
- V. Perrone, H. Shen, M. W. Seeger, C. Archambeau, and R. Jenatton. Learning search spaces for Bayesian optimization: Another view of hyperparameter transfer learning. In *Proc. NeurIPS*, volume 32, 2019.
- V. Picheny, H. Moss, L. Torossian, and N. Durrande. Bayesian quantile and expectile optimisation. In *Proc. UAI*, 2022.

- A. Ramachandran, S. Gupta, S. Rana, and S. Venkatesh. Information-theoretic transfer learning framework for Bayesian optimisation. In *Proc. ECML/PKDD*, pages 827–842, 2019.
- C. E. Rasmussen and C. K. I. Williams. *Gaussian processes for machine learning*. MIT Press, 2006.
- B. Shahriari, K. Swersky, Z. Wang, R. Adams, and N. de Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2015.
- J. Snoek, H. Larochelle, and R. Adams. Practical Bayesian optimization of machine learning algorithms. In *Proc. NIPS*, pages 2951–2959, 2012.
- N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proc. ICML*, pages 1015–1022, 2010.
- K. Swersky, J. Snoek, and R. P. Adams. Multi-task Bayesian optimization. In *Proc. NeurIPS*, volume 26, 2013.
- S. S. Tay, Q. P. Nguyen, C. S. Foo, and B. K. H. Low. No-regret sample-efficient Bayesian optimization for finding Nash equilibria with unknown utilities. In *Proc. AISTATS*, pages 3591–3619, 2023.
- L. Torossian, V. Picheny, and N. Durrande. Bayesian quantile and expectile optimisation. *arXiv preprint arXiv:2001.04833*, 2020.
- M. Volpp, L. P. Fröhlich, K. Fischer, A. Doerr, S. Falkner, F. Hutter, and C. Daniel. Meta-learning acquisition functions for transfer learning in Bayesian optimization. In *Proc. ICLR*, 2020.
- Z. Wang and S. Jegelka. Max-value entropy search for efficient Bayesian optimization. In *Proc. ICML*, pages 3627–3635, 2017.
- M. Wistuba, N. Schilling, and L. Schmidt-Thieme. Two-stage transfer surrogate model for automatic hyperparameter optimization. In *Proc. ECML/PKDD*, pages 199–214, 2016.
- M. Wistuba, N. Schilling, and L. Schmidt-Thieme. Scalable Gaussian process-based transfer surrogates for hyperparameter optimization. *Machine Learning*, 107(1):43–78, 2018.

A CONFIDENCE INTERVAL OF $\rho_f(\mathbf{x}_*; \alpha)$

In this section, we prove a confidence interval of $\rho_f(\mathbf{x}_*; \alpha)$:

$$\forall t \geq 1, \rho_f(\mathbf{x}_*; \alpha) \in [\rho_{l_t}(\mathbf{x}_t^-; \alpha), \rho_{u_t}(\mathbf{x}_t^+; \alpha)]$$

holds with probability $\geq 1 - \delta$.

Proof. Let us assume that $\rho_f(\mathbf{x}; \alpha)$ is in its confidence interval, i.e., the event $\forall t \geq 1, \forall \mathbf{x} \in \mathcal{X}, \rho_f(\mathbf{x}; \alpha) \in [\rho_{l_t}(\mathbf{x}; \alpha); \rho_{u_t}(\mathbf{x}; \alpha)]$ holds (which happens with probability $\geq 1 - \delta$). Furthermore, we define the following inputs in equation (5).

$$\mathbf{x}_t^- \triangleq \arg \max_{\mathbf{x} \in \mathcal{X}} \rho_{l_t}(\mathbf{x}; \alpha) \quad \mathbf{x}_t^+ \triangleq \arg \max_{\mathbf{x} \in \mathcal{X}} \rho_{u_t}(\mathbf{x}; \alpha).$$

Then, $\forall \mathbf{x} \in \mathcal{X}, \forall t \geq 1, \rho_f(\mathbf{x}_*; \alpha) \geq \rho_f(\mathbf{x}; \alpha) \geq \rho_{l_t}(\mathbf{x}; \alpha)$. Hence,

$$\begin{aligned} \forall t \geq 1, \rho_f(\mathbf{x}_*; \alpha) &\geq \max_{\mathbf{x} \in \mathcal{X}} \rho_{l_t}(\mathbf{x}; \alpha) \\ \text{i.e., } \forall t \geq 1, \rho_f(\mathbf{x}_*; \alpha) &\geq \rho_{l_t}(\mathbf{x}_t^-; \alpha). \end{aligned} \quad (12)$$

Furthermore,

$$\forall t \geq 1, \rho_f(\mathbf{x}_*; \alpha) \leq \rho_{u_t}(\mathbf{x}_*; \alpha) \leq \max_{\mathbf{x} \in \mathcal{X}} \rho_{u_t}(\mathbf{x}; \alpha) = \rho_{u_t}(\mathbf{x}_t^+; \alpha). \quad (13)$$

From equation (12) and equation (13),

$$\forall t \geq 1, \rho_f(\mathbf{x}_*; \alpha) \in [\rho_{l_t}(\mathbf{x}_t^-; \alpha), \rho_{u_t}(\mathbf{x}_t^+; \alpha)].$$

□

B PROOF OF LEMMA 3.1

In this section, we prove that when $\lambda \leq 1$ and $\eta \geq 1$, $\mathbf{x}_t^+ \in \mathcal{V}_t$.

Proof. To show that $\mathbf{x}_t^+ \in \mathcal{V}_t$, we show that \mathbf{x}_t^+ satisfies both OC equation (6) and UC equation (7) conditions.

- **OC:** We observe that when $\lambda \leq 1$,

$$\begin{aligned} \rho_{u_t}(\mathbf{x}_t^+; \alpha) &= \rho_{l_t}(\mathbf{x}_t^-; \alpha) + (\rho_{u_t}(\mathbf{x}_t^+; \alpha) - \rho_{l_t}(\mathbf{x}_t^-; \alpha)) \\ &\geq \rho_{l_t}(\mathbf{x}_t^-; \alpha) + \lambda (\rho_{u_t}(\mathbf{x}_t^+; \alpha) - \rho_{l_t}(\mathbf{x}_t^-; \alpha)). \end{aligned}$$

Hence, \mathbf{x}_t^+ satisfies OC equation (6) when $\lambda \leq 1$.

- **UC:** We observe that when $\eta \geq 1$,

$$\rho_{u_t}(\mathbf{x}_t^+; \alpha) - \rho_{l_t}(\mathbf{x}_t^+; \alpha) \geq \rho_{u_t}(\mathbf{x}_t^+; \alpha) - \rho_{l_t}(\mathbf{x}_t^-; \alpha) \geq \frac{1}{\eta} (\rho_{u_t}(\mathbf{x}_t^+; \alpha) - \rho_{l_t}(\mathbf{x}_t^-; \alpha)).$$

Hence, \mathbf{x}_t^+ satisfies UC equation (7) when $\eta \geq 1$.

□

C REVIEW OF THE UPPER CONFIDENCE BOUND OF THE $\rho_{u_t}(\mathbf{x}_t; \alpha) - \rho_{l_t}(\mathbf{x}_t; \alpha)$

Recall that the *lacing value* (LV), denoted as \mathbf{z}_{LV} w.r.t. \mathbf{x} , l_t , u_t , and α , is defined as the input in \mathcal{Z} that satisfies the following property (Nguyen et al., 2021b)

$$l_t(\mathbf{x}, \mathbf{z}_{LV}) \leq v_{l_t}(\mathbf{x}; \alpha) \leq v_{u_t}(\mathbf{x}; \alpha) \leq u_t(\mathbf{x}, \mathbf{z}_{LV}) \quad (14)$$

The existence of LV is proved in (Nguyen et al., 2021b). It plays an important role in the choice of \mathbf{z}_t given \mathbf{x}_t in V-UCB (Nguyen et al., 2021b) and CV-UCB (Nguyen et al., 2021a). Specifically, given the input query \mathbf{x}_t , the works of Nguyen et al. (2021b;a) select \mathbf{z}_t as follows:

- If we are optimizing VaR of $f(\mathbf{x}, \mathbf{Z})$, \mathbf{z}_t is selected as an LV w.r.t. \mathbf{x}_t, l_t, u_t , and α (Nguyen et al., 2021b).
- If we are optimizing CVaR of $f(\mathbf{x}, \mathbf{Z})$, \mathbf{z}_t is selected as an LV w.r.t. \mathbf{x}_t, l_t, u_t , and $\alpha_t \triangleq \arg \max_{\alpha' \in (0, \alpha]} v_{u_t}(\mathbf{x}_t; \alpha') - v_{l_t}(\mathbf{x}_t; \alpha')$ (Nguyen et al., 2021a).

We repeat the derivations in the works of Nguyen et al. (2021b) and Nguyen et al. (2021a) to show that with probability $\geq 1 - \delta$, for any input query $\mathbf{x}_t \in \mathcal{X}$ and the above choice of \mathbf{z}_t ,

$$\rho_{u_t}(\mathbf{x}_t; \alpha) - \rho_{l_t}(\mathbf{x}_t; \alpha) \leq 2\beta_t^{1/2} \sigma_t(\mathbf{x}_t, \mathbf{z}_t). \quad (15)$$

It is noted that equation (15) does not require any specific choice of \mathbf{x}_t . Due to the different choices of \mathbf{z}_t in the optimization of VaR and CVaR, they are treated separately as follows.

