

A GENERAL FRAMEWORK FOR USER-GUIDED BAYESIAN OPTIMIZATION

Anonymous authors

Paper under double-blind review

ABSTRACT

The optimization of expensive-to-evaluate black-box functions is prevalent in various scientific disciplines. Bayesian optimization is an automatic, general and sample-efficient method to solve these problems with minimal knowledge of the underlying function dynamics. However, the ability of Bayesian optimization to incorporate prior knowledge or beliefs about the function at hand in order to accelerate the optimization is limited, which reduces its appeal for knowledgeable practitioners with tight budgets. To allow domain experts to customize the optimization routine, we propose `CoLaBO`, the first Bayesian-principled framework for incorporating prior beliefs beyond the typical kernel structure, such as the likely location of the optimizer or the optimal value. The generality of `CoLaBO` makes it applicable across different Monte Carlo acquisition functions and types of user beliefs. We empirically demonstrate `CoLaBO`'s ability to substantially accelerate optimization when the prior information is accurate, and to retain approximately default performance when it is misleading.

1 INTRODUCTION

Bayesian Optimization (BO) is a well-established methodology for the optimization of expensive-to-evaluate black-box functions. Known for its sample efficiency, BO has been successfully applied to a variety of domains where laborious system tuning is prominent, such as hyperparameter optimization (Snoek et al., 2012; Bergstra et al., 2011b; Lindauer et al., 2022), neural architecture search (Ru et al., 2021; White et al., 2021), robotics (Calandra et al., 2014; Mayr et al., 2022), hardware design (Nardi et al., 2019; Ejje et al., 2022), and chemistry (Griffiths & Hernández-Lobato, 2020).

Typically employing a Gaussian Process (Rasmussen & Williams, 2006) (GP) surrogate model, BO allows the user to specify a prior over functions $p(f)$ via the Gaussian Process kernel, as well as an optional prior over its hyperparameters. Within the framework of the prior, the user can specify expected smoothness, output range and possible noise level of the function at hand, with the hopes of accelerating the optimization if accurate. However, the prior beliefs that can be specified within the framework of the kernel hyperparameters do not span the full range of beliefs that practitioners may possess. For example, practitioners may know which *parts of the input space* tend to work best (Oh et al., 2018; Perrone et al., 2019; Smith, 2018; Wang et al., 2019), know a range or upper bound on the optimal output (Jeong & Kim, 2021; Nguyen & Osborne, 2020) such as a maximal achievable accuracy of 100% , or other properties of the objective, such as preference relations between configurations (Huang et al., 2022). The limited ability of practitioners to interact and collaborate with the BO machinery (Kumar et al., 2022) thus runs the risk of failing to use valuable domain expertise, or alienating knowledgeable practitioners altogether.

While knowledge injection beyond what is natively supported by the GP kernel is crucial to further increase the efficiency of Bayesian optimization, so far no approach exists that allows the integration of knowledge regarding arbitrary properties of the function at hand. To address this gap, we propose an intuitive framework that is agnostic to the user's prior, and that effectively allows the user to reshape the Gaussian process at will to mimic their held beliefs. Our novel framework for *Collaborative Bayesian Optimization* (`CoLaBO`) diverges from the typical assumption of Gaussian posteriors, and

is applicable to any Monte Carlo acquisition function (Wilson et al., 2017; 2018; Balandat et al., 2020). Formally, we make the following contributions:

1. We introduce `COLaBO`, a framework for incorporating arbitrary user knowledge into Bayesian optimization in the form of an additional prior on the surrogate, orthogonal to the conventional prior on the kernel hyperparameters,
2. We demonstrate that the proposed framework is generally applicable to Monte Carlo acquisition functions and various priors over arbitrary properties of the objective function,
3. We empirically show that `COLaBO` substantially accelerates optimization when the prior is accurate, while being only marginally hindered by a severely misleading prior.

2 BACKGROUND

We outline Bayesian optimization, Gaussian Processes and Monte Carlo (MC) acquisition functions, as well as the concept of a prior over the optimum.

2.1 BAYESIAN OPTIMIZATION

We consider the problem of optimizing a black-box function f across a set of feasible inputs $\mathcal{X} \subset \mathbb{R}^d$:

$$\mathbf{x}^* \in \arg \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}). \quad (1)$$

We assume that $f(\mathbf{x})$ is expensive to evaluate and can potentially only be observed through a noise-corrupted estimate, $y_{\mathbf{x}}$, where $y_{\mathbf{x}} = f(\mathbf{x}) + \epsilon$, $\epsilon \sim \mathcal{N}(0, \sigma_{\epsilon}^2)$ for some noise level σ_{ϵ}^2 . In this setting, we wish to maximize f in an efficient manner. Bayesian optimization (BO) aims to globally maximize f by an initial design and thereafter sequentially choosing new points \mathbf{x}_n for some iteration n , creating the data $\mathcal{D}_n = \mathcal{D}_{n-1} \cup \{(\mathbf{x}_n, y_n)\}$ (Brochu et al., 2010; Shahriari et al., 2016; Garnett, 2022). After each new observation, BO constructs a probabilistic surrogate model $p(f|\mathcal{D}_n)$ (Snoek et al., 2012; Hutter et al., 2011; Bergstra et al., 2011a; Müller et al., 2023) and uses that surrogate to build an acquisition function $\alpha(\mathbf{x}; \mathcal{D}_n)$ (Jones et al., 1998; Srinivas et al., 2010; Hennig & Schuler, 2012a) which selects the next query.

2.2 GAUSSIAN PROCESSES

When constructing the surrogate, the most common choice is a *Gaussian process* (GP) (Rasmussen & Williams, 2006). The GP utilizes a covariance function k , which encodes a prior belief for the smoothness of f , and determines how previous observations influence prediction. Given observations \mathcal{D}_n at iteration n , the Gaussian posterior $p(f|\mathcal{D}_n)$ over the objective is characterized by the posterior mean $\mu_n(\mathbf{x}, \mathbf{x}')$ and (co-)variance $\Sigma_n(\mathbf{x}, \mathbf{x}')$ of the GP:

$$\begin{aligned} \mu_n(\mathbf{x}) &= \mathbf{k}_n(\mathbf{x})^\top (\mathbf{K}_n + \sigma_{\epsilon}^2 \mathbf{I})^{-1} \mathbf{y} \\ \Sigma_n(\mathbf{x}, \mathbf{x}') &= k(\mathbf{x}, \mathbf{x}') - \mathbf{k}_n(\mathbf{x})^\top (\mathbf{K}_n + \sigma_{\epsilon}^2 \mathbf{I})^{-1} \mathbf{k}_n(\mathbf{x}'), \end{aligned}$$

where $(\mathbf{K}_n)_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$, $\mathbf{k}_n(\mathbf{x}) = [k(\mathbf{x}, \mathbf{x}_1), \dots, k(\mathbf{x}, \mathbf{x}_n)]^\top$ and σ_{ϵ}^2 is the noise variance.

For various applications in BO and beyond, samples from the posterior are required either directly for optimization (Eriksson et al., 2019; Kandasamy et al., 2018) through Thompson sampling (Thompson, 1933), or to estimate auxiliary quantities of interest (Hernández-Lobato et al., 2015; Neiswanger et al., 2021; Hvarfner et al., 2023). For a finite set of k query locations ($\mathbf{X} = \mathbf{x}_1, \dots, \mathbf{x}_k$), the standard method of generating samples is via a location-scale transform of Gaussian random variables, $f(\mathbf{X}) = \mu_n(\mathbf{X}) + \mathbf{L}\epsilon$, where \mathbf{L} is the Cholesky decomposition of \mathbf{K} and $\epsilon \sim \mathcal{N}(0, \mathbf{I})$. Unfortunately, the standard approach is intrinsically non-scalable, incurring a $\mathcal{O}(k^3)$ cost due to the aforementioned matrix decomposition.

