

Bayesian Optimization with a Prior for the Optimum

Artur Souza¹, Luigi Nardi^{2,3}, Leonardo B. Oliveira¹,
Kunle Olukotun³, Marius Lindauer⁴, and Frank Hutter^{5,6}

¹ Universidade Federal de Minas Gerais, {arturluis, leob}@dcc.ufmg.br

² Lund University, luigi.nardi@cs.lth.se

³ Stanford University, {lnardi, kunle}@stanford.edu

⁴ Leibniz University Hannover, lindauer@tnt.uni-hannover.de

⁵ University of Freiburg, fh@cs.uni-freiburg.de

⁶ Bosch Center for Artificial Intelligence

Abstract. While Bayesian Optimization (BO) is a very popular method for optimizing expensive black-box functions, it fails to leverage the experience of domain experts. This causes BO to waste function evaluations on bad design choices (e.g., machine learning hyperparameters) that the expert already knows to work poorly. To address this issue, we introduce Bayesian Optimization with a Prior for the Optimum (BOPrO). BOPrO allows users to inject their knowledge into the optimization process in the form of priors about which parts of the input space will yield the best performance, rather than BO’s standard priors over functions, which are much less intuitive for users. BOPrO then combines these priors with BO’s standard probabilistic model to form a pseudo-posterior used to select which points to evaluate next. We show that BOPrO is around $6.67\times$ faster than state-of-the-art methods on a common suite of benchmarks, and achieves a new state-of-the-art performance on a real-world hardware design application. We also show that BOPrO converges faster even if the priors for the optimum are not entirely accurate and that it robustly recovers from misleading priors.

1 Introduction

Bayesian Optimization (BO) is a data-efficient method for the joint optimization of design choices that has gained great popularity in recent years. It is impacting a wide range of areas, including hyperparameter optimization [41,10], AutoML [20], robotics [5], computer vision [30], Computer Go [6], hardware design [31,23], and many others. It promises greater automation so as to increase both product quality and human productivity. As a result, BO is also established in large tech companies, e.g., Google [13] and Facebook [1].

Nevertheless, domain experts often have substantial prior knowledge that standard BO cannot easily incorporate so far [44]. Users can incorporate prior knowledge by narrowing the search space; however, this type of hard prior can lead to poor performance by missing important regions. BO also supports a prior

over functions $p(f)$, e.g., via a kernel function. However, this is not the prior domain experts have: they often know which ranges of hyperparameters tend to work best [36], and are able to specify a probability distribution $p_{\text{best}}(\mathbf{x})$ to quantify these priors; e.g., many users of the Adam optimizer [21] know that its best learning rate is often in the vicinity of 1×10^{-3} (give or take an order of magnitude), yet may not know the accuracy they can achieve in a new application. Similarly, Clarke *et al.* [7] derived neural network hyperparameter priors for image datasets based on their experience. In these cases, users know potentially good values for a new application, but cannot be certain about them.

As a result, many competent users instead revert to manual search, which can fully incorporate their prior knowledge. A recent survey showed that most NeurIPS 2019 and ICLR 2020 papers reported having tuned hyperparameters used manual search, with only a very small fraction using BO [4]. In order for BO to be adopted widely, and accelerate progress in the ML community by tuning hyperparameters faster and better, it is, therefore, crucial to devise a method that fully incorporates expert knowledge about the location of the optimum areas into BO. In this paper, we introduce Bayesian Optimization with a Prior for the Optimum (BOPrO), a BO variant that combines priors for the optimum with a probabilistic model of the observations made. Our technical contributions are:

- We introduce *Bayesian Optimization with a Prior over the Optimum*, short *BOPrO*, which allows users to inject priors that were previously difficult to inject into BO, such as Gaussian, exponential, multimodal, and multivariate priors for the location of the optimum. To ensure robustness against misleading priors, BOPrO gives more importance to the data-driven model as iterations progress, gradually forgetting the prior.
- BOPrO’s model bridges the gap between the well-established Tree-structured Parzen Estimator (TPE) methodology, which is based on Parzen kernel density estimators, and standard BO probabilistic models, such as Gaussian Processes (GPs). This is made possible by using the Probability of Improvement (PI) criterion to derive from BO’s standard posterior over functions $p(f | (\mathbf{x}_i, y_i)_{i=1}^t)$ the probability of an input \mathbf{x} leading to good function values.
- We demonstrate the effectiveness of BOPrO on a comprehensive set of synthetic benchmarks and real-world applications, showing that knowledge about the locality of an optimum helps BOPrO to achieve similar performance to current state-of-the-art on average $6.67\times$ faster on synthetic benchmarks and $1.59\times$ faster on a real-world application. BOPrO also achieves similar or better final performance on all benchmarks.

BOPrO is publicly available as part of the HyperMapper framework⁷.

⁷ <https://github.com/luinardi/hypermapper/wiki/prior-injection>

2 Background

2.1 Bayesian Optimization

Bayesian Optimization (BO) is an approach for optimizing an unknown function $f : \mathcal{X} \rightarrow \mathbb{R}$ that is expensive to evaluate over an input space \mathcal{X} . In this paper, we aim to minimize f , i.e., find $\mathbf{x}^* \in \arg \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$. BO approximates \mathbf{x}^* with an optimal sequence of evaluations $\mathbf{x}_1, \mathbf{x}_2, \dots \in \mathcal{X}$ that maximizes a utility metric, with each new \mathbf{x}_{t+1} depending on the previous function values y_1, y_2, \dots, y_t at $\mathbf{x}_1, \dots, \mathbf{x}_t$. BO achieves this by building a posterior on f based on the set of evaluated points. At each iteration, a new point is selected and evaluated based on the posterior, and the posterior is updated to include the new point $(\mathbf{x}_{t+1}, y_{t+1})$.

The points explored by BO are dictated by the acquisition function, which attributes an utility to each $\mathbf{x} \in \mathcal{X}$ by balancing the predicted value and uncertainty of the prediction for each \mathbf{x} [39]. In this work, as the acquisition function we choose Expected Improvement (EI) [29], which quantifies the expected improvement over the best function value found so far:

$$EI_{y_{inc}}(\mathbf{x}) := \int_{-\infty}^{\infty} \max(y_{inc} - y, 0) p(y|\mathbf{x}) dy, \quad (1)$$

where y_{inc} is the incumbent function value, i.e., the best objective function value found so far, and $p(y|\mathbf{x})$ is given by a probabilistic model, e.g., a GP. Alternatives to EI would be Probability of Improvement (PI) [24], upper-confidence bounds (UCB) [42], entropy-based methods (e.g. [17]), and knowledge gradient [45].

2.2 Tree-structured Parzen Estimator

The Tree-structured Parzen Estimator (TPE) method is a BO approach introduced by Bergstra *et al.* [3]. Whereas the standard probabilistic model in BO directly models $p(y|\mathbf{x})$, the TPE approach models $p(\mathbf{x}|y)$ and $p(y)$ instead.⁸ This is done by constructing two parametric densities, $g(\mathbf{x})$ and $b(\mathbf{x})$, which are computed using configurations with function value below and above a given threshold y^* , respectively. The separating threshold y^* is defined as a quantile γ of the observed function values. Thus, TPE builds a density $g(\mathbf{x})$ using the best γ configurations, and a density $b(\mathbf{x})$ using the remaining configurations. TPE uses the densities $g(\mathbf{x})$ and $b(\mathbf{x})$ to define $p(\mathbf{x}|y)$ as:

$$p(\mathbf{x}|y) = g(\mathbf{x})I(y < y^*) + b(\mathbf{x})(1 - I(y < y^*)), \quad (2)$$

where $I(y < y^*)$ is 1 when $y < y^*$ and 0 otherwise. Bergstra *et al.* [3] show that the parametrization of the generative model $p(\mathbf{x}, y) = p(\mathbf{x}|y)p(y)$ facilitates the computation of EI as it leads to $EI_{y^*}(\mathbf{x}) \propto g(\mathbf{x})/b(\mathbf{x})$ and, thus, to maximize EI , they need to simply find the configurations that maximize the ratio $g(\mathbf{x})/b(\mathbf{x})$.

⁸ Technically, the model does not parameterize $p(y)$, since it is computed based on the observed data points, which are heavily biased towards low values due to the optimization process. Instead, it parameterizes a dynamically changing $p_t(y)$, which helps to constantly challenge the model to yield better observations.