C.1 OPTIMIZATION OF VALUE-AT-RISK

As \mathbf{z}_t is selected as an LV w.r.t. \mathbf{x}_t, l_t, u_t , and α , it follows from the definition of LV equation (14) that

$$v_{u_t}(\mathbf{x}_t; \alpha) - v_{l_t}(\mathbf{x}_t; \alpha) \leq u_t(\mathbf{x}_t, \mathbf{z}_t) - l_t(\mathbf{x}_t, \mathbf{z}_t) = 2\beta_t^{1/2} \sigma_t(\mathbf{x}_t, \mathbf{z}_t). \quad (16)$$

C.2 OPTIMIZATION OF CONDITIONAL VALUE-AT-RISK

From the definition of CVaR in Sec. 2.1.

$$c_{u_t}(\mathbf{x}_t; \alpha) - c_{l_t}(\mathbf{x}_t; \alpha) = \frac{1}{\alpha} \int_0^\alpha v_{u_t}(\mathbf{x}_t; \alpha') - v_{l_t}(\mathbf{x}_t; \alpha') d\alpha'.$$

Recall that $\alpha_t \triangleq \arg \max_{\alpha' \in (0, \alpha]} v_{u_t}(\mathbf{x}_t; \alpha') - v_{l_t}(\mathbf{x}_t; \alpha')$,

$$\begin{aligned} \forall \alpha' \in (0, \alpha], v_{u_t}(\mathbf{x}_t; \alpha') - v_{l_t}(\mathbf{x}_t; \alpha') &\leq v_{u_t}(\mathbf{x}_t; \alpha_t) - v_{l_t}(\mathbf{x}_t; \alpha_t) \\ \text{i.e., } \frac{1}{\alpha} \int_0^\alpha v_{u_t}(\mathbf{x}_t; \alpha') - v_{l_t}(\mathbf{x}_t; \alpha') d\alpha' &\leq \frac{1}{\alpha} \int_0^\alpha v_{u_t}(\mathbf{x}_t; \alpha_t) - v_{l_t}(\mathbf{x}_t; \alpha_t) d\alpha' \\ &= v_{u_t}(\mathbf{x}_t; \alpha_t) - v_{l_t}(\mathbf{x}_t; \alpha_t) \end{aligned}$$

Therefore,

$$c_{u_t}(\mathbf{x}_t; \alpha) - c_{l_t}(\mathbf{x}_t; \alpha) \leq v_{u_t}(\mathbf{x}_t; \alpha_t) - v_{l_t}(\mathbf{x}_t; \alpha_t)$$

Since \mathbf{z}_t is selected as an LV w.r.t. \mathbf{x}_t, l_t, u_t , and α_t ,

$$v_{u_t}(\mathbf{x}_t; \alpha_t) - v_{l_t}(\mathbf{x}_t; \alpha_t) \leq u_t(\mathbf{x}_t, \mathbf{z}_t) - l_t(\mathbf{x}_t, \mathbf{z}_t).$$

Hence,

$$\begin{aligned} c_{u_t}(\mathbf{x}_t; \alpha) - c_{l_t}(\mathbf{x}_t; \alpha) &\leq u_t(\mathbf{x}_t, \mathbf{z}_t) - l_t(\mathbf{x}_t, \mathbf{z}_t) \\ &= 2\beta_t^{1/2} \sigma_t(\mathbf{x}_t, \mathbf{z}_t). \end{aligned} \quad (17)$$

As a result, equation (15) follows from equation (16) and equation (17).


D PROOF OF THEOREM 3.2

First, we prove an upper confidence bound of the instantaneous regret $r(\mathbf{x}_t) \triangleq \rho_f(\mathbf{x}_*; \alpha) - \rho_f(\mathbf{x}_t; \alpha)$ in the following lemma.

Lemma D.1. *Given $\mathbf{x}_t \in \mathcal{V}_t$, i.e., \mathbf{x}_t satisfies both OC equation (6) and UC equation (7), and \mathbf{z}_t is selected as an LV w.r.t. \mathbf{x}_t, l_t, u_t , and α (if optimizing VaR) or $\alpha_t \triangleq \arg \max_{\alpha' \in (0, \alpha]} v_{u_t}(\mathbf{x}_t; \alpha') - v_{l_t}(\mathbf{x}_t; \alpha')$ (if optimizing CVaR), then*

$$\forall t \geq 1, r(\mathbf{x}_t) \leq 2(\eta(1 - \lambda) + 1) \beta_t^{1/2} \sigma_t(\mathbf{x}_t, \mathbf{z}_t)$$

holds with probability $\geq 1 - \delta$.

Figure 4: Decomposition an upper confidence bound of $r(\mathbf{x}_t)$.

Proof. Assuming that $\forall \mathbf{x} \in \mathcal{X}, \forall t \geq 1, \rho_f(\mathbf{x}; \alpha) \in [\rho_{l_t}(\mathbf{x}; \alpha), \rho_{u_t}(\mathbf{x}; \alpha)]$ which happens with probability $\geq 1 - \delta$.

We decompose an upper confidence bound of the instantaneous regret $r(\mathbf{x}_t)$ into two parts.

$$\begin{aligned}
r(\mathbf{x}_t) &\triangleq \rho_f(\mathbf{x}_*; \alpha) - \rho_f(\mathbf{x}_t) \\
&= \rho_{u_t}(\mathbf{x}_*) - \rho_{l_t}(\mathbf{x}_t; \alpha) \\
&\leq \rho_{u_t}(\mathbf{x}_t^+; \alpha) - \rho_{l_t}(\mathbf{x}_t; \alpha) \\
&= \underbrace{[\rho_{u_t}(\mathbf{x}_t^+; \alpha) - \rho_{u_t}(\mathbf{x}_t; \alpha)]}_{\text{irreducible (B)}} + \underbrace{[\rho_{u_t}(\mathbf{x}_t; \alpha) - \rho_{l_t}(\mathbf{x}_t; \alpha)]}_{\text{reducible (A)}}. \tag{18}
\end{aligned}$$

Hence, an upper confidence bound $\rho_{u_t}(\mathbf{x}_t^+; \alpha) - \rho_{l_t}(\mathbf{x}_t; \alpha)$ of $r(\mathbf{x}_t)$ is decomposed into (A) a *reducible* part and (B) an *irreducible* part. (A) is called the reducible part because gathering observations at \mathbf{x}_t reduces the uncertainty of $\rho_f(\mathbf{x}_t; \alpha)$, hence, reducing (A), while (B) is called the irreducible part because it is likely that gathering observations at \mathbf{x}_t does not reduce the (B) (as $\rho_f(\mathbf{x}_t; \alpha) \leq \rho_{u_t}(\mathbf{x}_t; \alpha)$ with high probability). As a result, the primary difficulty is in bounding this irreducible part (A).

Fortunately, we select $\mathbf{x}_t \in \mathcal{V}_t$ such that it satisfies both (OC) equation (6) and (UC) equation (7) conditions, so the irreducible part (B) can be associated with the reducible part (A) as follows.

$$\begin{aligned}
&\underbrace{\rho_{u_t}(\mathbf{x}_t^+; \alpha) - \rho_{u_t}(\mathbf{x}_t; \alpha)}_{\text{irreducible (B)}} \\
&= [\rho_{u_t}(\mathbf{x}_t^+; \alpha) - \rho_{l_t}(\mathbf{x}_t^-; \alpha)] - [\rho_{u_t}(\mathbf{x}_t; \alpha) - \rho_{l_t}(\mathbf{x}_t^-; \alpha)] \\
&\leq [\rho_{u_t}(\mathbf{x}_t^+; \alpha) - \rho_{l_t}(\mathbf{x}_t^-; \alpha)] - \lambda [\rho_{u_t}(\mathbf{x}_t^+; \alpha) - \rho_{l_t}(\mathbf{x}_t^-; \alpha)] \quad \text{from (OC) equation (6)} \\
&= (1 - \lambda) [\rho_{u_t}(\mathbf{x}_t^+; \alpha) - \rho_{l_t}(\mathbf{x}_t^-; \alpha)] \\
&\leq \eta(1 - \lambda) \underbrace{[\rho_{u_t}(\mathbf{x}_t; \alpha) - \rho_{l_t}(\mathbf{x}_t; \alpha)]}_{\text{reducible (A)}} \quad \text{from (UC) equation (7) and } \lambda \leq 1. \tag{19}
\end{aligned}$$

Plug equation (19) into equation (18),

$$\begin{aligned}
r(\mathbf{x}_t) &\leq \eta(1 - \lambda) [\rho_{u_t}(\mathbf{x}_t; \alpha) - \rho_{l_t}(\mathbf{x}_t; \alpha)] + [\rho_{u_t}(\mathbf{x}_t; \alpha) - \rho_{l_t}(\mathbf{x}_t; \alpha)] \\
&= (\eta(1 - \lambda) + 1) [\rho_{u_t}(\mathbf{x}_t; \alpha) - \rho_{l_t}(\mathbf{x}_t; \alpha)] \\
&\leq 2(\eta(1 - \lambda) + 1) \beta_t^{1/2} \sigma_t(\mathbf{x}_t, \mathbf{z}_t) \quad \text{from equation (15)}.
\end{aligned}$$

□

Remark D.2. The reader may observe that one intention behind of (OC) and (UC) is to ensure the relationship between the irreducible part (B) and the reducible part (A) in equation (19). Then, instead of imposing 2 conditions (OC) and (UC), he/she may suggest we directly define an alternative versatile query set $\tilde{\mathcal{Q}}_t$ by imposing this relationship, e.g.,

$$\tilde{\mathcal{Q}}_t \triangleq \{\mathbf{x} \in \mathcal{X} \mid \rho_{u_t}(\mathbf{x}_t^+; \alpha) - \rho_{u_t}(\mathbf{x}; \alpha) \leq \nu(\rho_{u_t}(\mathbf{x}; \alpha) - \rho_{l_t}(\mathbf{x}; \alpha))\} \tag{20}$$

for $\nu \geq 0$. However, the above definition equation (20) of $\tilde{\mathcal{Q}}_t$ is insufficient for the goal of optimizing $\rho_f(\mathbf{x}; \alpha)$. This is because $\tilde{\mathcal{Q}}_t$ contains those inputs $\mathbf{x} \in \mathcal{X}$ that do not satisfies (OC) equation (6),

i.e., $\rho_{u_t}(\mathbf{x}; \alpha) < \rho_{l_t}(\mathbf{x}_t^-; \alpha)$, which are unlikely to be the maximizer of $\rho_f(\mathbf{x}; \alpha)$. Hence, sampling at these inputs is inefficient as they may not provide any information on the maximizer \mathbf{x}_* . Furthermore, while (OC) and (UC) are interpretable as described in Sec. 3.1, interpreting the definition of $\tilde{\mathcal{Q}}_t$ in equation (20) is more challenging. As a result, the versatile query set \mathcal{V}_t is defined based on (OC) and (UC) conditions which naturally lead to the relationship in equation (19).