2.3 DECOUPLED POSTERIOR SAMPLING

To remedy the issue of scalability in posterior sampling, $\mathcal{O}(k)$ weight-space approximations based on Random Fourier Features (RFF) (Rahimi & Recht, 2007) obtain approximate (continuous) function

draws $\tilde{f}(\mathbf{x}) = \sum_{i=1}^m w_i \phi_i(\mathbf{x})$, where $\phi_i(\mathbf{x}) = \frac{2}{\ell} (\boldsymbol{\psi}_i^\top \mathbf{x} + b_i)$. The random variables $w_i \sim \mathcal{N}(0, 1)$, $b_i \sim \mathcal{U}(0, 2\pi)$, and $\boldsymbol{\psi}_i$ are sampled proportional to the spectral density of k .

While achieving scalability, the seminal RFF approach by Rahimi & Recht (2007) suffers from the issue of variance starvation (Mutny & Krause, 2018; Wang et al., 2018; Wilson et al., 2020). As a remedy, Wilson et al. (2020) decouple the draw of functions from the approximate posterior $p(\tilde{f}|\mathcal{D})$ into a generally more accurate draw from the prior $p(\tilde{f})$, followed by a deterministic update determined by the data:

$$(\tilde{f}|\mathcal{D})(\mathbf{x}) \stackrel{d}{=} \underbrace{\tilde{f}(\mathbf{x})}_{\text{draw from prior}} + \underbrace{\mathbf{k}_n(\mathbf{x})^\top (\mathbf{K}_n + \sigma_\epsilon^2 \mathbf{I})^{-1} (\mathbf{y} - \tilde{f}(\mathbf{x}) - \boldsymbol{\epsilon})}_{\text{deterministic update}} \quad (2)$$

Eq. 2 deviates from the distribution-first approach that is typically prevalent in GPs in favor of a variable-first approach utilizing Matheron’s rule (Journel & Huijbregts, 1976).

2.4 MONTE CARLO ACQUISITION FUNCTIONS

Acquisition functions act on the surrogate model to quantify the utility of a point in the search space. They encode a trade-off between exploration and exploitation, typically using a greedy heuristic to do so. A simple and computationally cheap heuristic is Expected Improvement (EI) (Jones et al., 1998; Bull, 2011). For a noiseless function and a current best observation y_n^* , the EI acquisition function is $\alpha_{EI}(\mathbf{x}) = \mathbb{E}_{y_{\mathbf{x}}} [(y_n^* - y_{\mathbf{x}})^+]$. For noisy problem settings, a noise-adapted variant of EI (Letham et al., 2018) is frequently considered, where both the incumbent y_n^* and the upcoming query $y_{\mathbf{x}}$ are substituted for the non-observable noiseless incumbent f_n^* and noiseless upcoming query $f_{\mathbf{x}}$. Other frequently used acquisition functions are the Upper Confidence Bound (UCB) (Srinivas et al., 2012), Probability of Improvement (PI) (Kushner, 1964) and Knowledge Gradient (KG) (Frazier et al., 2009). Information-theoretic acquisition functions consider the mutual information to select the next query $\alpha_{MI}(\mathbf{x}) = I((\mathbf{x}, y_{\mathbf{x}}); *|\mathcal{D}_n)$, where $*$ can entail either the optimum \mathbf{x}_* as in (Predictive) Entropy Search (ES/PES) (Hennig & Schuler, 2012b; Hernández-Lobato et al., 2014), the optimal value f_* as in Max-value Entropy Search (MES) (Wang & Jegelka, 2017; Moss et al., 2021) or the tuple (\mathbf{x}_*, f_*) for Joint Entropy Search (JES) (Hvarfner et al., 2022a; Tu et al., 2022).

All the aforementioned acquisition functions compute expectations $\mathbb{E}_{f_{\mathbf{x}}}$ (or alternatively $\mathbb{E}_{y_{\mathbf{x}}}$) over some utility $u(f_{\mathbf{x}})$ of the output (Wilson et al., 2017; 2018), which typically have simple, or even closed-form, solutions for Gaussian posteriors. However, approximating the expectation through Monte Carlo integration has proven useful in the context of batch optimization (Wilson et al., 2018), efficient acquisition function approximation (Balandat et al., 2020), and non-Gaussian posteriors (Astudillo & Frazier, 2021). By sampling over possible outputs $f_{\mathbf{x}}$ and utilizing the reparametrization trick (Kingma & Welling, 2013; Rezende et al., 2014), utilities u can be easily computed across a larger set of applications and be optimized to greater accuracy.

2.5 PRIOR OVER THE OPTIMUM

A prior over the optimum (Souza et al., 2021; Hvarfner et al., 2022b; Mallik et al., 2023) is a user-specified belief $\pi : \mathcal{X} \rightarrow \mathbb{R}$ that represents the user-held subjective likelihood that a given configuration \mathbf{x} is optimal. Formally,

$$\pi(\mathbf{x}) = \mathbb{P} \left(\mathbf{x} = \arg \max_{\mathbf{x}'} f(\mathbf{x}') \right). \quad (3)$$

This prior is generally considered to be independent of observed data, but rather a result of previous experimentation or anecdotal evidence. Regions that the user expects to contain the optimum will typically have a high value, but this does not exclude the chance of the user belief $\pi(\mathbf{x})$ to be inaccurate, or even misleading. Lastly, we require π to be strictly positive in all of \mathcal{X} , which suggests that any point included in the search space may be optimal.

3 METHODOLOGY

We now introduce COLaBO, the first Bayesian-principled BO framework that flexibly allows users to *collaborate* with the optimizer by injecting prior knowledge about the objective that substantially

exceeds the type of prior knowledge natively supported by GPs. In Sec. 3.2, we introduce and derive a novel prior over function properties, which yields a surrogate model conditioned on the user belief π . Thereafter, in Sec. 3.3, we demonstrate how the hierarchical prior integrates with MC acquisition functions. Lastly, in Sec. 3.4, we introduce additional practical considerations to further enhance the performance of COLABO.

3.1 PRIOR OVER FUNCTION PROPERTIES

We consider the typical GP prior over functions $p(f) = \mathcal{GP}(\mu, \Sigma)$, where the characteristics of f , such as smoothness and output magnitude, are fully defined by the kernel k (and its associated hyperparameters θ , which are omitted for brevity). We seek to inject an additional, user-defined prior belief over f into the GP, such as the prior over the optimum in Sec. 2.5, $\pi(\mathbf{x}) = \mathbb{P}(\mathbf{x} = \arg \max_{\mathbf{x}'} f(\mathbf{x}'))$. We observe that π is an (unnormalized) density for function draws f , since f can be mapped to its $\arg \max$. Subsequently, it can serve as an input to $\pi(\mathbf{x})$. In turn, this observation reveals the posterior $p(f|\pi)$ through which $\pi(\mathbf{x})$ is defined. Formally,

$$\pi(\mathbf{x}) = \mathbb{P}\left(\mathbf{x} = \arg \max_{\mathbf{x}'} f(\mathbf{x}')\right) = \int_f \delta_*(\mathbf{x}|f)p(f|\pi)df \quad (4)$$

where the Dirac delta operator $\delta_*(\mathbf{x}|f)$ maps f to its $\arg \max$.

It is worthwhile to reflect on the meaning of $p(f|\pi)$ in greater detail. For a user-specified belief $\pi(f)$ over a *property* of f , such as the location of its $\arg \max \pi_{\mathbf{x}_*}(f)$ (Hvarfner et al., 2022b), the value of its optimum $\max \pi_{f_*}(f)$ (Nguyen & Osborne, 2020), or pairwise preference relations $\mathbf{x}' \succcurlyeq \mathbf{x}$, $\pi_{\succcurlyeq \mathbf{x}}(f)$ (Huang et al., 2022), the belief over such properties propagate to the belief over f itself. Concretely, if the user belief $\pi_{f_*}(f)$ asserts that the maximal value lies within $C_1 < \max f^* < C_2$, the resulting distribution over f should only contain functions whose \max falls within this range. Using rejection sampling, functions which disobey this criterion are filtered out, which yields the posterior $p(f|\pi_{f_*})$.