3 BO with a Prior for the Optimum

We now describe our BOPrO approach, which allows domain experts to inject prior knowledge about the locality of an optimum into the optimization. BOPrO combines this user-defined prior with a probabilistic model that captures the likelihood of the observed data $\mathcal{D}_t = (\mathbf{x}_i, y_i)_{i=1}^t$. BOPrO is independent of the probabilistic model being used; it can be freely combined with, e.g., Gaussian processes (GPs), random forests, or Bayesian NNs.

3.1 BOPrO Priors

BOPrO allows users to inject prior knowledge w.r.t. promising areas into BO. This is done via a prior distribution that informs where in the input space \mathcal{X} we expect to find good $f(\mathbf{x})$ values. A point is considered “good” if it leads to low function values, and potentially to a global optimum. We denote the prior distribution $P_g(\mathbf{x})$, where g denotes that this is a prior on good points and $\mathbf{x} \in \mathcal{X}$ is a given point. Examples of priors are shown in Figures 2 and 3, additional examples of continuous and discrete priors are shown in Appendices A and D, respectively. Similarly, we define a prior on where in the input space we expect to have “bad” points. Although we could have a user-defined probability distribution $P_b(\mathbf{x})$, we aim to keep the decision-making load on users low and thus, for simplicity, only require the definition of $P_g(\mathbf{x})$ and compute $P_b(\mathbf{x}) = 1 - P_g(\mathbf{x})$.⁹ $P_g(\mathbf{x})$ is normalized to $[0, 1]$ by min-max scaling before computing $P_b(\mathbf{x})$.

In practice, \mathbf{x} contains several dimensions but it is difficult for domain experts to provide a joint prior distribution $P_g(\mathbf{x})$ for all of them. However, users can typically easily specify, e.g., sketch out, univariate or bivariate distributions for continuous dimensions or provide a list of probabilities for discrete dimensions. In BOPrO, users can define a complex multivariate distribution, but we expect the standard use case to be users choosing to specify univariate distributions, implicitly assuming a prior that factors as $P_g(\mathbf{x}) = \prod_{i=1}^D P_g(x_i)$, where D is the number of dimensions in \mathcal{X} and x_i is the i -th input dimension of \mathbf{x} . To not assume unrealistically complex priors and to mimic our expected use case, we use factorized priors in our experiments; in Appendix E we show that these factorized priors can in fact lead to similar BO performance as multivariate priors.

3.2 Model

Whereas the standard probabilistic model in BO, e.g., a GP, quantifies $p(y|\mathbf{x})$ directly, that model is hard to combine with the prior $P_g(\mathbf{x})$. We therefore introduce a method to translate the standard probabilistic model $p(y|\mathbf{x})$ into a model that is easier to combine with this prior. Similar to the TPE work

⁹ We note that for continuous spaces, $P_b(\mathbf{x})$ is not a probability distribution as it does not integrate to 1 and therefore is only a pseudo-prior. For discrete spaces, we normalize $P_b(\mathbf{x})$ so that it sums to 1 and therefore is a proper distribution and prior.




Fig. 1: Our model is composed by a probabilistic model and the probability of improving over the threshold f_γ , i.e., right tail of the Gaussian. The black curve is the probabilistic model’s mean and the shaded area is the model’s variance.

described in Section 2.2, our generative model combines $p(\mathbf{x}|y)$ and $p(y)$ instead of directly modeling $p(y|\mathbf{x})$.

The computation we perform for this translation is to quantify the probability that a given input \mathbf{x} is “good” under our standard probabilistic model $p(y|\mathbf{x})$. As in TPE, we define configurations as “good” if their observed y -value is below a certain quantile γ of the observed function values (so that $p(y < f_\gamma) = \gamma$). We in addition exploit the fact that our standard probabilistic model $p(y|\mathbf{x})$ has a Gaussian form, and under this Gaussian prediction we can compute the probability $\mathcal{M}_g(\mathbf{x})$ of the function value lying below a certain quantile using the standard closed-form formula for PI [24]:

$$\mathcal{M}_g(\mathbf{x}) = p(f(\mathbf{x}) < f_\gamma | \mathbf{x}, \mathcal{D}_t) = \Phi \left(\frac{f_\gamma - \mu_{\mathbf{x}}}{\sigma_{\mathbf{x}}} \right), \quad (3)$$

where $\mathcal{D}_t = (\mathbf{x}_i, y_i)_{i=1}^t$ are the evaluated configurations, $\mu_{\mathbf{x}}$ and $\sigma_{\mathbf{x}}$ are the predictive mean and standard deviation of the probabilistic model at \mathbf{x} , and Φ is the standard normal CDF, see Figure 1. Note that there are two probabilistic models here:

1. The standard probabilistic model of BO, with a structural prior over functions $p(f)$, updated by data \mathcal{D}_t to yield a posterior over functions $p(f|\mathcal{D}_t)$, allowing us to quantify the probability $\mathcal{M}_g(\mathbf{x}) = p(f(\mathbf{x}) < f_\gamma | \mathbf{x}, \mathcal{D}_t)$ in Eq. (3).¹⁰
2. The TPE-like generative model that combines $p(y)$ and $p(\mathbf{x}|y)$ instead of directly modeling $p(y|\mathbf{x})$.

¹⁰ We note that the structural prior $p(f)$ and the optimum-prior $P_g(\mathbf{x})$ provide orthogonal ways to input prior knowledge. $p(f)$ specifies our expectations about the structure and smoothness of the function, whereas $P_g(\mathbf{x})$ specifies knowledge about the location of the optimum.

Eq. (3) bridges these two models by using the probability of improvement from BO’s standard probabilistic model as the probability $\mathcal{M}_g(\mathbf{x})$ in TPE’s model. Ultimately, this is a heuristic since there is no formal connection between the two probabilistic models. However, we believe that the use of BO’s familiar, theoretically sound framework of probabilistic modeling of $p(y|x)$, followed by the computation of the familiar PI formula is an intuitive choice for obtaining the probability of an input achieving at least a given performance threshold – exactly the term we need for TPE’s $\mathcal{M}_g(\mathbf{x})$. Similarly, we also define a probability $\mathcal{M}_b(\mathbf{x})$ of \mathbf{x} being bad as $\mathcal{M}_b(\mathbf{x}) = 1 - \mathcal{M}_g(\mathbf{x})$.

3.3 Pseudo-posterior

BOPrO combines the prior $P_g(\mathbf{x})$ in Section (3.1) and the model $\mathcal{M}_g(\mathbf{x})$ in Eq. (3) into a pseudo-posterior on “good” points. This pseudo-posterior represents the updated beliefs on where we can find good points, based on the prior and data that has been observed. The pseudo-posterior is computed as the product:

$$g(\mathbf{x}) \propto P_g(\mathbf{x})\mathcal{M}_g(\mathbf{x})^{\frac{t}{\beta}}, \quad (4)$$

where t is the current optimization iteration, β is an optimization hyperparameter, $\mathcal{M}_g(\mathbf{x})$ is defined in Eq. (3), and $P_g(\mathbf{x})$ is the prior defined in Sec 3.1, rescaled to $[0, 1]$ using min-max scaling. We note that this pseudo-posterior is not normalized, but this suffices for BOPrO to determine the next \mathbf{x}_t as the normalization constant cancels out (c.f. Section 3.5). Since $g(\mathbf{x})$ is not normalized and we include the exponent t/β in Eq. (4), we refer to $g(\mathbf{x})$ as a pseudo-posterior, to emphasize that it is not a standard posterior probability distribution.

The t/β fraction in Eq. (4) controls how much weight is given to $\mathcal{M}_g(\mathbf{x})$. As the optimization progresses, more weight is given to $\mathcal{M}_g(\mathbf{x})$ over $P_g(\mathbf{x})$. Intuitively, we put more emphasis on $\mathcal{M}_g(\mathbf{x})$ as it observes more data and becomes more accurate. We do this under the assumption that the model $\mathcal{M}_g(\mathbf{x})$ will eventually be better than the user at predicting where to find good points. This also allows to recover from misleading priors as we show in Section 4.1; similar to, and inspired by Bayesian models, the data ultimately washes out the prior. The β hyperparameter defines the balance between prior and model, with higher β values giving more importance to the prior.