Given Lemma D.1, with probability $\geq 1 - \delta$, the cumulative regret R_T is bounded by

$$\begin{aligned} R_T &\triangleq \sum_{t=1}^T r(\mathbf{x}_t) \\ &= 2(\eta(1 - \lambda) + 1) \sum_{t=1}^T \beta_t^{1/2} \sigma_t(\mathbf{x}_t, \mathbf{z}_t) \quad (\text{from Lemma D.1}) \\ &\leq 2(\eta(1 - \lambda) + 1) \beta_T^{1/2} \sum_{t=1}^T \sigma_t(\mathbf{x}_t, \mathbf{z}_t) \quad (\text{from the monotonicity of } \beta_t). \end{aligned}$$

To prove Theorem 3.2, it remains to show that

$$2\beta_T^{1/2} \sum_{t=1}^T \sigma_t(\mathbf{x}_t, \mathbf{z}_t) \leq \sqrt{C_1 T \beta_T \gamma_T} \quad (21)$$

where $C_1 \triangleq 8/\log(1 + \sigma_n^{-2})$, which follows directly from the work of Srinivas et al. (2010).

Specifically, using Cauchy-Schwarz inequality,

$$4\beta_T \left(\sum_{t=1}^T \sigma_t(\mathbf{x}_t, \mathbf{z}_t) \right)^2 \leq 4\beta_T T \sum_{t=1}^T \sigma_t^2(\mathbf{x}_t, \mathbf{z}_t).$$

Following Lemma 5.4 in (Srinivas et al., 2010), $\sum_{t=1}^T 4\beta_T \sigma_t^2(\mathbf{x}_t, \mathbf{z}_t) \leq C_1 \beta_T \gamma_T$ where $C_1 \triangleq 8/\log(1 + \sigma_n^{-2})$. Hence,

$$\begin{aligned} 4\beta_T \left(\sum_{t=1}^T \sigma_t(\mathbf{x}_t, \mathbf{z}_t) \right)^2 &\leq C_1 T \beta_T \gamma_T \\ 2\beta_T^{1/2} \sum_{t=1}^T \sigma_t(\mathbf{x}_t, \mathbf{z}_t) &\leq \sqrt{C_1 T \beta_T \gamma_T}. \end{aligned}$$

As a result,

$$R_T \leq (\eta(1 - \lambda) + 1) \sqrt{C_1 T \beta_T \gamma_T}.$$

E PROOF OF THEOREM 3.7

Recall that $\mathbf{x}_t^- \triangleq \arg \max_{\mathbf{x} \in \mathcal{X}} \rho_{l_t}(\mathbf{x}; \alpha)$ and at iteration t , we recommend $\mathbf{x}_t^* = \mathbf{x}_{\zeta(t)}^-$ where

$$\zeta(t) \triangleq \arg \max_{t' \in \{1, \dots, t\}} \rho_{l_{t'}}(\mathbf{x}_{t'}^-; \alpha). \quad (22)$$

We need to prove that with probability $\geq 1 - \delta$,

$$r(\mathbf{x}_t^*) = r(\mathbf{x}_{\zeta(t)}^-) \leq \eta \sqrt{C_1 \beta_t \gamma_t / t}.$$

Proof. From the definition of $\zeta(t)$ in equation (22),

$$\rho_{l_{\zeta(t)}}(\mathbf{x}_{\zeta(t)}^-; \alpha) \geq \rho_{l_{t'}}(\mathbf{x}_{t'}^-; \alpha) \quad \forall 1 \leq t' \leq t.$$

Hence,

$$\rho_{l_{\zeta(t)}}(\mathbf{x}_{\zeta(t)}^-; \alpha) \geq \frac{1}{t} \sum_{t'=1}^t \rho_{l_{t'}}(\mathbf{x}_{t'}^-; \alpha). \quad (23)$$

$$\begin{aligned} r(\mathbf{x}_{\zeta(t)}^-) &\triangleq \rho_f(\mathbf{x}_*; \alpha) - \rho_f(\mathbf{x}_{\zeta(t)}^-; \alpha) \\ &\leq \rho_f(\mathbf{x}_*; \alpha) - \rho_{l_{\zeta(t)}}(\mathbf{x}_{\zeta(t)}^-; \alpha) \\ &\leq \rho_f(\mathbf{x}_*; \alpha) - \frac{1}{t} \sum_{t'=1}^t \rho_{l_{t'}}(\mathbf{x}_{t'}^-; \alpha) \end{aligned} \quad (24)$$

$$\begin{aligned} &= \frac{1}{t} \sum_{t'=1}^t (\rho_f(\mathbf{x}_*; \alpha) - \rho_{l_{t'}}(\mathbf{x}_{t'}^-; \alpha)) \\ &\leq \frac{1}{t} \sum_{t'=1}^t (\rho_{u_{t'}}(\mathbf{x}_*; \alpha) - \rho_{l_{t'}}(\mathbf{x}_{t'}^-; \alpha)) \\ &\leq \frac{1}{t} \sum_{t'=1}^t (\rho_{u_{t'}}(\mathbf{x}_{t'}^+; \alpha) - \rho_{l_{t'}}(\mathbf{x}_{t'}^-; \alpha)) \\ &\leq \frac{1}{t} \sum_{t'=1}^t \min \left(\eta (\rho_{u_{t'}}(\mathbf{x}_{t'}; \alpha) - \rho_{l_{t'}}(\mathbf{x}_{t'}; \alpha)), \frac{1}{\lambda} (\rho_{u_{t'}}(\mathbf{x}_{t'}; \alpha) - \rho_{l_{t'}}(\mathbf{x}_{t'}^-; \alpha)) \right) \end{aligned} \quad (25)$$

$$\leq \frac{1}{t} \sum_{t'=1}^t \min \left(\eta (\rho_{u_{t'}}(\mathbf{x}_{t'}; \alpha) - \rho_{l_{t'}}(\mathbf{x}_{t'}; \alpha)), \frac{1}{\lambda} (\rho_{u_{t'}}(\mathbf{x}_{t'}; \alpha) - \rho_{l_{t'}}(\mathbf{x}_{t'}; \alpha)) \right) \quad (26)$$

$$= \frac{1}{t} \sum_{t'=1}^t \eta (\rho_{u_{t'}}(\mathbf{x}_{t'}; \alpha) - \rho_{l_{t'}}(\mathbf{x}_{t'}; \alpha)) \quad (27)$$

$$\leq \frac{1}{t} \sum_{t'=1}^t \eta \left(2\beta_{t'}^{1/2} \sigma_{t'}(\mathbf{x}_{t'}, \mathbf{z}_{t'}) \right) \quad (28)$$

$$\begin{aligned} &\leq \frac{1}{t} \sum_{t'=1}^t \eta \left(2\beta_t^{1/2} \sigma_{t'}(\mathbf{x}_{t'}, \mathbf{z}_{t'}) \right) \quad (\text{from the monotonicity of } \beta_{t'}) \\ &\leq \frac{\eta}{t} \sqrt{C_1 t \beta_t \gamma_t} \end{aligned} \quad (29)$$

$$= \eta \sqrt{C_1 \beta_t \gamma_t / t}.$$

where equation (24) is from equation (23); equation (25) is from OC equation (6) and UC equation (7); equation (26) is because $\rho_{l_{t'}}(\mathbf{x}_{t'}^-; \alpha) \geq \rho_{l_{t'}}(\mathbf{x}_{t'}; \alpha)$; equation (27) is because $\eta \in [1, 1/\lambda]$; equation (28) is from equation (15); equation (29) is from equation (21). \square

F META-VBO FOR VANILLA BO

F.1 BO WITH V-SET

In this section, we consider the vanilla Bayesian optimization (BO) of a blackbox function. The setting is similar to Sec. 2.1 except that there are not any environmental random variables.

In particular, let the blackbox function be denoted as $f : \mathcal{X} \rightarrow \mathbb{R}$. The evaluation of f is expensive and we can only observe noisy evaluation $y(\mathbf{x}) \triangleq f(\mathbf{x}) + \epsilon(\mathbf{x})$ where $\epsilon(\mathbf{x}) \sim \mathcal{N}(0, \sigma_n^2)$. We would like to find the maximizer

$$\mathbf{x}_* \triangleq \arg \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}). \quad (30)$$

BO is about designing a sequential strategy of selecting input query \mathbf{x}_t at iteration t , considering observations in the previous $t - 1$ iterations, denoted as $\mathbf{y}_{\mathbf{D}_t} \triangleq (y(\mathbf{x}_{t'})_{t'=1}^{t-1})$ where $\mathbf{D}_t \triangleq (\mathbf{x}_{t'})_{t'=1}^{t-1}$ to quickly discovery \mathbf{x}_* .