Having exemplified how user beliefs impact the prior over functions $p(f)$, we hereby introduce the formal definition of a user belief over a function property:

Definition 3.1 (User Belief over Functions). *The user belief over functions is $\pi(f) \propto p(f|\pi)/p(f)$.*

Intuitively, this definition states that the unnormalized density $\pi(f)$ specifies how likely a function f is to come from $p(f|\pi)$, relative to the vanilla $p(f)$. Moreover, it allows us to view π through the lens of a likelihood. Having defined the role of π and the posterior over functions it produces, a natural question arises: *How is the belief-weighted posterior $p(f|\pi)$ updated once observations \mathcal{D} are obtained?*

Since the data \mathcal{D} is independent of the prior (the data generation process is intrinsically unaffected by the belief held by the user), successive application of Bayes' rule yields the following posterior conditioned on \mathcal{D} and π ,

$$p(f|\mathcal{D}, \pi) = \frac{p(\mathcal{D}, \pi|f)p(f)}{p(\mathcal{D}, \pi)} = \frac{p(\mathcal{D}|f)p(\pi|f)p(f)}{p(\mathcal{D})p(\pi)} = \frac{p(f|\pi)}{p(f)}p(f|\mathcal{D}) \propto \pi(f)p(f|\mathcal{D}), \quad (5)$$

where the right side of the proportionality in Eq. 5 suggests an intuitive generation process for samples $(f|\mathcal{D}, \pi)$ to approximate the density $p(f|\mathcal{D}, \pi)$. Utilizing the pathwise update from Eq. 2, we note that given an approximate draw \tilde{f} from the prior, the subsequent data-dependent update is deterministic. Recalling Eq. 2 and assuming independence between π and \mathcal{D} , π only affects the draw from the prior, whereas \mathcal{D} only affects the update. Consequently, we obtain

$$(\tilde{f}|\mathcal{D}, \pi)(\mathbf{x}) \stackrel{d}{=} \underbrace{(\tilde{f}|\pi)(\mathbf{x})}_{\text{draw from prior}} + \underbrace{\mathbf{k}_n(\mathbf{x})^\top (\mathbf{K}_n + \sigma_\epsilon^2 \mathbf{I})^{-1} (\mathbf{y} - (\tilde{f}|\pi)(\mathbf{x}) - \epsilon)}_{\text{deterministic update}}, \quad (6)$$

where $(\tilde{f}|\pi) \sim p(f)\pi(\tilde{f})$ are once again obtained using rejection sampling on draws from $p(\tilde{f})$. Figure 1 displays this in detail: given the typical GP prior over functions *and* a user belief over the optimum, we obtain a posterior over functions $p(\tilde{f}|\pi)$ before having observed any data (top right). Samples from the approximate prior $p(\tilde{f})$ (light blue) are re-sampled proportionally to their

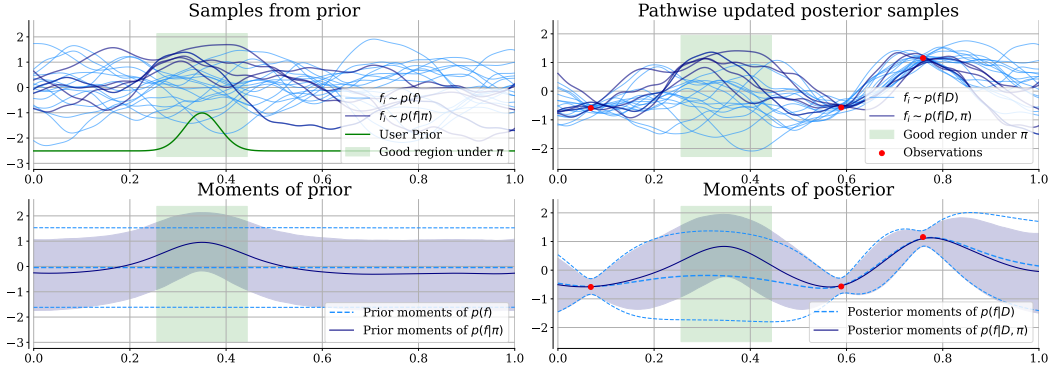


Figure 1: (Top left) Draws from the prior $p(f)$ (light blue) and the belief-weighted prior $p(f|\pi)$ whose members are likely to have their optimum within the green region. (Top right) Pathwise updated draws based on observed data. As the green region is distant from the observed data, samples are almost unaffected by the data in this region. (Bottom left) Exact mean and standard deviation ($\mu_{\mathbf{x}}, \sigma_{\mathbf{x}}$) of $p(f)$ and estimated mean and standard deviation of $p(f|\pi)$. (Bottom right) Exact $p(f|\mathcal{D})$ and estimated $p(f|\pi, \mathcal{D})$. As $p(f|\pi)$ constitutes of functions whose optimum is located within the green region the resulting model has a higher mean and lower variance within this region. Moreover, $p(f|\pi)$ globally displays lower upside variance compared to the vanilla GP.

probability of occurring under the prior $\pi_{\mathbf{x}_*}(\tilde{f})$ in green, leaving samples ($\tilde{f}|\pi_{\mathbf{x}_*}$) in navy blue, which are highly probable under $\pi_{\mathbf{x}_*}$. Once data is obtained, these samples are updates according to Eq. 6, which preserves the shape of the samples far away from observed data and yields the desired posterior.

3.2 PRIOR-WEIGHTED MONTE CARLO ACQUISITION FUNCTIONS

Naturally, neither the belief-weighted prior $p(f|\pi)$ nor the belief-weighted posterior $p(f|\mathcal{D}, \pi)$ have a closed-form expression. Both are inherently non-Gaussian for non-uniform beliefs. As such, we resort to MC acquisition functions to compute utilities that are amenable to BO. In the subsequent section, we focus on the prevalent acquisition functions EI, and MES.

Expected Improvement The computation of the Monte Carlo approximation of EI within the `ColaBO` framework requires only minor adaptations of the original MC acquisition function. By definition, MC-EI assigns utility u as $u_{\text{EI}}(f(\mathbf{x})) = \max(f_n^* - f(\mathbf{x}), 0)$, which yields

$$\alpha_{\text{EI}}(\mathbf{x}; \mathcal{D}) = \mathbb{E}_{f_{\mathbf{x}}|\mathcal{D}}[u_{\text{EI}}(f_{\mathbf{x}})] \approx \quad (7)$$

$$\sum_{\ell} \max(f_n^* - f_{\mathbf{x}}^{(\ell)}, 0), \quad f_{\mathbf{x}}^{(\ell)} \sim p(f(\mathbf{x})|\mathcal{D}). \quad (8)$$

Utilizing importance sampling, we can compute the MC-EI under the `ColaBO` posterior accordingly,

$$\alpha_{\text{EI}}(\mathbf{x}; \mathcal{D}, \pi) = \mathbb{E}_{f_{\mathbf{x}}|\mathcal{D}, \pi}[u_{\text{EI}}(f_{\mathbf{x}})] = \quad (9)$$

$$\int_f u_{\text{EI}}(f_{\mathbf{x}}) \pi(f) p(f|\mathcal{D}) df \approx \sum_{\ell} \pi(f^{(\ell)}) \max(f_n^* - f^{(\ell)}, 0), \quad f^{(\ell)} \sim p(f(\mathbf{x})|\mathcal{D}), \quad (10)$$

wherein samples in Eq. 10 are drawn from the prior, evaluated on π , and pathwise updated. In Figure 2, we demonstrate how `ColaBO`-EI differs from vanilla MC-EI for an identical posterior as in Figure 1. By computing α_{EI} from samples biased by $\pi_{\mathbf{x}_*}$, `ColaBO` substantially biases the search towards regions that are good under $\pi_{\mathbf{x}_*}$. Derivations for PI and KG are analogous to that of EI.




Figure 2: (Top) Draws from the posterior $p(f|\mathcal{D})$ (light blue) and the belief-weighted posterior $p(f|\pi, \mathcal{D})$ with a prior π located in the green region. (Bottom) Vanilla MC-EI and `ColaBO` MC-EI, resulting from computing the acquisition function from many sample draws from $p(f|\pi, \mathcal{D})$.