We note that directly computing Eq (4) can lead to numerical issues. Namely, the pseudo-posterior can reach extremely low values if the $P_g(\mathbf{x})$ and $\mathcal{M}_g(\mathbf{x})$ probabilities are low, especially as t/β grows. To prevent this, in practice, BOPrO uses the logarithm of the pseudo-posterior instead:

$$\log(g(\mathbf{x})) \propto \log(P_g(\mathbf{x})) + \frac{t}{\beta} \cdot \log(\mathcal{M}_g(\mathbf{x})). \quad (5)$$

Once again, we also define an analogous pseudo-posterior distribution on bad \mathbf{x} : $b(\mathbf{x}) \propto P_b(\mathbf{x})\mathcal{M}_b(\mathbf{x})^{\frac{t}{\beta}}$. We then use these quantities to define $p(\mathbf{x}|y)$ as follows:

$$p(\mathbf{x}|y) \propto \begin{cases} g(\mathbf{x}) & \text{if } y < f_\gamma \\ b(\mathbf{x}) & \text{if } y \geq f_\gamma. \end{cases} \quad (6)$$

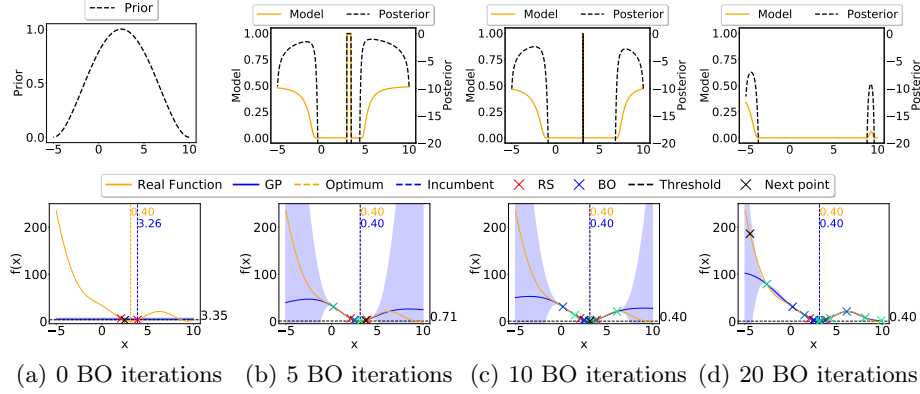


Fig. 2: Breakdown of the prior $P_g(\mathbf{x}) = \mathcal{B}(3, 3)$, the model $M_g(\mathbf{x})$, and the pseudo-posterior $g(\mathbf{x})$ (top row) on the 1D-Branin function (bottom row) and their evolution over the optimization iterations. In early iterations, the pseudo-posterior is high around the optimum, where both prior and model agree there are good points. In later iterations, it vanishes where the model is certain there will be no improvement and is high where there is uncertainty in the model.

3.4 Model and Pseudo-posterior Visualization

We visualize the prior $P_g(\mathbf{x})$, the model $M_g(\mathbf{x})$, and the pseudo-posterior $g(\mathbf{x})$ and their evolution over the optimization iterations for a 1D-Branin function. We define the 1D-Branin by setting the second dimension of the function to the global optimum $x_2 = 2.275$ and optimizing the first dimension. We use a Beta distribution prior $P_g(\mathbf{x}) = \mathcal{B}(3, 3)$, which resembles a truncated Gaussian centered close to the global optimum, and a GP as predictive model. We perform an initial design of $D + 1 = 2$ random points sampled from the prior and then run BOPrO for 20 iterations.

Figure 2 shows the optimization at different stages. Red crosses denote the initial design and blue/green crosses denote BOPrO samples, with green samples denoting later iterations. Figure 2a shows the initialization phase (bottom) and the Beta prior (top). After 5 iterations, in Figure 2b (top), the pseudo-posterior is high near the global minimum, around $x = \pi$, where both the prior $P_g(\mathbf{x})$ and the model $M_g(\mathbf{x})$ agree there are good points. After 10 iterations in Figure 2c (top), there are three regions with high pseudo-posterior. The middle region, where BOPrO is exploiting until the optimum is found, and two regions to the right and left, which will lead to future exploration as shown in Figure 2d (bottom) on the right and left of the global optimum in light green crosses. After 20 iterations, Figure 2d (top), the pseudo-posterior vanishes where the model $M_g(\mathbf{x})$ is certain there will be no improvement, but it is high where there is uncertainty in the GP.

3.5 Acquisition Function

We adopt the EI formulation used in Bergstra *et al.* [3] by replacing their Adaptive Parzen Estimators with our pseudo-posterior from Eq. (4), i.e.:

$$\begin{aligned} EI_{f_\gamma}(\mathbf{x}) &:= \int_{-\infty}^{\infty} \max(f_\gamma - y, 0) p(y|\mathbf{x}) dy = \int_{-\infty}^{f_\gamma} (f_\gamma - y) \frac{p(\mathbf{x}|y)p(y)}{p(\mathbf{x})} dy \\ &\propto \left(\gamma + \frac{b(\mathbf{x})}{g(\mathbf{x})} (1 - \gamma) \right)^{-1}. \end{aligned} \quad (7)$$

The full derivation of Eq. (7) is shown in Appendix B. Eq. (7) shows that to maximize improvement we would like points \mathbf{x} with high probability under $g(\mathbf{x})$ and low probability under $b(\mathbf{x})$, i.e., minimizing the ratio $b(\mathbf{x})/g(\mathbf{x})$. We note that the point that minimizes the ratio for our unnormalized pseudo-posteriors will be the same that minimizes the ratio for the normalized pseudo-posterior and, thus, computing the normalized pseudo-posteriors is unnecessary.

The dynamics of the BOPrO algorithm can be understood in terms of the following proposition (proof in Appendix B):

Proposition 1. *Given f_γ , $P_g(\mathbf{x})$, $P_b(\mathbf{x})$, $\mathcal{M}_g(\mathbf{x})$, $\mathcal{M}_b(\mathbf{x})$, $g(\mathbf{x})$, $b(\mathbf{x})$, $p(\mathbf{x}|y)$, and β as above, then*

$$\lim_{t \rightarrow \infty} \arg \max_{\mathbf{x} \in \mathcal{X}} EI_{f_\gamma}(\mathbf{x}) = \lim_{t \rightarrow \infty} \arg \max_{\mathbf{x} \in \mathcal{X}} \mathcal{M}_g(\mathbf{x}),$$

where EI_{f_γ} is the Expected Improvement acquisition function as defined in Eq. (7) and $\mathcal{M}_g(\mathbf{x})$ is as defined in Eq. (3).

In early BO iterations the prior for the optimum will have a predominant role, but in later BO iterations the model will grow more important, and as Proposition 1 shows, if BOPrO is run long enough the prior washes out and BOPrO *only* trusts the model $\mathcal{M}_g(\mathbf{x})$ informed by the data. Since $\mathcal{M}_g(\mathbf{x})$ is the Probability of Improvement (PI) on the probabilistic model $p(y|\mathbf{x})$ then, in the limit, maximizing $EI_{f_\gamma}(\mathbf{x})$ is equivalent to maximizing the PI acquisition function on the probabilistic model $p(y|\mathbf{x})$. In other words, for high values of t , BOPrO converges to standard BO with a PI acquisition function.

3.6 Putting It All Together

Algorithm 1 shows the BOPrO algorithm. In Line 3, BOPrO starts with a design of experiments (DoE) phase, where it randomly samples a number of points from the user-defined prior $P_g(\mathbf{x})$. After initialization, the BO loop starts at Line 4. In each loop iteration, BOPrO fits the models $\mathcal{M}_g(\mathbf{x})$ and $\mathcal{M}_b(\mathbf{x})$ on the previously evaluated points (Lines 5 and 6) and computes the pseudo-posteriors $g(\mathbf{x})$ and $b(\mathbf{x})$ (Lines 7 and 8 respectively). The EI acquisition function is computed next, using the pseudo-posteriors, and the point that maximizes EI is selected as the next point to evaluate at Line 9. The black-box function evaluation is performed at Line 10. This BO loop is repeated for a predefined number of iterations, according to the user-defined budget B .