To obtain the posterior belief of f given observations $\mathbf{y}_{\mathbf{D}_t}$, the blackbox function f is often modelled with a GP (Rasmussen and Williams, 2006). Let the posterior belief of $f(\mathbf{x})$ be denoted as $\mathcal{N}(\mu_t(\mathbf{x}), \sigma_t^2(\mathbf{x}))$.

Under the assumption that f belonging to the RKHS space with a norm bounded by B , we adopt the well-known confidence bound of $f(\mathbf{x})$ from the work of Chowdhury and Gopalan (2017). Let γ_t denote the maximum information gain about f that can be obtained from any set of $t - 1$ observations. Choosing $\delta \in (0, 1)$, $\beta_t = (B + \sigma_n \sqrt{2(\gamma_t + 1 + \log 1/\delta)})^2$, then the event that

$$\forall t \geq 1, \forall \mathbf{x} \in \mathcal{X}, l_t(\mathbf{x}) \leq f(\mathbf{x}) \leq u_t(\mathbf{x}) \quad (31)$$

holds with probability $\geq 1 - \delta$ where the lower and upper confidence bounds are defined as

$$l_t(\mathbf{x}) \triangleq \mu_t(\mathbf{x}) - \beta_t^{1/2} \sigma_t(\mathbf{x}) \quad \text{and} \quad u_t(\mathbf{x}) \triangleq \mu_t(\mathbf{x}) + \beta_t^{1/2} \sigma_t(\mathbf{x}), \text{ resp.} \quad (32)$$

It follows that

$$\forall t \geq 1, f(\mathbf{x}_*) \in [l_t(\mathbf{x}_t^-), u_t(\mathbf{x}_t^+)] \quad (33)$$

where $\mathbf{x}_t^- \triangleq \arg \max_{\mathbf{x} \in \mathcal{X}} l_t(\mathbf{x})$ and $\mathbf{x}_t^+ \triangleq \arg \max_{\mathbf{x} \in \mathcal{X}} u_t(\mathbf{x})$.

The *versatile query set* (V-set), denoted as \mathcal{V}_t , consists of those inputs $\mathbf{x} \in \mathcal{X}$ that satisfies the following overlapping condition (OC) and uncertainty condition (UC)

$$u_t(\mathbf{x}) \geq \max \left(\underbrace{l_t(\mathbf{x}_t^-) + \lambda(u_t(\mathbf{x}_t^+) - l_t(\mathbf{x}_t^-))}_{\text{overlapping condition}}, \underbrace{l_t(\mathbf{x}) + \frac{1}{\eta} (u_t(\mathbf{x}_t^+) - l_t(\mathbf{x}_t^-))}_{\text{uncertainty condition}} \right) \quad (34)$$

where $\lambda \in [0, 1]$ and $\eta \in [1, 1/\lambda]$. Similar to the proof of Theorem 3.2, we can obtain the following theorem on the cumulative regret of selecting the input query in V-set \mathcal{V}_t .

Theorem F.1. *Given the selection of \mathbf{x}_t as any input in \mathcal{V}_t , the algorithm achieves a sublinear cumulative regret:*

$$\forall T \geq 1, R_T \leq (\eta(1 - \lambda) + 1) \sqrt{C_1 T \beta_T \gamma_T} \quad (35)$$

holds with probability $\geq 1 - \delta$ where $C_1 \triangleq 8 / \log(1 + \sigma_n^{-2})$; $\gamma_T, \beta_T, \delta$ are defined in equation (31); λ and η are defined in equation (34).

F.2 META-VBO

Adopting the meta-BO setting in Sec. 3, we consider additional prior tasks $\tau \in \mathcal{T}$. The blackbox function of any prior task τ , denoted as $f_\tau : \mathcal{X} \rightarrow \mathbb{R}$, is also assumed to belong to an RKHS with a norm bounded by B . The optimization result of a prior task τ includes the GP posterior mean and variance of f_τ , denoted as $\hat{\mu}_\tau : \mathcal{X} \rightarrow \mathbb{R}$ and $\hat{\sigma}_\tau^2 : \mathcal{X} \rightarrow \mathbb{R}^+$, respectively. We also denote the lower and upper bound of the function evaluation $f_\tau(\mathbf{x})$ as $[\hat{l}_\tau(\mathbf{x}), \hat{u}_\tau(\mathbf{x})]$ where

$$\hat{l}_\tau(\mathbf{x}) \triangleq \hat{\mu}_\tau(\mathbf{x}) - \beta_{t_\tau}^{1/2} \hat{\sigma}_\tau(\mathbf{x}) \quad (36)$$

$$\hat{u}_\tau(\mathbf{x}) \triangleq \hat{\mu}_\tau(\mathbf{x}) + \beta_{t_\tau}^{1/2} \hat{\sigma}_\tau(\mathbf{x}) \quad (37)$$

and t_τ is the iteration where the BO procedure on prior task τ stops.

From Sec. 3.2, we utilize the prior tasks to assign different *priorities* to inputs in \mathcal{V}_t . We can approach this in two distinct ways.

1. We assume that the global maximizer of the current task is close to a *local maximizer* of a prior task. In particular, we assign higher priorities to inputs that have a higher chance of being a local maximizer (restricted to \mathcal{V}_t) of prior tasks. Under this assumption, we introduce the notion of the *probable local maximizer set*, denoted as $\mathcal{P}_{\tau|\mathcal{V}_t}$, in a similar manner to that in Sec. 3.2:

$$\mathcal{P}_{\tau|\mathcal{V}_t} \triangleq \{\mathbf{x} \in \mathcal{V}_t \mid \hat{u}_\tau(\mathbf{x}) \geq \max_{\mathbf{x}' \in \mathcal{V}_t} \hat{l}_\tau(\mathbf{x}')\}. \quad (38)$$

2. We assume that the global maximizer of the current task is close to a *global maximizer* of a prior task. In particular, we assign higher priorities to inputs that have a higher chance of being a global maximizer (restricted to \mathcal{V}_t) of prior tasks. Under this assumption, we introduce the notion of *the probable global maximizer set*, denoted as $\mathcal{P}_{\tau|\mathcal{X}}$, in a similar manner to that in Sec. 3.2:

$$\mathcal{P}_{\tau|\mathcal{X}} \triangleq \{\mathbf{x} \in \mathcal{X} \mid \hat{u}_{\tau}(\mathbf{x}) \geq \max_{\mathbf{x}' \in \mathcal{X}} \hat{l}_{\tau}(\mathbf{x}')\}. \quad (39)$$

Let us denote the probable maximizer set as \mathcal{P}_{τ} which can be set to be either the probable local maximizer set $\mathcal{P}_{\tau|\mathcal{V}_t}$ or the global maximizer set $\mathcal{P}_{\tau|\mathcal{X}}$ depending on the assumption one may adopt. Then, the priority of an input $\mathbf{x} \in \mathcal{V}_t$, denoted as $\pi_t(\mathbf{x})$, is measured by the number of the probable maximizer sets that \mathbf{x} belongs to, i.e.,

$$\forall \mathbf{x} \in \mathcal{V}_t, \pi_t(\mathbf{x}) \triangleq \sum_{\tau \in \mathcal{T}} \mathbb{1}_{\mathbf{x} \in \mathcal{P}_{\tau}}. \quad (40)$$

The higher $\pi_t(\mathbf{x})$ is, the more likely that \mathbf{x} is in the probable local maximizer set of a prior task. Hence, we select the input query \mathbf{x}_t as

$$\mathbf{x}_t \triangleq \max_{\mathbf{x} \in \mathcal{V}_{t,\max}} u_t(\mathbf{x}) \text{ where } \mathcal{V}_{t,\max} \triangleq \{\mathbf{x} \in \mathcal{V}_t \mid \pi_t(\mathbf{x}) = \max_{\mathbf{x}' \in \mathcal{V}_t} \pi_t(\mathbf{x}')\}. \quad (41)$$

Following Sec. 3.3, we propose to recommend $\mathbf{x}_t^* \triangleq \mathbf{x}_{\zeta(t)}^-$ as an approximation to the optimal solution where $\zeta(t) \triangleq \arg \max_{t' \in \{1, \dots, t\}} l_{t'}(\mathbf{x}_{t'}^-)$.

Theorem F.2. *Let $\zeta(t) \triangleq \arg \max_{t' \in \{1, \dots, t\}} l_{t'}(\mathbf{x}_{t'}^-)$, then, by recommending $\mathbf{x}_t^* = \mathbf{x}_{\zeta(t)}^-$, the instantaneous regret at \mathbf{x}_t^* is bounded: $r(\mathbf{x}_t^*) \leq \eta \sqrt{C_1 \beta_t \gamma_t / t}$ with probability $\geq 1 - \delta$.*

G META-VBO FOR ADVERSARIALLY ROBUST BAYESIAN OPTIMIZATION

Let us consider the adversarially robust BO problem from the work of Bogunovic et al. (2018) that is reformulated in (Nguyen et al., 2021b). It is to find the maximizer $\mathbf{x}_* \triangleq \arg \max_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{z} \in \mathcal{Z}} f(\mathbf{x}, \mathbf{z})$. Furthermore, the work of (Nguyen et al., 2021b), casts this problem as a BO of VaR problem for $\alpha \rightarrow 0^+$ and a uniform distribution of \mathbf{Z} . Let us denote $v_f(\mathbf{x}; 0^+) \triangleq \min_{\mathbf{z} \in \mathcal{Z}} f(\mathbf{x}, \mathbf{z})$. Then, by replacing $\rho_f(\mathbf{x}; \alpha)$ with $v_f(\mathbf{x}; 0^+)$, equivalently, $\min_{\mathbf{z} \in \mathcal{Z}} f(\mathbf{x}, \mathbf{z})$, we obtain a meta-BO solution to the adversarially robust BO.