Algorithm 1 CoLaBO iteration

-
- 1: **Input:** User prior π , number of function samples L , current data \mathcal{D} , tempering factor β
 - 2: **Output:** Next query location \mathbf{x}' .
 - 3: **for** $\ell \in \{1, \dots, L\}$ **do**
 - 4: $\pi^{(\ell)} = \pi(\tilde{f}^{(\ell)}; \beta, n)$, $\tilde{f}^{(\ell)} \sim p(\tilde{f})$ \triangleright Sample functions and evaluate on π
 - 5: $(\tilde{f}^{(\ell)}|\mathcal{D}) = \text{PathwiseUpdate}(\tilde{f}^{(\ell)}, \mathcal{D})$ \triangleright Per-sample update as in Eq. 6
 - 6: **end for**
 - 7: $p(\tilde{f}|\mathcal{D}, \pi) \approx \sum_{\ell} \pi^{(\ell)}(\tilde{f}^{(\ell)}|\mathcal{D})$ \triangleright Form MC estimate of posterior
 - 8: $\mathbf{x}' = \arg \max_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{p(\tilde{f}|\mathcal{D}, \pi)}[u(\tilde{f}_{\mathbf{x}})]$ \triangleright Maximize MC acquisition
-

Max-Value Entropy Search We derive a CoLaBO-MES acquisition function by first considering the definition of the entropy, $\mathbf{H}[p(y_{\mathbf{x}}|\mathcal{D})] = \mathbb{E}_{y_{\mathbf{x}}|\mathcal{D}}[-\log p(y_{\mathbf{x}}|\mathcal{D})]$. When considering the belief-weighted posterior, we further condition the posterior on π and obtain

$$\alpha_{\text{MES}}(\mathbf{x}) = \mathbb{E}_{f_{*}|\mathcal{D}, \pi} [\mathbb{E}_{y_{\mathbf{x}}|\mathcal{D}, \pi, f_{*}} [\log p(y_{\mathbf{x}}|\mathcal{D}, \pi, f_{*})]] - \mathbb{E}_{y_{\mathbf{x}}|\mathcal{D}, \pi} [\log p(y_{\mathbf{x}}|\mathcal{D}, \pi)] \quad (11)$$

$$\propto \mathbb{E}_{f_{*}|\mathcal{D}, \pi} [\mathbb{E}_{f_{\mathbf{x}}|\mathcal{D}, \pi} [\mathbb{E}_{y_{\mathbf{x}}|f_{\mathbf{x}}} [\log p(y_{\mathbf{x}}|f_{\mathbf{x}}, \pi, f_{*})]]] - \mathbb{E}_{f_{\mathbf{x}}|\mathcal{D}, \pi} [\mathbb{E}_{y_{\mathbf{x}}|f_{\mathbf{x}}} [\log p(y_{\mathbf{x}}|f_{\mathbf{x}}, \pi)]] \quad (12)$$

$$\approx \frac{1}{Z_J} \sum_{j=1}^J \sum_{\ell=1}^L \sum_{k=1}^K \log p(y_{\mathbf{x}}^{(k)}|f_{\mathbf{x}}^{(\ell)}, f_{*}^{(j)}) \pi(f^{(\ell)}) \pi(f^{(j)}) - \sum_{\ell=1}^L \sum_{k=1}^K \log p(y_{\mathbf{x}}^{(k)}|f_{\mathbf{x}}^{(\ell)}) \pi(f^{(\ell)}), \quad (13)$$

where Z_J is a normalizing constant $\sum_J \pi(f^{(j)})$ brought on by sampling optimal values, $y_{\mathbf{x}}|f_{\mathbf{x}}$ can trivially be obtained by sampling Gaussian noise $\varepsilon \sim \mathcal{N}(0, \sigma_{\varepsilon}^2)$ to a noiseless observation $f_{\mathbf{x}}|\mathcal{D}$ in the innermost expectation, and $f_{\mathbf{x}}$ and f_{*} are obtained through the pathwise sampling procedure outlined in Eq. 6. The samples are evaluated on $p((y_{\mathbf{x}}|f_{\mathbf{x}}), (y_{\mathbf{x}}|f_{\mathbf{x}}, f_{*}))$. As evident by Eq. 13, π notably affects the posterior distribution of both the observations $y_{\mathbf{x}}$ and the optimal values f_{*} . PES and JES are derived analogously. However, these acquisition function require conditioning on additional, simulated data and consequently, additional pathwise updates, to compute.

3.3 UPDATING THE USER BELIEF

Intuitively, a reasonable user should increasingly trust the data over their prior belief, as more data is acquired. For this purpose, we utilize a generalized likelihood (Grünwald, 2012; Friel & Pettitt, 2008) (a.k.a. power posterior or tempered likelihood) framework, commonly used in presence of model misspecification, to initially increase the impact of π on the posterior predictive, and gradually decrease its impact over time. As such, we define the *tempered user belief* as

$$\pi(f; \beta, n) = \pi(f)^{\beta/n}, \quad (14)$$

where β specifies the tempering rate and n the iteration number. A similar approach of diminishing the impact of π over time has been adopted in prior works (Hvarfner et al., 2022b; Souza et al., 2021). We note, however, that CoLaBO differs in that the decay acts on the *model* in a principled Bayesian manner, as opposed to heuristically on the *acquisition function* in prior work.

3.4 PRACTICAL CONSIDERATIONS

CoLaBO introduces additional flexibility to MC-based BO acquisition functions. The CoLaBO framework deviates from vanilla (Q)MC acquisition functions (Wilson et al., 2017; Balandat et al., 2020) by utilizing approximate sample functions from the posterior, as opposed to pointwise draws from the posterior predictive and the reparametrization trick (Kingma & Welling, 2013; Rezende et al., 2014). CoLaBO holds three shortcomings not prevalent in vanilla MC acquisition functions: (1) it cannot utilize Quasi-MC in the draws from the predictive posterior (only in the RFF weights), (2) it cannot fix the base samples (Balandat et al., 2020) drawn from the posterior for acquisition function consistency across the search space, and (3) the RFF approximation of $p(f)$ introduces bias. In Sec. 4.1, we empirically display the impact of these shortcomings. Notably, while acquisition function optimization no longer enjoys improved accuracy resulting from reparametrization, the acquisition function can still be effectively optimized due to the fact that CoLaBO backpropagates through quantities computed as sums of smooth functions.

4 RESULTS

We evaluate the performance of `ColaBO` on various tasks, using priors over the optimum π_{x_*} obtained from known optima on synthetic tasks, as well as from prior work (Mallik et al., 2023) on realistic tasks. We consider two variants of `ColaBO`: one using `LogEI` (Ament et al., 2023), a numerically stable, smoothed `logsumexp` transformation of EI with analogous derivation, and one variant using `MES`. We benchmark against the vanilla variants of each acquisition function, as well as π_{BO} (Hvarfner et al., 2022b) and sampling from π_{x_*} . All acquisition functions are implemented in `BoTorch` (Balandat et al., 2020) using a Matérn-5/2 kernel (Matérn, 1960) and MAP hyperparameter estimation. Unless stated otherwise, `ColaBO` and π_{BO} are initialized with the mode of the prior followed by 2 Sobol samples, whereas `LogEI` and `MES` are conventionally initialized with $D + 1$ Sobol samples. For all experiments, we run `ColaBO` with $\beta = N/2.5$, where N is the total number of iterations, in order to make the prior influence approximately equal across experiments with differing budgets. We investigate the sensitivity to β in Appendix B.3. The complete experimental setup is presented in detail in Appendix A, and our code is publicly available at <https://anonymous.4open.science/r/colabo-6712/>.

4.1 APPROXIMATION QUALITY OF THE COLABO FRAMEWORK

Firstly, we demonstrate the approximation quality of `ColaBO` *without* user priors to assert its accuracy compared to a vanilla MC acquisition function. To facilitate comparison, we randomly sample 10 points on the Hartmann (3D) function, and optimize `LogEI` with a large budget. We subsequently optimize `ColaBO-LogEI` on the same set of points and compare the $\arg \max$ to the solution found by the gold standard. Figure 3 displays the (\log_{10}) Euclidian distance between the $\arg \max$ of `LogEI` and its `ColaBO` variant. We note that, for small amounts (≤ 256) of posterior samples, the error induced by RFF bias is relatively low, which is evidenced by all RFF variants being roughly equal in distance to the true acquisition function optimizer.




Figure 3: Mean and 1/4 standard deviation of MC-induced errors of `ColaBO-LogEI` relative vanilla `LogEI` as measured by the distance to the $\arg \max$ of the acquisition function on Hartmann (3D) on 10 randomly sampled points for 40 seeds.