Algorithm 1 BOPrO Algorithm. \mathcal{D}_t keeps track of all function evaluations so far: $(\mathbf{x}_i, y_i)_{i=1}^t$.

```

1: Input: Input space  $\mathcal{X}$ , user-defined prior distribution  $P_g(\mathbf{x})$ , quantile  $\gamma$ , initial
   model weight  $\beta$ , and BO budget  $B$ .
2: Output: Optimized point  $\mathbf{x}_{inc}$ .
3:  $\mathcal{D}_1 \leftarrow \text{initialize}(\mathcal{X})$ 
4: for  $t = 1$  to  $B$  do
5:    $\mathcal{M}_g(\mathbf{x}) \leftarrow \text{fit\_model\_good}(\mathcal{D}_t)$  ▷ see Eq. (3)
6:    $\mathcal{M}_b(\mathbf{x}) \leftarrow \text{fit\_model\_bad}(\mathcal{D}_t)$ 
7:    $g(\mathbf{x}) \leftarrow P_g(\mathbf{x}) \cdot \mathcal{M}_g(\mathbf{x})^{\frac{t}{\beta}}$  ▷ see Eq. (4)
8:    $b(\mathbf{x}) \leftarrow P_b(\mathbf{x}) \cdot \mathcal{M}_b(\mathbf{x})^{\frac{t}{\beta}}$ 
9:    $\mathbf{x}_t \in \arg \max_{\mathbf{x} \in \mathcal{X}} EI_{f_\gamma}(\mathbf{x})$  ▷ see Eq. (7)
10:   $y_t \leftarrow f(\mathbf{x}_t)$ 
11:   $\mathcal{D}_{t+1} \leftarrow \mathcal{D}_t \cup (\mathbf{x}_t, y_t)$ 
12: end for
13:  $\mathbf{x}_{inc} \leftarrow \text{compute\_best}(\mathcal{D}_{t+1})$ 
14: return  $\mathbf{x}_{inc}$ 
    
```

4 Experiments

We implement both Gaussian processes (GPs) and random forests (RFs) as predictive models and use GPs in all experiments, except for our real-world experiments (Section 4.3), where we use RFs for a fair comparison. We set the model weight $\beta = 10$ and the model quantile to $\gamma = 0.05$, see our sensitivity hyperparameter study in Appendices I and J. We initialize BOPrO with $D + 1$ random samples from the prior. We optimize our EI acquisition function using a combination of multi-start local search [19] and CMA-ES [15]. We consider four synthetic benchmarks: Branin, SVM, FC-Net, and XGBoost, which are 2, 2, 6, and 8 dimensional, respectively. The last three are part of the Profet benchmarks [22], generated by a generative model built using performance data on OpenML or UCI datasets. See Appendix C for more details.

4.1 Prior Forgetting

We first show that BOPrO can recover from a misleading prior, thanks to our model $\mathcal{M}_g(\mathbf{x})$ and the t/β parameter in Eq. (4). As BO progresses, $\mathcal{M}_g(\mathbf{x})$ becomes more accurate and receives more weight, guiding optimization away from the wrong prior and towards better function values.

Figure 3 shows BOPrO on the 1D Branin function with an exponential prior. Columns (b), (c), and (d) show BOPrO after $D + 1 = 2$ initial samples and 0, 10, 20 BO iterations, respectively. After initialization, as shown in Column (b), the pseudo-posterior is nearly identical to the exponential prior and guides BOPrO towards the region of the space on the right, which is towards the local optimum. This happens until the model $\mathcal{M}_g(\mathbf{x})$ becomes certain there will be no more improvement from sampling that region (Columns (c) and (d)). After

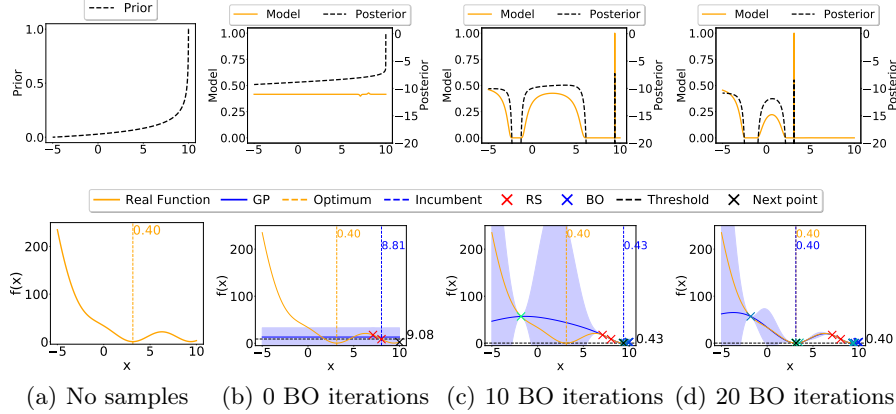


Fig. 3: BOPrO on the 1D Branin function. The leftmost column shows the exponential prior. The other columns show the model and the log pseudo-posterior after 0 (initialization only), 10, and 20 BO iterations. BOPrO forgets the wrong prior on the local optimum and converges to the global optimum.

that, $\mathcal{M}_g(\mathbf{x})$ guides the pseudo-posterior towards exploring regions with high uncertainty. Once the global minimum region is found, the pseudo-posterior starts balancing exploiting the global minimum and exploring regions with high uncertainty, as shown in Figure 3d (bottom). Notably, the pseudo-posterior after $x > 4$ falls to 0 in Figure 3d (top), as the model $\mathcal{M}_g(\mathbf{x})$ is certain there will be no improvement from sampling the region of the local optimum. We provide additional examples of forgetting in Appendix A, and a comparison of BOPrO with misleading priors, no prior, and correct priors in Appendix F.

4.2 Comparison Against Strong Baselines

We build two priors for the optimum in a controlled way and evaluate BOPrO’s performance with these different prior strengths. We emphasize that in practice, manual priors would be based on the domain experts’ expertise on their applications; here, we only use artificial priors to guarantee that our prior is not biased by our own expertise for the benchmarks we used. In practice, users will manually define these priors like in our real-world experiments (Section 4.3).

Our synthetic priors take the form of Gaussian distributions centered near the optimum. For each input $x \in \mathcal{X}$, we inject a prior of the form $\mathcal{N}(\mu_x, \sigma_x^2)$, where μ_x is sampled from a Gaussian centered at the optimum value x_{opt} ¹¹ for that parameter $\mu_x \sim \mathcal{N}(x_{opt}, \sigma_x^2)$, and σ_x is a hyperparameter of our experimental setup determining the prior’s strength. For each run of BOPrO, we sample new μ_x ’s. This setup provides us with a synthetic prior that is close to the optimum,

¹¹ If the optimum for a benchmark is not known, we approximate it using the best value found during previous BO experiments.




Fig. 4: Log regret comparison of BOPrO with weak and strong priors, sampling from the strong prior, 10,000 \times random search (RS), and Spearmint (mean \pm one std on 30 repetitions). We run each benchmark for 200 iterations.

but not exactly centered at it, and allows us to control the strength of the prior by σ_x . We use two prior strengths in our experiments: a strong prior, computed with $\sigma_x = 0.01$, and a weak prior, computed with $\sigma_x = 0.1$.

Figure 4 compares BOPrO to other optimizers using the log simple regret on 30 runs on the synthetic benchmarks. We compare the results of BOPrO with weak and strong priors to 10,000 \times random search (RS, i.e., for each BO sample we draw 10,000 uniform random samples), sampling from the strong prior only, and Spearmint [41], a well-adopted BO approach using GPs and EI. In Appendix G, we also compare BOPrO to TPE, SMAC, and TuRBO [19,26,9]. Also, in Appendix H, we compare BOPrO to other baselines with the same prior initialization and show that the performance of the baselines remains similar.

BOPrO with a strong prior for the optimum beats 10,000 \times RS and BOPrO with a weak prior on all benchmarks. It also outperforms the performance of sampling from the strong prior; this is expected because the prior sampling cannot focus on the real location of the optimum. Both methods are identical in the initialization phase because they both sample from the prior in that phase.

BOPrO with a strong prior is also more sample efficient and finds better or similar results to Spearmint on all benchmarks. Importantly, in all our experiments, BOPrO with a good prior consistently shows tremendous speedups in the early phases of the optimization process, requiring on average only 15 iterations to reach the performance that Spearmint reaches after 100 iterations (6.67 \times faster). Thus, BOPrO makes use of the best of both worlds, leveraging prior knowledge and efficient optimization based on BO.