Specifically, from the following confidence bounds of $\min_{\mathbf{z} \in \mathcal{Z}} f(\mathbf{x}, \mathbf{z})$

$$\min_{\mathbf{z} \in \mathcal{Z}} l_t(\mathbf{x}, \mathbf{z}) \leq \min_{\mathbf{z} \in \mathcal{Z}} f(\mathbf{x}, \mathbf{z}) \leq \min_{\mathbf{z} \in \mathcal{Z}} u_t(\mathbf{x}, \mathbf{z}),$$

we obtain the confidence bounds of $\min_{\mathbf{z} \in \mathcal{Z}} f(\mathbf{x}_*, \mathbf{z})$:

$$\min_{\mathbf{z} \in \mathcal{Z}} l_t(\mathbf{x}_t^-, \mathbf{z}) \leq \min_{\mathbf{z} \in \mathcal{Z}} f(\mathbf{x}_*, \mathbf{z}) \leq \min_{\mathbf{z} \in \mathcal{Z}} u_t(\mathbf{x}_t^+, \mathbf{z})$$

where $\mathbf{x}_t^- \triangleq \arg \max_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{z} \in \mathcal{Z}} l_t(\mathbf{x}, \mathbf{z})$ and $\mathbf{x}_t^+ \triangleq \arg \max_{\mathbf{x} \in \mathcal{X}} u_t(\mathbf{x}, \mathbf{z})$.

The V-set \mathcal{V}_t is defined as the set of inputs $\mathbf{x} \in \mathcal{X}$ that satisfy:

- Overlapping condition (OC):

$$\min_{\mathbf{z} \in \mathcal{Z}} u_t(\mathbf{x}, \mathbf{z}) \geq \min_{\mathbf{z} \in \mathcal{Z}} l_t(\mathbf{x}_t^-, \mathbf{z}) + \lambda (\min_{\mathbf{z} \in \mathcal{Z}} u_t(\mathbf{x}_t^+, \mathbf{z}) - \min_{\mathbf{z} \in \mathcal{Z}} l_t(\mathbf{x}_t^-, \mathbf{z}))$$

where $\lambda \in [0, 1]$.

- Uncertainty condition (UC):

$$\min_{\mathbf{z} \in \mathcal{Z}} u_t(\mathbf{x}, \mathbf{z}) - \min_{\mathbf{z} \in \mathcal{Z}} l_t(\mathbf{x}, \mathbf{z}) \geq \frac{1}{\eta} (\min_{\mathbf{z} \in \mathcal{Z}} u_t(\mathbf{x}_t^+, \mathbf{z}) - \min_{\mathbf{z} \in \mathcal{Z}} l_t(\mathbf{x}_t^-, \mathbf{z}))$$

where $\eta \in [1, 1/\lambda]$.

Algorithm 1 Meta-BO of Risk Measure ρ and risk level α **Require:** $\mathcal{X}, \alpha, \lambda, \eta, (\mathbf{D}_1, \mathbf{y}_{\mathbf{D}_1}), (\rho_{\hat{l}_\tau}, \rho_{\hat{u}_\tau})_{\tau \in \mathcal{T}}$

- 1: Compute GP posterior given $(\mathbf{D}_1, \mathbf{y}_{\mathbf{D}_1})$: μ_1, σ_1 .
- 2: **for** t in $1, 2, \dots, T$ **do**
- 3: Compute $\phi_t^- \triangleq \rho_{l_t}(\mathbf{x}_t^-; \alpha)$ and $\phi_t^+ \triangleq \rho_{u_t}(\mathbf{x}_t^+; \alpha)$.
- 4: Compute the uncertainty of $\rho_f(\mathbf{x}_*; \alpha)$: $\text{CI}_* \triangleq \phi_t^+ - \phi_t^-$.
- 5: Find $\mathcal{V}_t \triangleq \{\mathbf{x} \in \mathcal{X} \mid \rho_{u_t}(\mathbf{x}; \alpha) \geq \phi_t^- + \lambda \text{CI}_* \wedge \rho_{u_t}(\mathbf{x}; \alpha) - \rho_{l_t}(\mathbf{x}; \alpha) \geq \eta^{-1} \text{CI}_*\}$.
- 6: For $\tau \in \mathcal{T}$, compute $\phi_{\tau,t}^- \triangleq \max_{\mathbf{x} \in \mathcal{V}_t} \rho_{\hat{l}_\tau}(\mathbf{x}; \alpha)$.
- 7: Select $\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{V}_t} \left(\sum_{\tau \in \mathcal{T}} \mathbb{1}_{\rho_{\hat{u}_\tau}(\mathbf{x}; \alpha) \geq \phi_{\tau,t}^-}, \rho_{u_t}(\mathbf{x}; \alpha) \right)$.
- 8: Select \mathbf{z}_t as a lacing value w.r.t. \mathbf{x}_t, l_t, u_t , and α (Nguyen et al., 2021b;a).
- 9: $\mathbf{D}_{t+1} = \mathbf{D}_t \cup \{(\mathbf{x}_t, \mathbf{z}_t)\}$ and $\mathbf{y}_{\mathbf{D}_{t+1}} = \mathbf{y}_{\mathbf{D}_t} \cup \{y_t(\mathbf{x}_t, \mathbf{z}_t)\}$.
- 10: Compute GP posterior given $(\mathbf{D}_{t+1}, \mathbf{y}_{\mathbf{D}_{t+1}})$: μ_{t+1} and σ_{t+1} .
- 11: **end for**

Given \mathbf{x}_t chosen as any input from \mathcal{V}_t , we select $\mathbf{z}_t = \arg \min_{\mathbf{z} \in \mathcal{Z}} l_t(\mathbf{x}_t, \mathbf{z})$ (Bogunovic et al., 2018). This choice of \mathbf{z}_t is an LV w.r.t. \mathbf{x}_t, l_t, u_t , and $\alpha \rightarrow 0^+$ since it satisfies the definition of LV (Nguyen et al., 2021b):

$$l_t(\mathbf{x}_t, \mathbf{z}_t) = \min_{\mathbf{z} \in \mathcal{Z}} l_t(\mathbf{x}_t, \mathbf{z}) \leq \min_{\mathbf{z} \in \mathcal{Z}} u_t(\mathbf{x}_t, \mathbf{z}) \leq u_t(\mathbf{x}_t, \mathbf{z}_t)$$

The priorities assigned to inputs \mathcal{V}_t follow our discussion in Sec. 3.2. In particular, the probable local maximizer set is constructed as follows.

$$\mathcal{P}_\tau | \mathcal{V}_t = \{\mathbf{x} \in \mathcal{V}_t \mid \min_{\mathbf{z} \in \mathcal{Z}} \hat{u}_\tau(\mathbf{x}, \mathbf{z}) \geq \max_{\mathbf{x}' \in \mathcal{V}_t} \min_{\mathbf{z} \in \mathcal{Z}} \hat{l}_\tau(\mathbf{x}', \mathbf{z})\}.$$

Then, the input query \mathbf{x}_t is selected as

$$\mathbf{x}_t = \max_{\mathbf{x} \in \mathcal{V}_{t,\max}} \min_{\mathbf{z} \in \mathcal{Z}} u_t(\mathbf{x}, \mathbf{z})$$

where $\mathcal{V}_{t,\max}$ is defined in Sec. 3.2.

Regarding the recommended input, we follow the recommendation strategy in Sec. 3.3 to choose $\mathbf{x}_t^* = \bar{\mathbf{x}}_{\zeta(t)}$ where $\zeta(t) = \arg \max_{t' \in \{1, \dots, t\}} \min_{\mathbf{z} \in \mathcal{Z}} l_{t'}(\mathbf{x}_{t'}, \mathbf{z})$.

H PSEUDOCODE

The pseudocode of our proposed algorithm is shown in Algorithm 1. It is noted that $\rho_{\hat{l}_\tau}$ and $\rho_{\hat{u}_\tau}$ represent functions for calculating the lower and upper bounds of the risk measure, respectively. In line 7, the comparison between tuples $\left(\sum_{\tau \in \mathcal{T}} \mathbb{1}_{\rho_{\hat{u}_\tau}(\mathbf{x}; \alpha) \geq \phi_{\tau,t}^-}, \rho_{u_t}(\mathbf{x}; \alpha) \right)$ is executed elementwise. This comparison involves assessing the second position if the first position results in a tie, akin to the comparison operations in Python.

We demonstrate that the additional computational complexity of our proposed algorithm is not significant when compared to V-UCB, which does not utilize any prior tasks. Assuming finite domains \mathcal{X} and \mathcal{Z} to leverage the computational complexity presented in Nguyen et al. (2021b), each iteration of V-UCB has a complexity of $\mathcal{O}(|\mathcal{Z}| |\mathcal{X}| (|\mathbf{D}_t|^2 + \log |\mathcal{Z}|))$ (Nguyen et al., 2021b). Both lines 3 and 5 in our proposed Algorithm 1 share the same computational complexity of $\mathcal{O}(|\mathcal{Z}| |\mathcal{X}| (|\mathbf{D}_t|^2 + \log |\mathcal{Z}|))$ as it requires the computation of VaR as in V-UCB. Line 6, involving iteration over prior tasks, incurs an additional $\mathcal{O}(|\mathcal{T}| |\mathcal{Z}| |\mathcal{X}| (|\mathbf{D}_t|^2 + \log |\mathcal{Z}|))$. Consequently, compared to V-UCB, our algorithm introduces only an extra linear dependence on the number of prior tasks.

I ADDITIONAL EXPERIMENTS

I.1 ILLUSTRATION OF VERSATILE QUERY SET AND PROBABLE LOCAL MAXIMIZER SET

We visualize the versatile query set (V-set) \mathcal{Q}_t and the probable local maximizer sets $\mathcal{P}_\tau | \mathcal{Q}_t$ when we maximize the VaR of a blackbox function given a set of prior tasks in 3 different scenarios.