4.2 SYNTHETIC FUNCTIONS WITH KNOWN PRIORS

We adapt a similar evaluation protocol to Hvarfner et al. (2022b), and evaluate `ColaBO` for two types of user beliefs for synthetic tasks: well-located and poorly located priors over the optimal location, designed to emulate a well-informed and poorly-informed practitioner, respectively. The well-located prior is offset by a small (10%) amount from the optimum, and the poorly located prior is maximally offset, while retaining its mode inside the search space. Complete details on the priors can be found in Appendix A.3. On well-located priors, both `ColaBO-LogEI` and `ColaBO-MES` demonstrate substantially improved performance relative to their vanilla counterparts, comparable to π_{BO} on all benchmarks. On poorly located priors, `ColaBO` demonstrates superior robustness, recovering the performance of the vanilla acquisition function within the maximal budget of $20D$ iterations and clearly outperforming π_{BO} , which more frequently misled by the poor prior. In Appendix B.1, we also demonstrate `ColaBO` utilizing (accurate) beliefs over the optimal value: similarly to Figure 4, `ColaBO` yields increased efficiency relative to baselines, albeit not as substantial. Moreover, we demonstrate its usage with batch evaluations on well-located priors in Sec. B.2, showing that the drop in performance from batching evaluations is marginal at worst.

4.3 HYPERPARAMETER TUNING TASKS

Lastly, we evaluate `ColaBO` on three $4D$ deep learning HPO tasks from the PD1 (Wang et al., 2023) benchmarking suite. While the optima for these tasks are ultimately unknown, we utilize the priors




Figure 4: Performance on synthetic functions with well-located priors. Both ColaBO-LogEI and ColaBO-MES offer drastic speed-ups over their vanilla variants, and offer similar performance to π BO. The ranking of ColaBO acquisition functions are generally consistent with their respective vanilla variants. This is most prominent on Rosenbrock (6D), where ColaBO-MES struggles similarly to vanilla MES.




Figure 5: Performance on poorly located priors. ColaBO acquisition functions are more robust than π BO, as it frequently recovers the performance of the vanilla acquisition function before the total budget is depleted. ColaBO-LogEI struggles marginally on Hartmann (6D), whereas ColaBO-MES recovers or surpasses baseline performance on all tasks.

provided in MF-Prior-Bench¹ (Mallik et al., 2023), which are intended to provide a good starting point for further optimization. To emulate a realistic HPO setting, we consider a smaller optimization budget of $10D$ iterations, and initialize all methods that utilize user beliefs with only one DoE sample, that being the mode of the prior. The two ColaBO variants perform best in this evaluation, producing the best terminal performance on two tasks (WMT, LM1B), with all methods being tied on the third (CIFAR). ColaBO also finds highly competitive configurations in as few as ten evaluations. π BO offers similarly competitive performance, but does not offer the rapid initial progress of ColaBO, especially on LM1B. Crucially, ColaBO demonstrates substantial and consistent speed-ups compared to its vanilla counterparts, surpassing the terminal performance of the baseline within a third of the budget on both WMT and LM1B.

¹<https://github.com/automl/mf-prior-bench>




Figure 6: Performance on the 4D PD1 hyperparameter tuning tasks of various deep learning pipelines. ColaBO drastically accelerates optimization initially, finding configurations with close to terminal performance in ten evaluations or less. π BO offers similarly competitive performance, but does not offer the rapid initial progress of ColaBO on CIFAR and LM1B. The y-axis is broken to increase visibility.

5 RELATED WORK

In BO, auxiliary prior information can be conveyed in multiple ways. We outline meta learning/transfer learning for BO based on data from previous experiments, and data-less approaches, where exotic priors based on acquired knowledge are used.

Learning from Previous Experiments Transfer learning and meta learning for BO aims to automatically extract and use knowledge from prior executions of BO by pre-training the model on data acquired from previous executions (Swersky et al., 2013; Wistuba et al., 2015; Perrone et al., 2019; Feurer et al., 2015; 2018; Rothfuss et al., 2021a;b; Wistuba & Grabocka, 2021; Feurer et al., 2022). Typically, meta- and transfer learning exploit relevant previous data for training the GP for the current task while retaining predictive uncertainty to account for imperfect task correlation.

Expert Priors over Function Optimum Few previous works have proposed to inject explicit prior distributions over the location of an optimum into BO. In these cases, users explicitly define a prior that encodes their beliefs on where the optimum is more likely to be located. Bergstra et al. (2011a) suggest an approach that supports prior beliefs from a fixed set of distributions, which affects the very initial stage of optimization. However, this approach cannot be combined with standard acquisition functions. BOPrO (Souza et al., 2021) employs a similar structure that combines the user-provided prior distribution with a data-driven model into a pseudo-posterior. From the pseudo-posterior, configurations are selected using the EI acquisition function, using the formulation in Bergstra et al. (2011a). π BO (Hvarfner et al., 2022b) suggests a general-purpose prior-weighted acquisition function, where the influence of the prior decreases over time. They provide convergence guarantees for when the framework is applied to the EI acquisition function. While effective, none of these approaches act on the surrogate model in a Bayesian-principled fashion, but strictly as heuristics. Moreover, they solely focus on priors over optimal inputs π_{x_*} , thus offering less utility than Colabo.

Priors over Optimal Value Similarly few works have addressed the issue of auxiliary knowledge of the optimal value. Both Jeong & Kim (2021) and Nguyen & Osborne (2020) propose altering the GP and accompanying it with tailored acquisition functions. Jeong & Kim (2021) employ variational inference, proposing distinct variational families depending on the type of knowledge pertaining to the optimal value. Nguyen & Osborne (2020) use a parabolic transformation of the output space to ensure an upper bound is preserved. Unlike Colabo, neither of these methods is general enough to accompany arbitrary user priors to guide the optimization.

6 CONCLUSION, LIMITATIONS AND FUTURE WORK

We presented Colabo, a flexible BO framework that allows practitioners to inject beliefs over function properties in a Bayesian-principled manner, allowing for increased efficiency in the BO procedure. Colabo works across a collection of MC acquisition functions, inheriting their flexibility in batch optimization and ability to work with non-Gaussian posteriors. It demonstrates competitive performance for well-located priors, using them to substantially accelerate optimization. Moreover, it retains approximately baseline performance when applied to detrimental priors, demonstrating greater robustness than π BO.

Colabo crucially relies on multiple steps of MC. While flexible, this approach drives computational expense in order to assert sufficient accuracy, requiring tens of seconds per evaluation to achieve desired accuracy, depending on the size of the benchmark. Moreover, obtaining draws from π_{x_*} scales exponentially in the dimensionality of the prior. While practitioners are unlikely to specify priors over more than a handful of variables, Colabo may become impractical when priors of higher dimensionality are employed.

Paths for future work could involve more accurate and efficient sampling procedures (Lin et al., 2023) from the belief-weighted prior, as well as variational (Titsias, 2009) or pre-trainedp Müller et al. (2022); Müller et al. (2023) approaches to obtain a representative belief-biased model with an analytical posterior. This would likely bring down the runtime of Colabo and broaden its potential use. Lastly, applying Colabo to multi-fidelity optimization (Kandasamy et al., 2016; Mallik et al., 2023) offers an additional avenue for increased efficiency which would further increase its viability on costly deep learning pipelines.