4.3 The Spatial Use-case

We next apply BOPrO to the **Spatial** [23] real-world application. **Spatial** is a programming language and corresponding compiler for the design of application accelerators, i.e., FPGAs. We apply BOPrO to three **Spatial** benchmarks, namely, 7D shallow and deep CNNs, and a 10D molecular dynamics grid application (MD Grid). We compare the performance of BOPrO to RS, manual optimization, and HyperMapper [31], the current state-of-the-art BO solution




Fig. 5: Log regret comparison of random search (RS), HyperMapper, BOPrO, and manual optimization on *Spatial*. The line and shaded regions show mean and stdev after 30 repetitions. Vertical lines denote the end of the initialization phase.

for *Spatial*. For a fair comparison between BOPrO and HyperMapper, since HyperMapper uses RFs as its surrogate model, here, we also use RFs in BOPrO. The manual optimization and the prior for BOPrO were provided by an unbiased *Spatial* developer, who is not an author of this paper. The priors were provided once and kept unchanged for the whole project. More details on the setup, including the priors used, are presented in Appendix D.

Figure 5 shows the log regret on the *Spatial* benchmarks. BOPrO vastly outperforms RS in all benchmarks. BOPrO is also able to leverage the expert’s prior and outperforms the expert in all benchmarks ($2.73\times$, $1.05\times$, and $10.41\times$ speedup for shallow CNN, deep CNN, and MD Grid, respectively). In the MD Grid benchmark, BOPrO achieves better performance than HyperMapper in the early stages of optimization ($2.98\times$ speedup after the first 10 iterations, see the plot inset), and achieves better final performance ($1.21\times$ speedup). For context, this is a significant improvement in the FPGA field, where a 10% improvement could qualify for acceptance in a top-tier conference. In the CNN benchmarks, BOPrO converges to the minima regions faster than HyperMapper ($1.18\times$ and $2\times$ faster for shallow and deep, respectively). Thus, BOPrO leverages both the expert’s prior knowledge and BO to provide a new state of the art for *Spatial*.

5 Related Work

TPE by Bergstra *et al.* [3], the default optimizer in the HyperOpt package [2], supports limited hand-designed priors in the form of normal or log-normal distributions. We make three technical contributions that make BOPrO more flexible than TPE. First, BOPrO allow more flexible priors; second, BOPrO is agnostic to the probabilistic model used, allowing the use of more sample-efficient models (e.g., we use GPs and RFs in our experiments); and third, inspired by Bayesian models, BOPrO gives more importance to the data as iterations progress. We also show that BOPrO outperforms HyperOpt’s TPE in Appendix G.

In parallel work, Li *et al.* [25] allow users to specify priors via a probability distribution. Their two-level approach samples a number of configurations by

maximizing samples from a GP posterior and then chooses the configuration with the highest prior as the next to evaluate. In contrast, BOPrO leverages the information from the prior more directly; is agnostic to the probabilistic model used, which is important for applications with many discrete variables like our real-world application, where RFs outperform GPs; and provably recovers from misspecified priors, while in their approach the prior never gets washed out.

The work of Ramachandran *et al.* [37] also supports priors in the form of probability distributions. Their work uses the probability integral transform to warp the search space, stretching regions where the prior has high probability, and shrinking others. Once again, compared to their approach, BOPrO is agnostic to the probabilistic model used and directly controls the balance between prior and model via the β hyperparameter. Additionally, BOPrO’s probabilistic model is fitted independently from the prior, which ensures it is not biased by the prior, while their approach fits the model to a warped version of the space, transformed by the prior, making it difficult to recover from misleading priors.

The optimization tools SMAC [19] and iRace [28] also support simple hand-designed priors, e.g. log-transformations. However, these are not properly reflected in the predictive models and both cannot explicitly recover from bad priors.

Oh *et al.* [33] and Siivola *et al.* [40] propose structural priors for high-dimensional problems. They assume that users always center the search space at regions they expect to be good and then develop BO approaches that favor configurations near the center. However, this is a rigid assumption about optimum locality, which does not allow users to freely specify their priors. Similarly, Shahriari *et al.* [38] focus on unbounded search spaces. The priors in their work are not about good regions of the space, but rather a regularizer that penalizes configurations based on their distance to the user-defined search space. The priors are automatically derived from the search space and not provided by users.

Our work also relates to meta-learning for BO [43], where BO is applied to many similar optimization problems in a sequence such that knowledge about the general problem structure can be exploited in future optimization problems. In contrast to meta-learning, BOPrO allows human experts to explicitly specify their priors. Furthermore, BOPrO does not depend on meta-features [11] and incorporates the human’s prior instead of information from different experiments [27].

6 Conclusions and Future Work

We have proposed a BO variant, BOPrO, that allows users to inject their expert knowledge into the optimization in the form of priors about which parts of the input space will yield the best performance. These are different than standard priors over functions which are much less intuitive for users. So far, BO failed to leverage the experience of domain experts, not only causing inefficiency but also driving users away from adopting BO because they could not exploit their years of knowledge in optimizing their black-box functions. BOPrO addresses this issue and we therefore expect it to facilitate the adoption of BO. We showed that BOPrO is $6.67\times$ more sample efficient than strong BO baselines, and $10,000\times$

faster than random search, on a common suite of benchmarks and achieves a new state-of-the-art performance on a real-world hardware design application. We also showed that BOPrO converges faster and robustly recovers from misleading priors. In future work, we will study how our approach can be used to leverage prior knowledge from meta-learning. Bringing these two worlds together will likely boost the performance of BO even further.

7 Acknowledgments

We thank Matthew Feldman for **Spatial** support. Luigi Nardi and Kunle Olukotun were supported in part by affiliate members and other supporters of the Stanford DAWN project — Ant Financial, Facebook, Google, Intel, Microsoft, NEC, SAP, Teradata, and VMware. Luigi Nardi was also partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. Artur Souza and Leonardo B. Oliveira were supported by CAPES, CNPq, and FAPEMIG. Frank Hutter acknowledges support by the European Research Council (ERC) under the European Union Horizon 2020 research and innovation programme through grant no. 716721. The computations were also enabled by resources provided by the Swedish National Infrastructure for Computing (SNIC) at LUNARC partially funded by the Swedish Research Council through grant agreement no. 2018-05973.

References

1. Balandat, M., Karrer, B., Jiang, D., Daulton, S., Letham, B., Wilson, A.G., Bakshy, E.: BoTorch: A framework for efficient Monte-Carlo Bayesian optimization. In: *Advances in Neural Information Processing Systems* (2020)
2. Bergstra, J., Yamins, D., Cox, D.D.: Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In: *International Conference on Machine Learning* (2013)
3. Bergstra, J.S., Bardenet, R., Bengio, Y., Kégl, B.: Algorithms for hyper-parameter optimization. In: *Advances in Neural Information Processing Systems* (2011)
4. Bouthillier, X., Varoquaux, G.: Survey of machine-learning experimental methods at NeurIPS2019 and ICLR2020. Research report, Inria Saclay Ile de France (Jan 2020), <https://hal.archives-ouvertes.fr/hal-02447823>
5. Calandra, R., Seyfarth, A., Peters, J., Deisenroth, M.P.: Bayesian optimization for learning gaits under uncertainty. *Annals of Mathematics and Artificial Intelligence* **76**(1-2), 5–23 (2016)
6. Chen, Y., Huang, A., Wang, Z., Antonoglou, I., Schrittwieser, J., Silver, D., de Freitas, N.: Bayesian optimization in alphago. *CoRR* **abs/1812.06855** (2018)
7. Clarke, A., McMahon, B., Menon, P., Patel, K.: Optimizing hyperparams for image datasets in fastai. <https://www.platform.ai/post/optimizing-hyperparams-for-image-datasets-in-fastai> (2020)
8. Dixon, L.C.W.: The global optimization problem: an introduction. *Toward global optimization* **2**, 1–15 (1978)