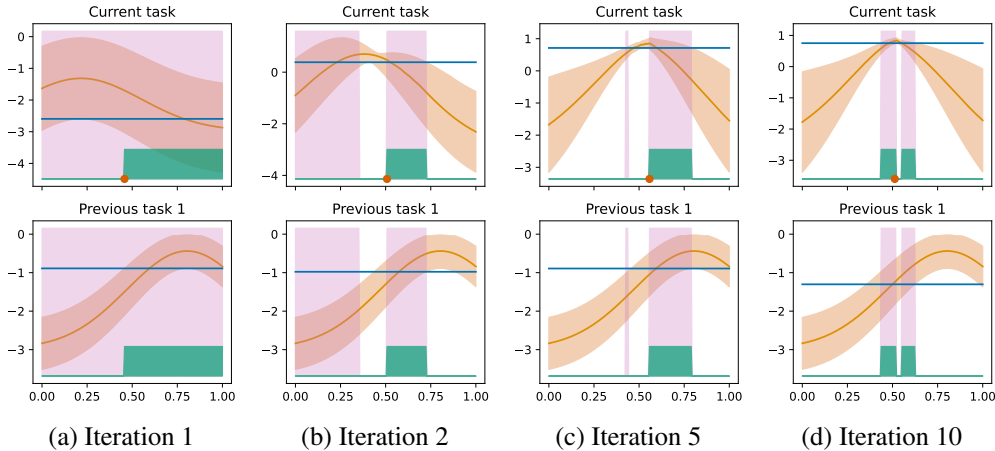


Figure 5: Visualization of the query set \mathcal{Q}_t and the probable local maximizer sets $\mathcal{P}_{\tau|\mathcal{Q}_t}$ when the prior task has poorly-estimated VaR. The figures in each column share the same x -axis. The top figure displays the GP posterior mean (the orange line) along with the lower and upper bounds (the orange shaded area) for the blackbox function of the current task. The bottom figure displays the same information for the prior tasks. The blue line in both figures indicates the maximum lower bounds, while the pink shaded area represents the V-set \mathcal{Q}_t . Although \mathcal{Q}_t is defined based on the bounds of the blackbox function of the current task, it is also visualized in the bottom figures to aid in visualizing the probable local maximizer set $\mathcal{P}_{\tau|\mathcal{Q}_t}$. The green shaded areas in the bottom figures illustrate the probable local maximizer set $\mathcal{P}_{\tau|\mathcal{Q}_t}$ for the prior task. In the top figures, the green shaded areas indicate the count of probable local maximizer sets to which an input belongs (a higher green area corresponds to a larger number of probable local maximizer sets). The red dot shows the input query \mathbf{x}_t .

1. **Figure 5: prior task with poorly-estimated VaR.** The set of prior tasks consists of a single task that differs from the current task. Furthermore, the VaR of the prior task is not well-estimated due to sparse observations in the prior task, as shown by the considerable gap between the lower and upper bounds in the bottom figures. During the initial BO iterations, when the VaR of the current task remains uncertain across the input domain, the preceding task steers the input query towards a region where it is likely that the VaR of prior tasks attains high evaluations. However, as the posterior distribution of the VaR in the current task refines, such as at iteration 10, the influence of the prior task on the input query diminishes. Hence, we observe that as \mathcal{Q}_t reduces, the estimation of VaR in the prior task becomes too inaccurate to pinpoint promising inputs within \mathcal{Q}_t . Additionally, this figure highlights instances where \mathcal{Q}_t may become disjointed. For example, at iteration 2 (Figure 5b), the gap between two pink shaded regions is due to the uncertainty condition in Equation equation (7).
2. **Figure 6: Prior task with well-estimated VaR.** In this scenario, the blackbox function in the prior task is similar to that in Figure 5. However, the estimation of VaR for the prior task is more accurate, indicated by a smaller gap between lower and upper bounds in the bottom figures, due to an increased number of observations in the prior task. With a well-estimated VaR in the prior task, the probable local maximizer sets $\mathcal{P}_{\tau|\mathcal{Q}_t}$ (the green shaded areas) are small even at the first BO iteration. However, as the V-set \mathcal{Q}_t reduces, the influence of the prior task on the current optimization task reduces, such as at iteration 15.
3. **Figure 7: Multiple prior tasks.** There are 3 prior tasks: both prior task 1 & 2 differ from the current task but prior task 2 shares similarities in a local region around the maximizer; prior task 3 bears similarity to the current task. The varying heights of the green shaded areas in the top figures demonstrate the different priorities assigned to inputs within \mathcal{Q}_t . In particular, this illustration shows that a substantial proportion of useful prior tasks leads to a priority assignment that identifies promising inputs within \mathcal{Q}_t .




Figure 6: Visualization of the query set \mathcal{V}_t and the probable local maximizer sets $\mathcal{P}_{\tau|\mathcal{V}_t}$ when the prior task has well-estimated VaR. The plot notations are similar to those in Figure 5. In this context, there exists a single prior task where the VaR maximizer is positioned around $x = 0.75$, contrasting with the VaR maximizer of the current task, which is situated approximately at $x = 0.5$.

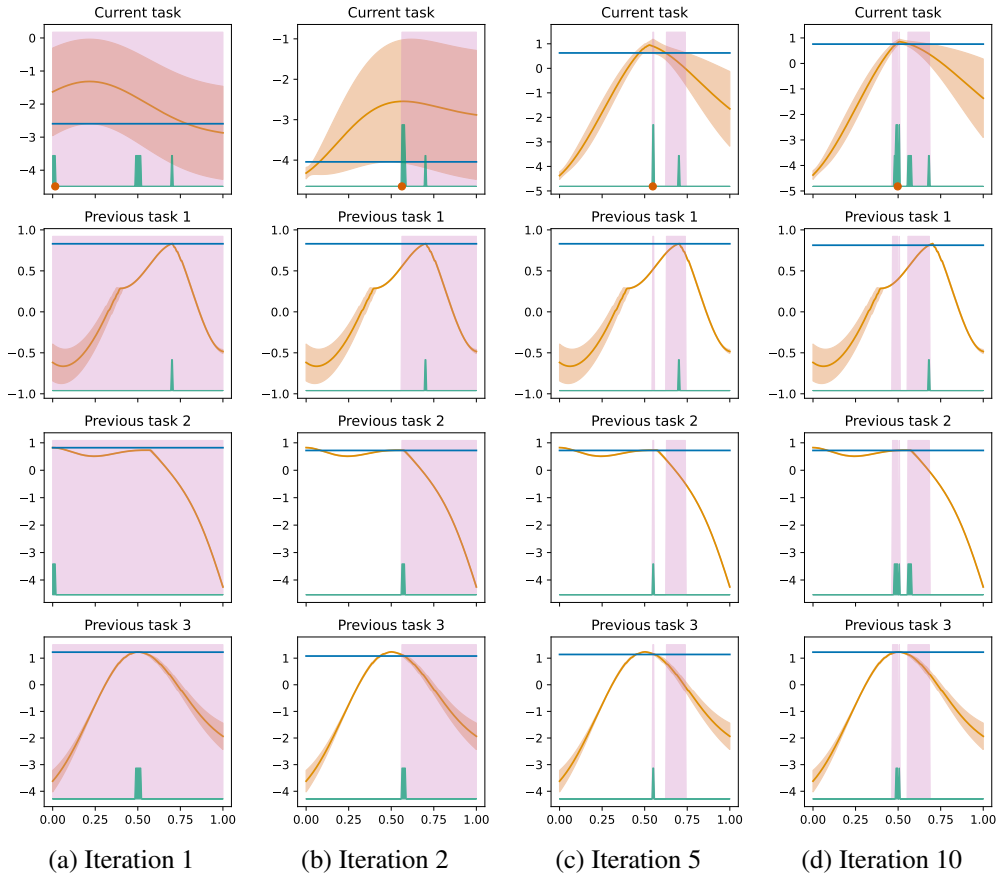


Figure 7: Visualization of the query set \mathcal{V}_t and the probable local maximizer sets $\mathcal{P}_{\tau|\mathcal{V}_t}$ when there are multiple prior tasks. The plot notations are similar to those in Figure 5.

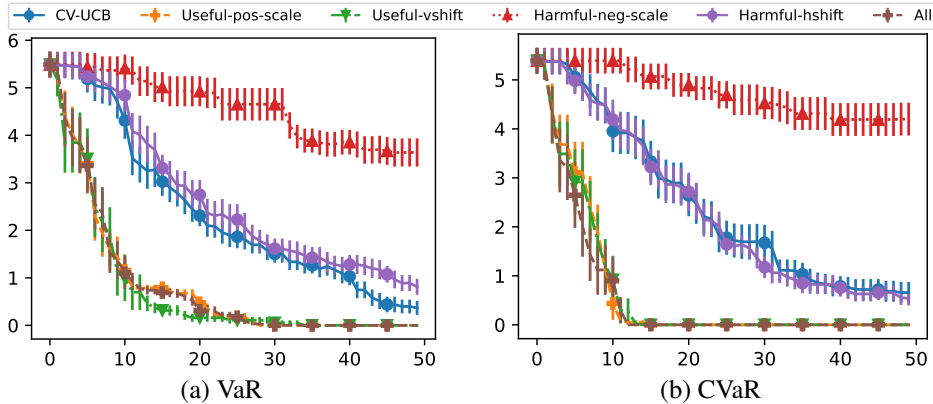


Figure 8: Plots of the regret at the recommended input against the BO iteration in optimizing (a) VaR and (b) CVaR of a blackbox function with 9-dimensional inputs.