REFERENCES

- Sebastian Ament, Samuel Daulton, David Eriksson, Maximilian Balandat, and Eytan Bakshy. Unexpected improvements to expected improvement for bayesian optimization. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=lvvAG6j9PE>.
- Raul Astudillo and Peter Frazier. Bayesian optimization of function networks. *Advances in neural information processing systems*, 34:14463–14475, 2021.
- M. Balandat, B. Karrer, D. R. Jiang, S. Daulton, B. Letham, A. G. Wilson, and E. Bakshy. Botorch: A framework for efficient monte-carlo bayesian optimization. In *Advances in Neural Information Processing Systems*, 2020. URL <http://arxiv.org/abs/1910.06403>.
- J. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl. Algorithms for hyper-parameter optimization. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Weinberger (eds.), *Proceedings of the 25th International Conference on Advances in Neural Information Processing Systems (NeurIPS’11)*, pp. 2546–2554, 2011a.
- James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for Hyper-Parameter Optimization. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 24. Curran Associates, Inc., 2011b.
- E. Brochu, V. Cora, and N. de Freitas. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv:1012.2599v1 [cs.LG]*, 2010.
- Adam D. Bull. Convergence rates of efficient global optimization algorithms. 12:2879–2904, 2011.
- R. Calandra, N. Gopalan, A. Seyfarth, J. Peters, and M. Deisenroth. Bayesian gait optimization for bipedal locomotion. In P. Pardalos and M. Resende (eds.), *Proceedings of the Eighth International Conference on Learning and Intelligent Optimization (LION’14)*, 2014.
- Adel Ejjeh, Leon Medvinsky, Aaron Councilman, Hemang Nehra, Suraj Sharma, Vikram Adve, Luigi Nardi, Eriko Nurvitadhi, and Rob A Rutenbar. Hpv2fpga: Enabling true hardware-agnostic fpga programming. In *Proceedings of the 33rd IEEE International Conference on Application-specific Systems, Architectures, and Processors*, 2022.
- David Eriksson, Michael Pearce, Jacob Gardner, Ryan D Turner, and Matthias Poloczek. Scalable global optimization via local Bayesian optimization. In *Advances in Neural Information Processing Systems*, pp. 5496–5507, 2019. URL <http://papers.nips.cc/paper/8788-scalable-global-optimization-via-local-bayesian-optimization.pdf>.
- M. Feurer, Jost Tobias Springenberg, and F. Hutter. Initializing bayesian hyperparameter optimization via meta-learning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pp. 1128–1135, 2015.
- M. Feurer, B. Letham, F. Hutter, and E. Bakshy. Practical transfer learning for bayesian optimization. *ArXiv abs/1802.02219*, 2018.
- Matthias Feurer, Benjamin Letham, Frank Hutter, and Eytan Bakshy. Practical transfer learning for Bayesian optimization. *arXiv preprint 1802.02219*, 2022.
- Peter Frazier, Warren Powell, and Savas Dayanik. The knowledge-gradient policy for correlated normal beliefs. *INFORMS journal on Computing*, 21(4):599–613, 2009.
- N. Friel and A. N. Pettitt. Marginal likelihood estimation via power posteriors. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 70(3):589–607, 2008. ISSN 13697412, 14679868. URL <http://www.jstor.org/stable/20203843>.
- R. Garnett. *Bayesian Optimization*. Cambridge University Press, 2022. Available for free at <https://bayesoptbook.com/>.

- Ryan-Rhys Griffiths and José Miguel Hernández-Lobato. Constrained bayesian optimization for automatic chemical design using variational autoencoders. *Chemical Science*, 2020.
- Peter Grünwald. The safe bayesian. In Nader H. Bshouty, Gilles Stoltz, Nicolas Vayatis, and Thomas Zeugmann (eds.), *Algorithmic Learning Theory*, pp. 169–183, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. ISBN 978-3-642-34106-9.
- P. Hennig and C. Schuler. Entropy search for information-efficient global optimization. 98888(1): 1809–1837, 2012a.
- P. Hennig and C. J. Schuler. Entropy search for information-efficient global optimization. *Journal of Machine Learning Research*, 13(1):1809–1837, June 2012b. ISSN 1532-4435.
- J. M. Hernández-Lobato, M. W. Hoffman, and Z. Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In *Advances in Neural Information Processing Systems*, 2014. URL <https://proceedings.neurips.cc/paper/2014/file/069d3bb002acd8d7dd095917f9efe4cb-Paper.pdf>.
- José Miguel Hernández-Lobato, Michael Gelbart, Matthew Hoffman, Ryan Adams, and Zoubin Ghahramani. Predictive entropy search for bayesian optimization with unknown constraints. In *International conference on machine learning*, pp. 1699–1707. PMLR, 2015.
- Daolang Huang, Louis Filstroff, Petrus Mikkola, Runkai Zheng, and Samuel Kaski. Bayesian optimization augmented with actively elicited expert knowledge, 2022.
- F. Hutter, H. Hoos, and K. Leyton-Brown. Sequential model-based optimization for general algorithm configuration. In C. Coello (ed.), *Proceedings of the Fifth International Conference on Learning and Intelligent Optimization (LION'11)*, volume 6683, pp. 507–523, 2011.
- Carl Hvarfner, Frank Hutter, and Luigi Nardi. Joint entropy search for maximally-informed bayesian optimization. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, 2022a.
- Carl Hvarfner, Danny Stoll, Artur Souza, Marius Lindauer, Frank Hutter, and Luigi Nardi. PiBO: Augmenting Acquisition Functions with User Beliefs for Bayesian Optimization. In *International Conference on Learning Representations*, 2022b.
- Carl Hvarfner, Erik Hellsten, Frank Hutter, and Luigi Nardi. Self-correcting bayesian optimization through bayesian active learning. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=dX9MjUtP1A>.
- Taewon Jeong and Heeyoung Kim. Objective bound conditional gaussian process for bayesian optimization. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 4819–4828. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/jeong21a.html>.
- D. Jones, M. Schonlau, and W. Welch. Efficient global optimization of expensive black box functions. 13:455–492, 1998.
- A G Journel and C J Huijbregts. Mining geostatistics, Jan 1976.
- K. Kandasamy, G. Dasarathy, J. Oliva, J. Schneider, and B. Póczos. Gaussian Process Bandit Optimisation with Multi-fidelity Evaluations. In D. Lee, M. Sugiyama, U. von Luxburg, I. Guyon, and R. Garnett (eds.), *Proceedings of the 30th International Conference on Advances in Neural Information Processing Systems (NeurIPS'16)*, pp. 992–1000, 2016.
- K. Kandasamy, A. Krishnamurthy, J. Schneider, and B. Póczos. Parallelised Bayesian optimisation via Thompson sampling. In A. Storkey and F Perez-Cruz (eds.), *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 84, pp. 133–142. Proceedings of Machine Learning Research, 2018.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2013. URL <https://arxiv.org/abs/1312.6114>.

- Arun Kumar, Santu Rana, Alistair Shilton, and Svetha Venkatesh. Human-ai collaborative bayesian optimisation. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 16233–16245. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/6751611b394a3464cea53eed91cf163c-Paper-Conference.pdf.
- H. J. Kushner. A New Method of Locating the Maximum Point of an Arbitrary Multippeak Curve in the Presence of Noise. *Journal of Basic Engineering*, 86(1):97–106, 03 1964. ISSN 0021-9223. doi: 10.1115/1.3653121. URL <https://doi.org/10.1115/1.3653121>.
- B. Letham, K. Brian, G. Ottoni, and E. Bakshy. Constrained Bayesian optimization with noisy experiments. *Bayesian Analysis*, 2018.
- Jihao Andreas Lin, Padhy Shreyas Janz David Hernández-Lobato José Miguel Antorán, Javier, and Alexander Terenin. Sampling from gaussian process posteriors using stochastic gradient descent. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=Sf9goJtTCE>.
- Marius Lindauer, Katharina Eggenberger, Matthias Feurer, André Biedenkapp, Difan Deng, Carolin Benjamins, Tim Ruhkopf, René Sass, and Frank Hutter. Smac3: A versatile bayesian optimization package for hyperparameter optimization. *Journal of Machine Learning Research*, 23(54):1–9, 2022. URL <http://jmlr.org/papers/v23/21-0888.html>.
- Neeratyoy Mallik, Edward Bergman, Carl Hvarfner, Danny Stoll, Maciej Janowski, Marius Lindauer, Luigi Nardi, and Frank Hutter. Priorband: Practical hyperparameter optimization in the age of deep learning. *arXiv preprint 2306.12370*, 2023.
- B. Matérn. Spatial variation. *Meddelanden fran Statens Skogsforskningsinstitut*, 1960.
- Matthias Mayr, Carl Hvarfner, Konstantinos Chatzilygeroudis, Luigi Nardi, and Volker Krueger. Learning skill-based industrial robot tasks with user priors. *IEEE 18th International Conference on Automation Science and Engineering*, 2022. URL <https://arxiv.org/abs/2208.01605>.
- Henry B. Moss, David S. Leslie, Javier Gonzalez, and Paul Rayson. Gibbon: General-purpose information-based bayesian optimisation. *Journal of Machine Learning Research*, 22(235):1–49, 2021. URL <http://jmlr.org/papers/v22/21-0120.html>.
- Samuel Müller, Noah Hollmann, Sebastian Pineda Arango, Josif Grabocka, and Frank Hutter. Transformers can do bayesian inference. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=KSugKcbNf9>.
- Samuel Müller, Matthias Feurer, Noah Hollmann, and Frank Hutter. PFNs4BO: In-context learning for Bayesian optimization. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 25444–25470. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/muller23a.html>.
- Mojmir Mutny and Andreas Krause. Efficient high dimensional bayesian optimization with additivity and quadrature fourier features. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/4e5046fc8d6a97d18a5f54beaed54dea-Paper.pdf.
- L. Nardi, D. Koeplinger, and K. Olukotun. Practical design space exploration. In *2019 IEEE 27th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*, pp. 347–358. IEEE, 2019.
- Willie Neiswanger, Ke Alexander Wang, and Stefano Ermon. Bayesian algorithm execution: Estimating computable properties of black-box functions using mutual information. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8005–8015. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/neiswanger21a.html>.