9. Eriksson, D., Pearce, M., Gardner, J.R., Turner, R., Poloczek, M.: Scalable global optimization via local bayesian optimization. In: *Advances in Neural Information Processing Systems* (2019)
10. Falkner, S., Klein, A., Hutter, F.: BOHB: robust and efficient hyperparameter optimization at scale. In: *International Conference on Machine Learning* (2018)
11. Feurer, M., Springenberg, J.T., Hutter, F.: Initializing bayesian hyperparameter optimization via meta-learning. In: *AAAI Conference on Artificial Intelligence* (2015)
12. Gardner, J.R., Kusner, M.J., Xu, Z.E., Weinberger, K.Q., Cunningham, J.P.: Bayesian optimization with inequality constraints. In: *International Conference on Machine Learning, ICML* (2014)
13. Golovin, D., Solnik, B., Moitra, S., Kochanski, G., Karro, J., Sculley, D.: Google vizier: A service for black-box optimization. In: *SIGKDD International Conference on Knowledge Discovery and Data Mining* (2017)
14. GPy: GPy: A gaussian process framework in python. <http://github.com/SheffieldML/GPy> (since 2012)
15. Hansen, N., Ostermeier, A.: Adapting arbitrary normal mutation distributions in evolution strategies: the covariance matrix adaptation. In: *Proceedings of IEEE International Conference on Evolutionary Computation* (1996)
16. Hansen, N., Akimoto, Y., Baudis, P.: CMA-ES/pycma on Github
17. Hernández-Lobato, J.M., Hoffman, M.W., Ghahramani, Z.: Predictive entropy search for efficient global optimization of black-box functions. In: *Advances in Neural Information Processing Systems* (2014)
18. Hutter, F., Xu, L., Hoos, H., Leyton-Brown, K.: Algorithm runtime prediction: Methods & evaluation. *Artificial Intelligence* **206**, 79–111 (2014)
19. Hutter, F., Hoos, H.H., Leyton-Brown, K.: Sequential model-based optimization for general algorithm configuration. In: *Learning and intelligent optimization conference* (2011)
20. Hutter, F., Kotthoff, L., Vanschoren, J. (eds.): *Automated Machine Learning: Methods, Systems, Challenges*. Springer (2018), available at <http://automl.org/book>.
21. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: *International Conference on Learning Representations* (2015)
22. Klein, A., Dai, Z., Hutter, F., Lawrence, N.D., Gonzalez, J.: Meta-surrogate benchmarking for hyperparameter optimization. In: *Advances in Neural Information Processing Systems* (2019)
23. Koeplinger, D., Feldman, M., Prabhakar, R., Zhang, Y., Hadjis, S., Fiszal, R., Zhao, T., Nardi, L., Pedram, A., Kozyrakis, C., Olukotun, K.: Spatial: A Language and Compiler for Application Accelerators. In: *SIGPLAN Conference on Programming Language Design and Implementation* (2018)
24. Kushner, H.J.: A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *Journal of Basic Engineering* **86**(1), 97–106 (1964)
25. Li, C., Gupta, S., Rana, S., Nguyen, V., Robles-Kelly, A., Venkatesh, S.: Incorporating expert prior knowledge into experimental design via posterior sampling. *arXiv preprint arXiv:2002.11256* (2020)
26. Lindauer, M., Eggensperger, K., Feurer, M., Falkner, S., Biedenkapp, A., Hutter, F.: Smac v3: Algorithm configuration in python. github.com/automl/SMAC3 (2017)
27. Lindauer, M., Hutter, F.: Warmstarting of model-based algorithm configuration. In: *AAAI Conference on Artificial Intelligence* (2018)

28. López-Ibáñez, M., Dubois-Lacoste, J., Pérez Cáceres, L., Stützle, T., Birattari, M.: The irace package: Iterated racing for automatic algorithm configuration. *Operations Research Perspectives* **3**, 43–58 (2016)
29. Mockus, J., Tiesis, V., Zilinskas, A.: The application of bayesian methods for seeking the extremum. *Towards global optimization* **2**(117-129), 2 (1978)
30. Nardi, L., Bodin, B., Saeedi, S., Vespa, E., Davison, A.J., Kelly, P.H.: Algorithmic performance-accuracy trade-off in 3d vision applications using hypermapper. In: *International Parallel and Distributed Processing Symposium Workshops* (2017)
31. Nardi, L., Koeplinger, D., Olukotun, K.: Practical design space exploration. In: *International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems* (2019)
32. Neal, R.M.: *Bayesian learning for neural networks*, vol. 118. Springer Science & Business Media (2012)
33. Oh, C., Gavves, E., Welling, M.: BOCK : Bayesian optimization with cylindrical kernels. In: *International Conference on Machine Learning* (2018)
34. Paleyes, A., Pullin, M., Mahsereci, M., Lawrence, N., González, J.: Emulation of physical processes with emukit. In: *Workshop on Machine Learning and the Physical Sciences, NeurIPS* (2019)
35. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011)
36. Perrone, V., Shen, H., Seeger, M., Archambeau, C., Jenatton, R.: Learning search spaces for bayesian optimization: Another view of hyperparameter transfer learning. In: *Advances in Neural Information Processing Systems* (2019)
37. Ramachandran, A., Gupta, S., Rana, S., Li, C., Venkatesh, S.: Incorporating expert prior in bayesian optimisation via space warping. *Knowledge-Based Systems* **195**, 105663 (2020)
38. Shahriari, B., Bouchard-Côté, A., Freitas, N.: Unbounded bayesian optimization via regularization. In: *Artificial intelligence and statistics*. pp. 1168–1176 (2016)
39. Shahriari, B., Swersky, K., Wang, Z., Adams, R.P., De Freitas, N.: Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE* **104**(1), 148–175 (2015)
40. Siivola, E., Vehtari, A., Vanhatalo, J., González, J., Andersen, M.R.: Correcting boundary over-exploration deficiencies in bayesian optimization with virtual derivative sign observations. In: *International Workshop on Machine Learning for Signal Processing* (2018)
41. Snoek, J., Larochelle, H., Adams, R.P.: Practical bayesian optimization of machine learning algorithms. In: *Advances in Neural Information Processing Systems* (2012)
42. Srinivas, N., Krause, A., Kakade, S.M., Seeger, M.W.: Gaussian process optimization in the bandit setting: No regret and experimental design. In: *International Conference on Machine Learning* (2010)
43. Vanschoren, J.: *Meta-learning*. In: *Automated Machine Learning - Methods, Systems, Challenges*, pp. 35–61. Springer (2019)
44. Wang, Q., Ming, Y., Jin, Z., Shen, Q., Liu, D., Smith, M.J., Veeramachaneni, K., Qu, H.: Atmseer: Increasing transparency and controllability in automated machine learning. In: *CHI Conference on Human Factors in Computing Systems* (2019)
45. Wu, J., Poloczek, M., Wilson, A.G., Frazier, P.I.: Bayesian optimization with gradients. In: *Advances in Neural Information Processing Systems* **30** (2017)

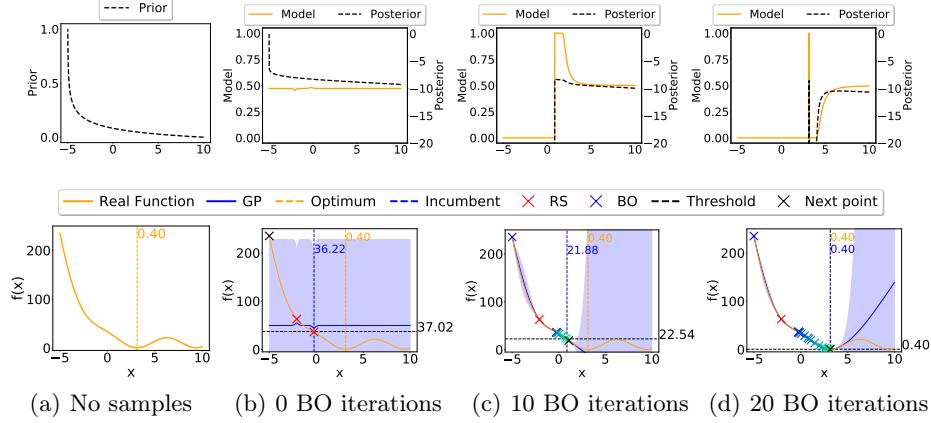


Fig. 6: BOPrO on the 1D Branin function with a decay prior. The leftmost column shows the log pseudo-posterior before any samples are evaluated, in this case, the pseudo-posterior is equal to the decay prior. The other columns show the model and pseudo-posterior after 0 (only random samples), 10, and 20 BO iterations. 2 random samples are used to initialize the GP model.