I.2 EXPERIMENTS WITH 9-DIMENSIONAL INPUTS

In this experiment, we empirically demonstrate the effectiveness of our algorithm in optimizing the VaR and CVaR for a blackbox function with 9-dimensional inputs. We create this blackbox function by summing three instances of the Hartmann-3D function. It is important to note that, from the perspective of the BO algorithm, this function is a blackbox, and the underlying construction is unknown. Our experiments involve a diverse set \mathcal{T} of prior tasks, incorporating transformations such as positive and negative scaling, horizontal and vertical shifting, as detailed in Section 4. The outcomes are illustrated in Figure 8. When \mathcal{T} comprises only useful tasks from $\mathcal{T}_{\text{useful-pos-scale}}$ and $\mathcal{T}_{\text{useful-hshift}}$, our proposed meta-BO algorithm outperforms V-UCB (in Figure 8a) and CV-UCB (in Figure 8b). In the extreme scenario when \mathcal{T} contains harmful tasks from $\mathcal{T}_{\text{harmful-neg-scale}}$ and $\mathcal{T}_{\text{harmful-hshift}}$, our algorithm may not surpass V-UCB or CV-UCB, but it reduces the regret.

I.3 CONDITIONAL VALUE-AT-RISK EXPERIMENTS

The Gaussian curve $f_{\text{Gaussian}}(x, z)$ is defined as follows.

$$f_{\text{Gaussian}}(x, z) = 5 \exp \left(- \begin{pmatrix} x \\ z \end{pmatrix}^{\top} \begin{pmatrix} 0.05 & 0 \\ 0 & 0.5 \end{pmatrix} \begin{pmatrix} x \\ z \end{pmatrix} \right) \quad (42)$$

where $\mathcal{X} \times \mathcal{Z}$ consists of 10000 data points in $[0, 1]^2$. The distribution of Z is $\mathcal{N}(0.5, 0.09)$.

Other synthetic functions are obtained from <https://www.sfu.ca/~ssurjano>. The domains $\mathcal{X} \times \mathcal{Z}$ of the Branin-Hoo, the Goldstein-Price, the six-hump camel consist of 10000 data points in $[0, 1]^2$. The distributions of Z in these experiments are $\mathcal{N}(0.1, 0.2)$, $\mathcal{N}(0.1, 0.4)$, and $\mathcal{N}(0.1, 0.2)$, respectively. The domain $\mathcal{X} \times \mathcal{Z}$ of the Hartmann-3D function consists of 20000 data points in $[0, 1]^3$, and that of the Hartmann-6D function consists of 40000 data points in $[0, 1]^6$. The distributions of Z in the Hartmann-3D and Hartmann-6D experiments are $\mathcal{N}(0.9, 0.2)$ and $\mathcal{N}(0.5, 0.4)$, respectively. The portfolio optimization and the yacht hydrodynamics experiments are adopted from (Cakmak et al., 2020; Nguyen et al., 2021a). The domains $\mathcal{X} \times \mathcal{Z}$ of the two experiments consist of 30000 and 40000 data points, respectively. While the distribution of Z is uniform in the portfolio optimization experiment, it is $\mathcal{N}(0.5, 0.1)$ in the yacht hydrodynamics experiment. Figs. 9-14 show the current task and several prior tasks in the experiments with the Branin-Hoo, the Goldstein-Price, and the six-hump camel functions. While prior tasks in $\mathcal{T}_{\text{useful-vshift}}$ (i.e., vertical shifting) and $\mathcal{T}_{\text{useful-pos-scale}}$ (i.e., positive scaling) share common maximizers of VaR and CVaR of the current task, those in $\mathcal{T}_{\text{harmful-hshift}}$ (i.e., horizontal-x-shifting) and $\mathcal{T}_{\text{harmful-neg-scale}}$ (i.e., negative scaling) differ that of the current task.

Fig. 15 shows the average the standard error of $r(\mathbf{x}_t^*)$ over 30 random repetitions of the BO of CVaR experiments. These experiments share the same sets of prior tasks with the BO of VaR experiments in Sec. 4: $\mathcal{T}_{\text{useful-pos-scale}}$, $\mathcal{T}_{\text{useful-vshift}}$, $\mathcal{T}_{\text{harmful-neg-scale}}$, $\mathcal{T}_{\text{harmful-hshift}}$, and \mathcal{T}_{all} . The

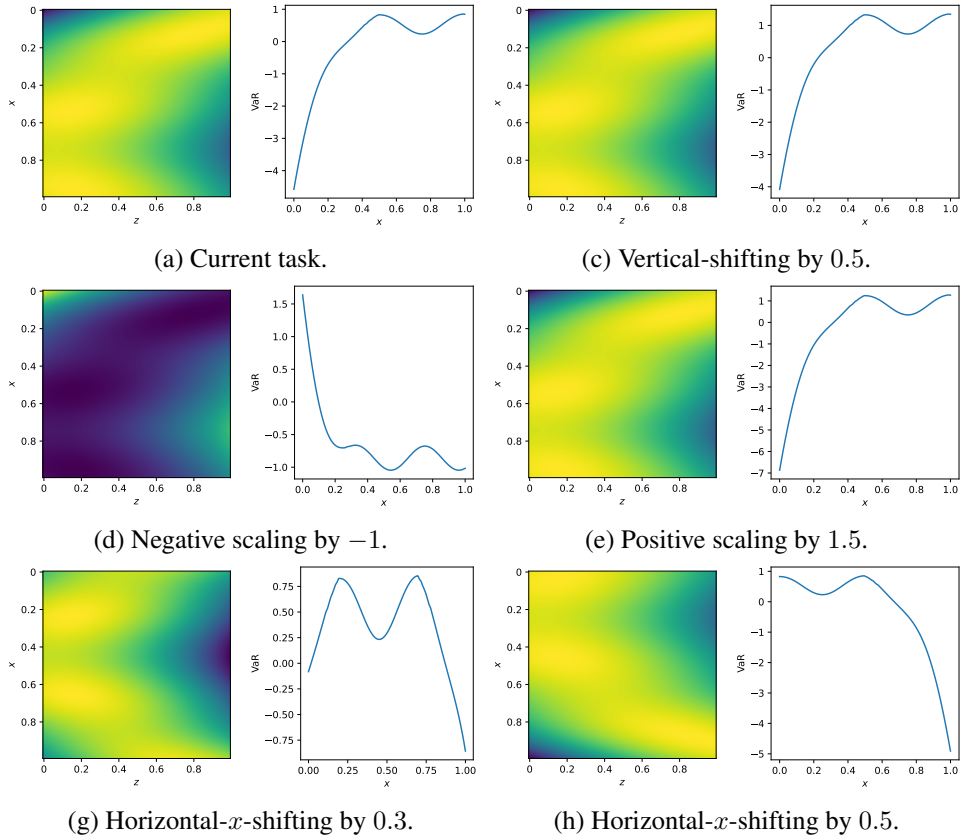


Figure 9: Current and prior BO of VaR tasks in Branin-Hoo experiments. The heatmaps show the blackbox function $f(\mathbf{x}, \mathbf{z})$.

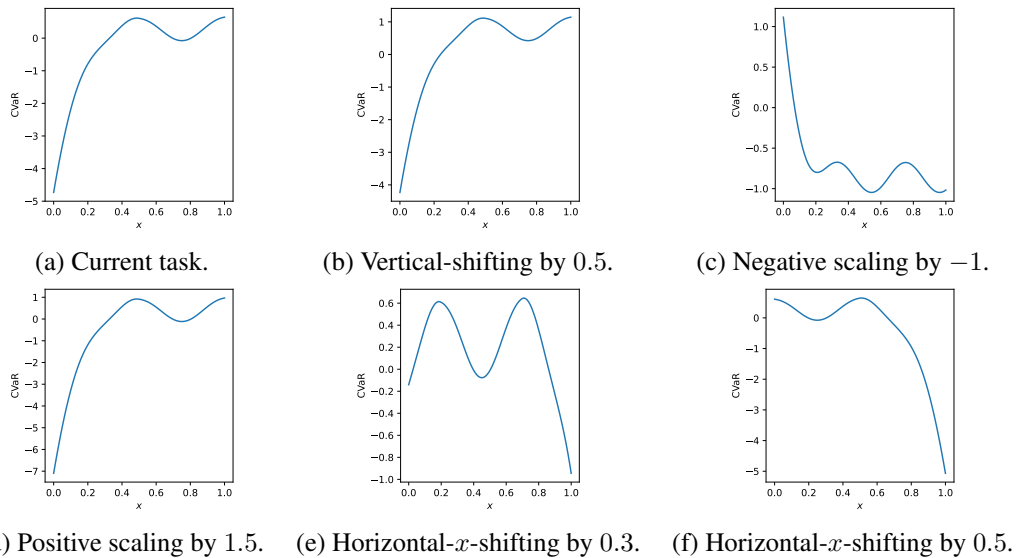


Figure 10: Current and prior BO of CVaR tasks in Branin-Hoo experiments.