- Vu Nguyen and Michael A. Osborne. Knowing the what but not the where in Bayesian optimization. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 7317–7326. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/nguyen20d.html>.
- C. Oh, E. Gavves, and M. Welling. BOCK : Bayesian optimization with cylindrical kernels. In *International Conference on Machine Learning*, pp. 3865–3874, 2018.
- V. Perrone, H. Shen, M. Seeger, C. Archambeau, and R. Jenatton. Learning search spaces for bayesian optimization: Another view of hyperparameter transfer learning. In *Advances in Neural Information Processing Systems*, 2019.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In J. Platt, D. Koller, Y. Singer, and S. Roweis (eds.), *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2007. URL https://proceedings.neurips.cc/paper_files/paper/2007/file/013a006f03dbc5392effeb8f18fda755-Paper.pdf.
- C. Rasmussen and C. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In Eric P. Xing and Tony Jebara (eds.), *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pp. 1278–1286, Beijing, China, 22–24 Jun 2014. PMLR. URL <https://proceedings.mlr.press/v32/rezende14.html>.
- Jonas Rothfuss, Vincent Fortuin, Martin Josifoski, and Andreas Krause. Pacoh: Bayes-optimal meta-learning with pac-guarantees. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 9116–9126, 2021a.
- Jonas Rothfuss, Dominique Heyn, Jinfan Chen, and Andreas Krause. Meta-learning reliable priors in the function space. In *Advances in Neural Information Processing Systems*, volume 34, 2021b.
- Binxin Ru, Xingchen Wan, Xiaowen Dong, and Michael Osborne. Interpretable neural architecture search via bayesian optimisation with weisfeiler-lehman kernels. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=j9Rv7qdXjd>.
- 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, 2016.
- L. Smith. A disciplined approach to neural network hyper-parameters: Part 1–learning rate, batch size, momentum, and weight decay. *arXiv preprint arXiv:1803.09820*, 2018.
- J. Snoek, H. Larochelle, and R. Adams. Practical Bayesian optimization of machine learning algorithms. In P. Bartlett, F. Pereira, C. Burges, L. Bottou, and K. Weinberger (eds.), *Proceedings of the 26th International Conference on Advances in Neural Information Processing Systems (NeurIPS’12)*, pp. 2960–2968, 2012.
- A. Souza, L. Nardi, L. Oliveira, K. Olukotun, M. Lindauer, and F. Hutter. Bayesian optimization with a prior for the optimum. In *Machine Learning and Knowledge Discovery in Databases. Research Track - European Conference, ECML PKDD 2021, Bilbao, Spain, September 13-17, 2021, Proceedings, Part III*, volume 12977 of *Lecture Notes in Computer Science*, pp. 265–296. Springer, 2021.
- N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In J. Fürnkranz and T. Joachims (eds.), *Proceedings of the 27th International Conference on Machine Learning (ICML’10)*, pp. 1015–1022. Omnipress, 2010.
- N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger. Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, 58(5):3250–3265, May 2012. ISSN 1557-9654. doi: 10.1109/tit.2011.2182033. URL <http://dx.doi.org/10.1109/TIT.2011.2182033>.

- K. Swersky, J. Snoek, and R. Adams. Multi-task Bayesian optimization. In C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger (eds.), *Proceedings of the 27th International Conference on Advances in Neural Information Processing Systems (NeurIPS'13)*, pp. 2004–2012, 2013.
- W. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- Michalis Titsias. Variational learning of inducing variables in sparse gaussian processes. In David van Dyk and Max Welling (eds.), *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, volume 5 of *Proceedings of Machine Learning Research*, pp. 567–574, Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA, 16–18 Apr 2009. PMLR. URL <https://proceedings.mlr.press/v5/titsias09a.html>.
- Ben Tu, Axel Gandy, Nikolas Kantas, and Behrang Shafei. Joint entropy search for multi-objective bayesian optimization. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=ZChgD80oGds>.
- Q. Wang, Y. Ming, Z. Jin, Q. Shen, D. Liu, M. J. Smith, K. Veeramachaneni, and H. Qu. Atmseer: Increasing transparency and controllability in automated machine learning. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, pp. 1–12. Association for Computing Machinery, 2019.
- Zi Wang and Stefanie Jegelka. Max-value entropy search for efficient bayesian optimization. In *International Conference on Machine Learning (ICML)*, 2017.
- Zi Wang, Clement Gehring, Pushmeet Kohli, and Stefanie Jegelka. Batched large-scale bayesian optimization in high-dimensional spaces. In Amos Storkey and Fernando Perez-Cruz (eds.), *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pp. 745–754. PMLR, 09–11 Apr 2018. URL <https://proceedings.mlr.press/v84/wang18c.html>.
- Zi Wang, George E. Dahl, Kevin Swersky, Chansoo Lee, Zachary Nado, Justin Gilmer, Jasper Snoek, and Zoubin Ghahramani. Pre-trained Gaussian processes for Bayesian optimization. *arXiv preprint arXiv:2109.08215*, 2023.
- C. White, W. Neiswanger, and Y. Savani. BANANAS: Bayesian optimization with neural architectures for neural architecture search. In Q. Yang, K. Leyton-Brown, and Mausam (eds.), *Proceedings of the Thirty-Fifth Conference on Artificial Intelligence (AAAI'21)*, pp. 10293–10301. Association for the Advancement of Artificial Intelligence, AAAI Press, 2021.
- James Wilson, Frank Hutter, and Marc Deisenroth. Maximizing acquisition functions for bayesian optimization. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL <https://proceedings.neurips.cc/paper/2018/file/498f2c21688f6451d9f5fd09d53edda7-Paper.pdf>.
- James T. Wilson, Riccardo Moriconi, Frank Hutter, and Marc Peter Deisenroth. The reparameterization trick for acquisition functions, 2017. URL <https://arxiv.org/abs/1712.00424>.
- James T. Wilson, Viacheslav Borovitskiy, Alexander Terenin, Peter Mostowsky, and Marc Peter Deisenroth. Efficiently sampling functions from gaussian process posteriors. In *International Conference on Machine Learning*, 2020. URL <https://arxiv.org/abs/2002.09309>.
- M. Wistuba, N. Schilling, and L. Schmidt-Thieme. Hyperparameter search space pruning - A new component for sequential model-based hyperparameter optimization. In A. Appice, P. Rodrigues, V. Costa, J. Gama, A. Jorge, and C. Soares (eds.), *Machine Learning and Knowledge Discovery in Databases (ECML/PKDD'15)*, volume 9285, pp. 104–119, 2015.
- Martin Wistuba and Josif Grabocka. Few-shot bayesian optimization with deep kernel surrogates. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=bJxgv5C3sYc>.