A Prior Forgetting Supplementary Experiments

In this section, we show additional evidence that BOPrO can recover from wrongly defined priors so to complement section 4.1. Figure 6 shows BOPrO on the 1D Branin function as in Figure 3 but with a decay prior. Column (a) of Figure 6 shows the decay prior and the 1D Branin function. This prior emphasizes the wrong belief that the optimum is likely located on the left side around $x = -5$ while the optimum is located at the orange dashed line. Columns (b), (c), and (d) of Figure 6 show BOPrO on the 1D Branin after $D + 1 = 2$ initial samples and 0, 10, and 20 BO iterations, respectively. In the beginning of BO, as shown in column (b), the pseudo-posterior is nearly identical to the prior and guides BOPrO towards the left region of the space. As more points are sampled, the model becomes more accurate and starts guiding the pseudo-posterior away from the wrong prior (column (c)). Notably, the pseudo-posterior before $x = 0$ falls to 0, as the predictive model is certain there will be no improvement from sampling this region. After 20 iterations, BOPrO finds the optimum region, despite the poor start (column (d)). The peak in the pseudo-posterior in column (d) shows BOPrO will continue to exploit the optimum region as it is not certain if the exact optimum has been found. The pseudo-posterior is also high in the high uncertainty region after $x = 4$, showing BOPrO will explore that region after it finds the optimum.

Figure 7 shows BOPrO on the standard 2D Branin function. We use exponential priors for both dimensions, which guides optimization towards a region with only poor performing high function values. 7a shows the prior and 7b shows

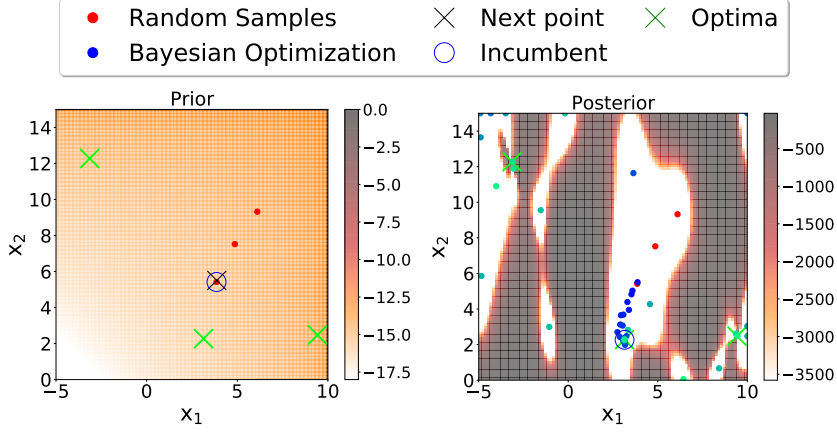


Fig. 7: BOPrO on the Branin function with exponential priors for both dimensions. (a) shows the log pseudo-posterior before any samples are evaluated, in this case, the pseudo-posterior is equal to the prior; the green crosses are the optima. (b) shows the result of optimization after 3 initialization samples drawn from the prior and 50 BO iterations. The dots in (b) show the points explored by BOPrO, with greener points denoting later iterations. The colored heatmap shows the log of the pseudo-posterior $g(\mathbf{x})$.

optimization results after $D + 1 = 3$ initialization samples and 50 BO iterations. Note that, once again, optimization begins near the region incentivized by the prior, but moves away from the prior and towards the optima as BO progresses. After 50 BO iterations, BOPrO finds all three optima regions of the Branin.

B Mathematical Derivations

B.1 EI Derivation

Here, we provide a full derivation of Eq. (7):

$$EI_{f_\gamma}(\mathbf{x}) := \int_{-\infty}^{\infty} \max(f_\gamma - y, 0) p(y|\mathbf{x}) dy = \int_{-\infty}^{f_\gamma} (f_\gamma - y) \frac{p(\mathbf{x}|y)p(y)}{p(\mathbf{x})} dy.$$

As defined in Section 3.2, $p(y < f_\gamma) = \gamma$ and γ is a quantile of the observed objective values $\{y^{(i)}\}$. Then $p(\mathbf{x}) = \int_{\mathbb{R}} p(\mathbf{x}|y)p(y) dy = \gamma g(\mathbf{x}) + (1 - \gamma)b(\mathbf{x})$, where $g(\mathbf{x})$ and $b(\mathbf{x})$ are the posteriors introduced in Section 3.3. Therefore

$$\begin{aligned} \int_{-\infty}^{f_\gamma} (f_\gamma - y) p(\mathbf{x}|y)p(y) dy &= g(\mathbf{x}) \int_{-\infty}^{f_\gamma} (f_\gamma - y) p(y) dy \\ &= \gamma f_\gamma g(\mathbf{x}) - g(\mathbf{x}) \int_{-\infty}^{f_\gamma} y p(y) dy, \end{aligned} \quad (8)$$

so that finally

$$EI_{f_\gamma}(\mathbf{x}) = \frac{\gamma f_\gamma g(\mathbf{x}) - g(\mathbf{x}) \int_{-\infty}^{f_\gamma} yp(y)dy}{\gamma g(\mathbf{x}) + (1-\gamma)b(\mathbf{x})} \propto \left(\gamma + \frac{b(\mathbf{x})}{g(\mathbf{x})}(1-\gamma) \right)^{-1}. \quad (9)$$

B.2 Proof of Proposition 1

Here, we provide the proof of Proposition 1:

$$\lim_{t \rightarrow \infty} \arg \max_{\mathbf{x} \in \mathcal{X}} EI_{f_\gamma}(\mathbf{x}) \quad (10)$$

$$= \lim_{t \rightarrow \infty} \arg \max_{\mathbf{x} \in \mathcal{X}} \int_{-\infty}^{f_\gamma} (f_\gamma - y)p(\mathbf{x}|y)p(y)dy \quad (11)$$

$$= \lim_{t \rightarrow \infty} \arg \max_{\mathbf{x} \in \mathcal{X}} g(\mathbf{x}) \int_{-\infty}^{f_\gamma} (f_\gamma - y)p(y)dy \quad (12)$$

$$= \lim_{t \rightarrow \infty} \arg \max_{\mathbf{x} \in \mathcal{X}} \left(\gamma f_\gamma g(\mathbf{x}) - g(\mathbf{x}) \int_{-\infty}^{f_\gamma} yp(y)dy \right) \quad (13)$$

$$= \lim_{t \rightarrow \infty} \arg \max_{\mathbf{x} \in \mathcal{X}} \frac{\gamma f_\gamma g(\mathbf{x}) - g(\mathbf{x}) \int_{-\infty}^{f_\gamma} yp(y)dy}{\gamma g(\mathbf{x}) + (1-\gamma)b(\mathbf{x})} \quad (14)$$

which, from Eq. (9), is equal to:

$$= \lim_{t \rightarrow \infty} \arg \max_{\mathbf{x} \in \mathcal{X}} \left(\gamma + \frac{b(\mathbf{x})}{g(\mathbf{x})}(1-\gamma) \right)^{-1} \quad (15)$$

we can take Eq. (15) to the power of $\frac{1}{t}$ without changing the expression, since the argument that maximizes EI does not change:

$$= \lim_{t \rightarrow \infty} \arg \max_{\mathbf{x} \in \mathcal{X}} \left(\gamma + \frac{b(\mathbf{x})}{g(\mathbf{x})}(1-\gamma) \right)^{-\frac{1}{t}} \quad (16)$$

substituting $g(\mathbf{x})$ and $b(\mathbf{x})$ using their definitions in Section 3.3:

$$= \lim_{t \rightarrow \infty} \arg \max_{\mathbf{x} \in \mathcal{X}} \left(\gamma + \frac{P_b(\mathbf{x})\mathcal{M}_b(\mathbf{x})^{\frac{t}{\beta}}}{P_g(\mathbf{x})\mathcal{M}_g(\mathbf{x})^{\frac{t}{\beta}}}(1-\gamma) \right)^{-\frac{1}{t}} \quad (17)$$