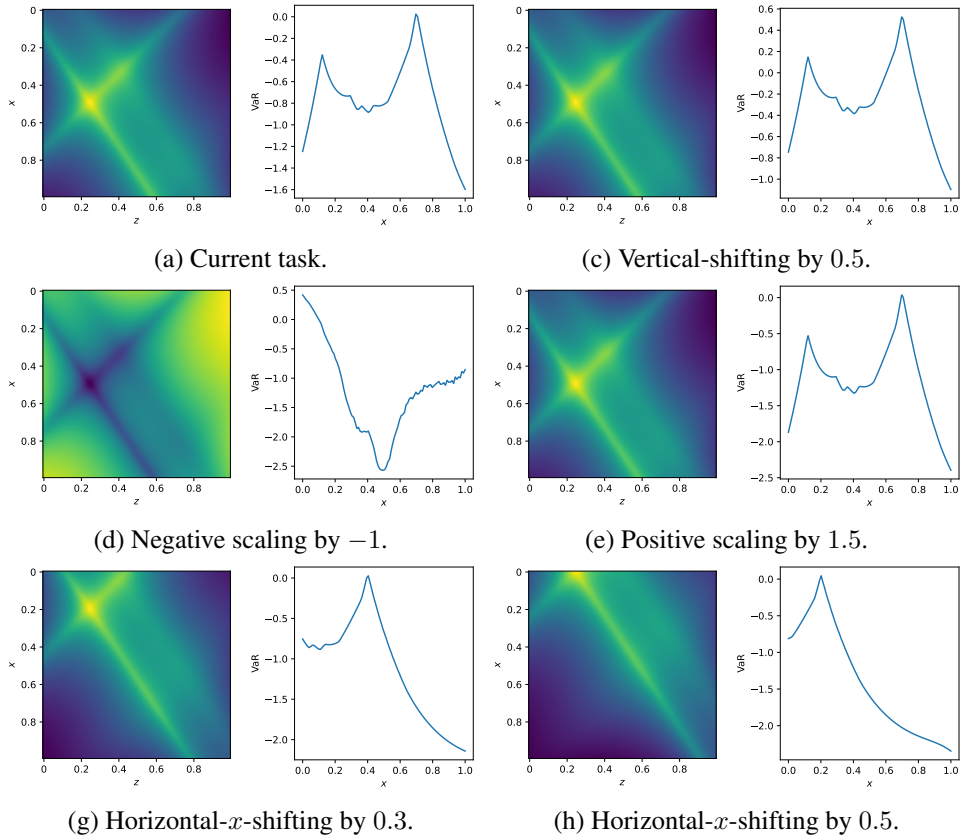


Figure 11: Current and prior BO of VaR tasks in Goldstein-Price experiments. The heatmaps show the blackbox function $f(\mathbf{x}, \mathbf{z})$.




Figure 12: Current and prior BO of CVaR tasks in Goldstein-Price experiments.

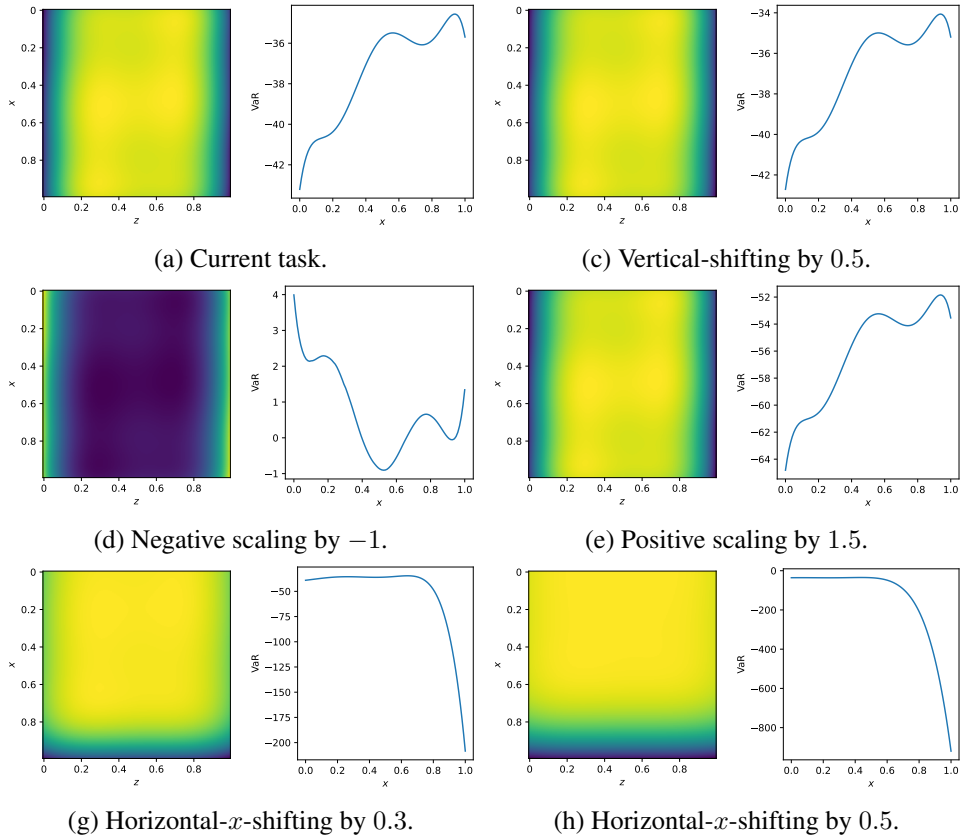


Figure 13: Current and prior BO of VaR tasks in six-hump camel experiments. The heatmaps show the blackbox function $f(\mathbf{x}, \mathbf{z})$.

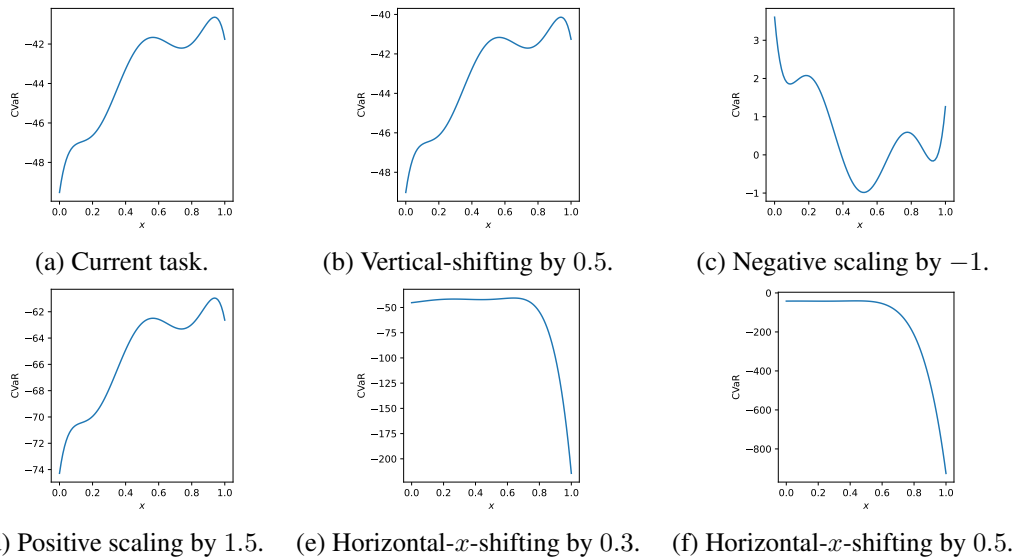


Figure 14: Current and prior BO of CVaR tasks in six-hump camel experiments.

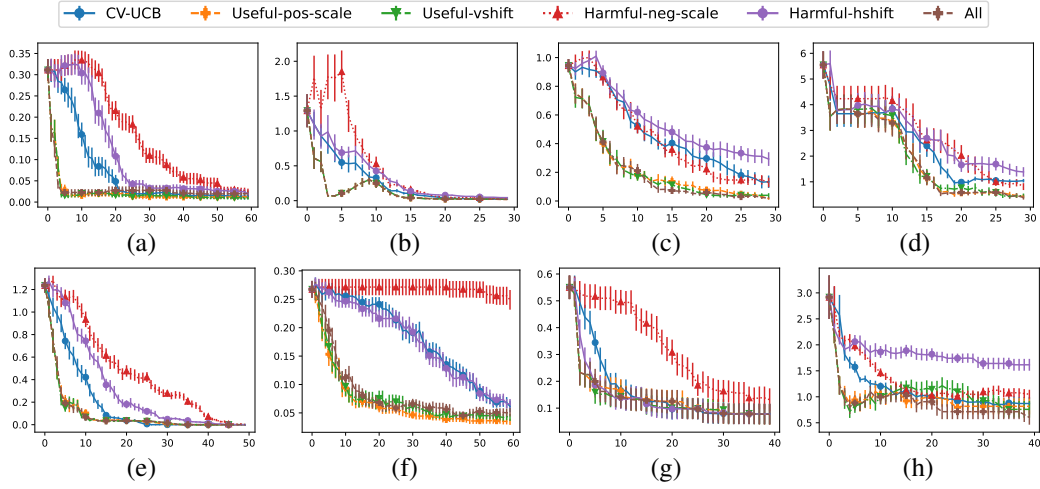


Figure 15: Meta-BO of CVaR experiments on (a) a Gaussian curve, (b) the Branin-Hoo function, (c) the Goldstein-Price function, (d) the six-hump camel function, (e) the Hartmann-3D function, (f) the Hartmann-6D function, (g) the yacht hydrodynamics dataset, and (h) the portfolio optimization dataset.

risk level is set at $\alpha = 0.1$ for all experiments. We choose $\lambda = 0$ and $\eta = 1$ as explained in Sec. 3.3. In Fig. 15, we observe that when \mathcal{T} contains only prior useful tasks in $\mathcal{T}_{\text{useful-pos-scale}}$ and $\mathcal{T}_{\text{useful-vshift}}$, our proposed meta-BO of CVaR outperforms CV-UCB (Nguyen et al., 2021a) which does not utilize any prior tasks. Furthermore, when \mathcal{T} contains both harmful and useful tasks in \mathcal{T}_{all} , our meta-BO of CVaR is still able to exploit prior useful tasks and is robust against prior harmful tasks to outperforms CV-UCB. When \mathcal{T} consists of only harmful tasks in $\mathcal{T}_{\text{harmful-neg-scale}}$ and $\mathcal{T}_{\text{hamrful-hsfhit}}$, our proposed algorithm does not outperform CV-UCB, but it still converges to a low-regret solution except for the experiment with the Hartmann-6D function. It is noted that the curves of useful-pos-scale, useful-vshift, and all overlap each other in most of the plots in Fig. 15.