A EXPERIMENTAL SETUP

A.1 MODEL

We outline the model used and the budget allocated to the various MC approximations involved with COLA_{BO}. For all experiments, we utilize MAP estimation of the hyperparameters, and update the hyperparameters at every iteration of BO. All hyperparameters - lengthscale, outputscale and observation noise ($\theta = \{\ell, \sigma_\epsilon^2, \sigma_f^2\}$) are given conventional $\mathcal{LN}(0, 1)$ prior, applied on normalized inputs and standardized outputs. Furthermore, we fit the constant c of the mean function, assigning it a $\mathcal{N}(0, 1)$ prior as well. In Tab. 1, we display the parameters of the MC approximations for various tasks. *No. f* is the maximal number of functions used in the MC computation of the acquisition function. *No. Resamples* is the number of initial posterior draws maximally used for the re-sampling of functions from the posterior $p(f|\pi)$. Lastly, *No. f_** is the number of optimal values used in the computation of COLA_{BO}-MES.

Task	No. f	No. RFFs	No. Resamples	No. f_*
Synthetic Good	768	2048	$1.5 * 10^3$	32
Synthetic Bad	768	2048	$1.5 * 10^5$	32
PD1	512	4096	$2 * 10^5$	32
Appendix	512	1024	10^5	32

Table 1: Budget-related parameters of the Monte Carlo approximations for all tasks.

A.2 BENCHMARKS

We outline the benchmarks used, their search spaces and the amount of synthetic noise added. When adding noise, we intend for the ratio of noise variance to total output range to be approximately equal across benchmarks.

Task	Dimensionality	σ_ϵ	Search space
Hartmann (4D)	4	0.25	$[0, 1]^D$
Levy (5D)	5	0.5	$[-5, 5]^D$
Hartmann (6D)	6	0.25	$[0, 1]^D$
Rosenbrock (6D)	6	5	$[-2.048, 2.048]^D$
Stybtang (7D)	7	1	$[-4, 4]^D$

Table 2: Benchmarks used for the Bayesian optimization experiments.

A.3 PRIORS

For synthetic benchmarks, the approximate optima of all included functions can be obtained in advance. Thus, the correctness of the prior is ultimately known in advance. For a function of dimensionality d with optimum at \mathbf{x}_* , the well-located prior is constructed by using a sampling an offset direction ϵ and scaling the offset by a dimensionality- and quality-specific term $c(d, q) = q\sqrt{d}$. For the well-located prior on synthetic tasks, we use $q = 0.1$, which implies that the optimum is located 10% of the distance across the search space away from the optimum, and construct a Gaussian prior as

$$\pi_{\mathbf{x}_*}(\mathbf{x}) \sim \mathcal{N}(\mathbf{x}_* + c_d \epsilon / \|\epsilon\|, \sigma_s), \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}). \quad (15)$$

with $\sigma_s = 25\%$ for all tasks and prior qualities. For our 20 runs of the well-located prior, this procedure yields us 20 unique priors per quality type, with identical offsets from the true optimum. No priors with a mode outside the search space were allowed, such priors were simply replaced. For the misinformed priors, we set $q = 1$, guaranteeing that the mode of the prior will be outside




Figure 7: ColaBO with priors over the optimal value. Terminal performance substantially increases on 3 out of 5 benchmarks (Levy, Hartmann (6D), Stybtang), and is approximately preserved on the final two. ColaBO-MES improves marginally more than ColaBO-LogEI when utilizing a prior π_{f_*} over the optimal value.

of the search space, and subsequently relocating to the edge of the search space by its shortest path. Priors for all tasks are displayed in Tab. 3. For the PD1 tasks, the location for the priors were obtained from MF-Prior-Bench (<https://github.com/automl/mf-prior-bench>). However, these priors require offsetting in order to not be too strong, thus making subsequent BO obsolete. PD1 priors are provided in $[0, 1]$ -normalized space for simplicity.

Task	Location	Offset, good	Offset, bad	σ_s
Hartmann (4D)	$[0.19, 0.19, 0.56, 0.26]$	$0.1\sqrt{D}$	max	0.25
Levy (5D)	$[1]^D$	$1\sqrt{D}$	max	2.5
Hartmann (6D)	$[0.20, 0.15, 0.48, 0.28, 0.31, 0.66]$	$0.1\sqrt{D}$	max	0.25
Rosenbrock (6D)	$[1]^D$	$0.4096\sqrt{D}$	max	1.024
Stybtang (7D)	$[-2.9]^D$	$0.8\sqrt{D}$	max	2
PD1-WMT	$[0.90, 0.69, 0.02, 0.97]$	$0.05\sqrt{D}$	N/A	0.25
PD1-CIFAR	$[1, 0.80, 0.0, 0.0]$	$0.05\sqrt{D}$	N/A	0.25
PD1-LM1B	$[0.91, 0.67, 0.36, 0.85]$	$0.05\sqrt{D}$	N/A	0.25

Table 3: π_{x_*} for synthetic BO tasks of both prior qualities and PD1.

B ADDITIONAL EXPERIMENTS

We provide complementary experiments to those introduced in the main paper. Firstly, we display results when ColaBO is used with a prior π_{f_*} over the optimal value in Sec. B.1. In Sec. B.2, we demonstrate ColaBO’s extensibility to batch evaluations, seamlessly extending the work of (?). Lastly, we perform a sensitivity analysis on the tempering parameter β in Sec. B.3.

B.1 MAX-VALUE PRIORS

We evaluate ColaBO with priors over the optimal value π_{f_*} in Fig. 7. For each task, we place a Gaussian prior over the optimal value, centering it exactly at the optimal value. Notably, such a prior substantially influences the exploration-exploitation trade-off; if the prior suggests that the incumbent has a value close to the optimal one, we are encouraged to exploit as samples with well-above-optimal values in exploratory will be discarded. Conversely, we are heavily encouraged to explore if the current best observation holds a value that we believe is far from optimal. On Hartmann (6D), we can see this behavior at play. Initial performance is poorer for ColaBO than their respective baselines, presumably due to above-average exploration, but terminal performance is better.

B.2 BATCH EVALUATIONS

We evaluate ColaBO on batch evaluations, utilizing the sequential greedy technique for MC acquisition functions from Wilson et al. (2018). Drop-off from sequential to batch evaluations is not evident from the plots, as ordering between sequential and batch varies with the benchmark. While unpredictable, we speculate that the altered exploration-exploitation trade-off provided by the batched acquisition function is occasionally beneficial in the presence of auxiliary user beliefs π_{x_*} .




Figure 8: $q = 1$ (sequential) and $q = 3$ batch evaluation on a subset of synthetic functions with well-located priors for ColaBO-LogEI and ColaBO-MES. Total function evaluations are plotted for both sequential and batched variants, leaving them with the same number of total function evaluations.

B.3 β SENSITIVITY ANALYSIS

We consider the effect of the β hyperparameter of ColaBO introduced in Section 3.3, controlling the tempering of the prior. At first glance, a larger beta should force the optimizer to retain queries close to the prior mode for longer. However, having the optimum of function draws at the mode of the prior does not preclude other regions of the search space from being almost as good. As such, a very large β can still allow for queries that are distant from the good region under the prior. However, a large beta does force very local queries early on, and require a larger amount of data to retract back to a conventional posterior. In Figs. 11, 9, 12 and 10, we display the performance of ColaBO-LogEI and ColaBO-MES for three values of beta and both well-located and misleading priors for a subset of the benchmarks from Sec. 4.2.




Figure 9: β ablation on synthetic functions with well-located priors for ColaBO-MES. The magnitude of beta does not substantially impact performance, as rankings are not consistent across benchmarks. As the prior is not perfectly located, a smaller value of β can reasonably yield a better terminal performance.




Figure 10: β ablation on synthetic functions with poorly located priors for ColaBO-MES. ColaBO-MES displays very good robustness to poorly located priors, recovering approximately default performance for any β .




Figure 11: β ablation on synthetic functions with well-located priors for ColaBO-LogEI. Similarly to ColaBO-MES, The magnitude of beta does not substantially impact performance, as rankings are not consistent across benchmarks. As the prior is not perfectly located, a smaller value of β can yield a better terminal performance.




Figure 12: β ablation on synthetic functions with poorly located priors for ColaBO-LogEI. The method displays good robustness to poorly located priors on two out of three benchmarks, recovering approximately default performance for any β . On Hartmann (6D), recovery is slower, and baseline performance is not recouped within the optimization budget for any of the tested β .