$$= \lim_{t \rightarrow \infty} \arg \max_{\mathbf{x} \in \mathcal{X}} \left(\frac{P_b(\mathbf{x})\mathcal{M}_b(\mathbf{x})^{\frac{t}{\beta}}}{P_g(\mathbf{x})\mathcal{M}_g(\mathbf{x})^{\frac{t}{\beta}}}(1-\gamma) \right)^{-\frac{1}{t}} \quad (18)$$

$$= \lim_{t \rightarrow \infty} \arg \max_{\mathbf{x} \in \mathcal{X}} \left(\frac{P_b(\mathbf{x})}{P_g(\mathbf{x})} \right)^{-\frac{1}{t}} \left(\frac{\mathcal{M}_b(\mathbf{x})^{\frac{t}{\beta}}}{\mathcal{M}_g(\mathbf{x})^{\frac{t}{\beta}}} \right)^{-\frac{1}{t}} (1-\gamma)^{-\frac{1}{t}} \quad (19)$$

$$= \lim_{t \rightarrow \infty} \arg \max_{\mathbf{x} \in \mathcal{X}} \left(\frac{P_b(\mathbf{x})}{P_g(\mathbf{x})} \right)^{-\frac{1}{t}} \left(\frac{\mathcal{M}_b(\mathbf{x})}{\mathcal{M}_g(\mathbf{x})} \right)^{-\frac{1}{\beta}} (1 - \gamma)^{-\frac{1}{t}} \quad (20)$$

$$= \arg \max_{\mathbf{x} \in \mathcal{X}} \left(\frac{\mathcal{M}_b(\mathbf{x})}{\mathcal{M}_g(\mathbf{x})} \right)^{-\frac{1}{\beta}} \quad (21)$$

$$= \arg \max_{\mathbf{x} \in \mathcal{X}} \left(\frac{1 - \mathcal{M}_g(\mathbf{x})}{\mathcal{M}_g(\mathbf{x})} \right)^{-\frac{1}{\beta}} \quad (22)$$

$$= \arg \max_{\mathbf{x} \in \mathcal{X}} \left(\frac{1}{\mathcal{M}_g(\mathbf{x})} - 1 \right)^{-\frac{1}{\beta}} \quad (23)$$

$$= \arg \max_{\mathbf{x} \in \mathcal{X}} (\mathcal{M}_g(\mathbf{x}))^{\frac{1}{\beta}} \quad (24)$$

$$= \arg \max_{\mathbf{x} \in \mathcal{X}} \mathcal{M}_g(\mathbf{x}) \quad (25)$$

This shows that as iterations progress, the model grows more important. If BOPrO is run long enough, the prior washes out and BOPrO only trusts the probabilistic model. Since $\mathcal{M}_g(\mathbf{x})$ is the Probability of Improvement (PI) on the probabilistic model $p(y|\mathbf{x})$ then, in the limit, maximizing the acquisition function $EI_{f_\gamma}(\mathbf{x})$ is equivalent to maximizing the PI acquisition function on the probabilistic model $p(y|\mathbf{x})$. In other words, for high values of t , BOPrO converges to standard BO with a PI acquisition function.

C Experimental Setup

We use a combination of publicly available implementations for our predictive models. For our Gaussian Process (GP) model, we use GPy's [14] GP implementation with the Matérn5/2 kernel. We use different length-scales for each input dimension, learned via Automatic Relevance Determination (ARD) [32]. For our Random Forests (RF), we use scikit-learn's RF implementation [35]. We set the fraction of features per split to 0.5, the minimum number of samples for a split to 5, and disable bagging. We also adapt our RF implementation to use the same split selection approach as Hutter *et al.* [18].

For our constrained Bayesian Optimization (cBO) approach, we use scikit-learn's RF classifier, trained on previously explored configurations, to predict the probability of a configuration being feasible. We then weight our EI acquisition function by this probability of feasibility, as proposed by Gardner *et al.* [12]. We normalize our EI acquisition function before considering the probability of feasibility, to ensure both values are in the same range. This cBO implementation is used in the Spatial use-case as in Nardi *et al.* [31].

For all experiments, we set the model weight hyperparameter to $\beta = 10$ and the model quantile to $\gamma = 0.05$, see Appendices J and I. Before starting the main BO loop, BOPrO is initialized by random sampling $D + 1$ points from the prior, where D is the number of input variables. We use the public implementation

Table 1: Search spaces for our synthetic benchmarks. For the Profet benchmarks, we report the original ranges and whether or not a log scale was used.

Benchmark	Parameter name	Parameter values	Log scale
Branin	x_1	$[-5, 10]$	-
	x_2	$[0, 15]$	-
SVM	C	$[e^{-10}, e^{10}]$	✓
	γ	$[e^{-10}, e^{10}]$	✓
FCNet	learning rate	$[10^{-6}, 10^{-1}]$	✓
	batch size	$[8, 128]$	✓
	units layer 1	$[16, 512]$	✓
	units layer 2	$[16, 512]$	✓
	dropout rate l1	$[0.0, 0.99]$	-
	dropout rate l2	$[0.0, 0.99]$	-
XGBoost	learning rate	$[10^{-6}, 10^{-1}]$	✓
	gamma	$[0, 2]$	-
	L1 regularization	$[10^{-5}, 10^3]$	✓
	L2 regularization	$[10^{-5}, 10^3]$	✓
	number of estimators	$[10, 500]$	-
	subsampling	$[0.1, 1]$	-
	maximum depth	$[1, 15]$	-
	minimum child weight	$[0, 20]$	-

of Spearmint¹², which by default uses 2 random samples for initialization. We normalize our synthetic priors before computing the pseudo-posterior, to ensure they are in the same range as our model. We also implement interleaving which randomly samples a point to explore during BO with a 10% chance.

We optimize our EI acquisition function using a combination of a multi-start local search and CMA-ES [15]. Our multi-start local search is similar to the one used in SMAC [19]. Namely, we start local searches on the 10 best points evaluated in previous BO iterations, on the 10 best performing points from a set of 10,000 random samples, on the 10 best performing points from 10,000 random samples drawn from the prior, and on the mode of the prior. To compute the neighbors of each of these 31 total points, we normalize the range of each parameter to $[0, 1]$ and randomly sample four neighbors from a truncated Gaussian centered at the original value and with standard deviation $\sigma = 0.1$. For CMA-ES, we use the public implementation of pycma [16]. We run pycma with two starting points, one at the incumbent and one at the mode of the prior. For both initializations we set $\sigma_0 = 0.2$. We only use CMA-ES for our continuous search space benchmarks.

We use four synthetic benchmarks in our experiments.

Branin. The Branin function is a well-known synthetic benchmark for optimization problems [8]. The Branin function has two input dimensions and three global minima.

¹² <https://github.com/HIPS/Spearmint>

Table 2: Search space, priors, and expert configuration for the Shallow CNN application. The default value for each parameter is shown in bold.

Parameter name	Type	Values	Expert	Prior
LP	Ordinal	[1 , 4, 8, 16, 32]	16	[0.4, 0.065, 0.07, 0.065, 0.4]
P1	Ordinal	[1 , 2, 3, 4]	1	[0.1, 0.3, 0.3, 0.3]
SP	Ordinal	[1 , 4, 8, 16, 32]	16	[0.4, 0.065, 0.07, 0.065, 0.4]
P2	Ordinal	[1 , 2, 3, 4]	4	[0.1, 0.3, 0.3, 0.3]
P3	Ordinal	[1 , 2, ..., 31, 32]	1	[0.1, 0.1, 0.033, 0.1, 0.021, 0.021, 0.021, 0.1, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021, 0.021]
P4	Ordinal	[1 , 2, ..., 47, 48]	4	[0.08, 0.0809, 0.0137, 0.1, 0.0137, 0.0137, 0.0137, 0.1, 0.0137, 0.0137, 0.0137, 0.05, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137, 0.0137]
x276	Categorical	[false, true]	true	[0.1, 0.9]

SVM. A hyperparameter-optimization benchmark in 2D based on Profet [22]. This benchmark is generated by a generative meta-model built using a set of SVM classification models trained on 16 OpenML tasks. The benchmark has two input parameters, corresponding to SVM hyperparameters.

FC-Net. A hyperparameter and architecture optimization benchmark in 6D based on Profet. The FC-Net benchmark is generated by a generative meta-model built using a set of feed-forward neural networks trained on the same 16 OpenML tasks as the SVM benchmark. The benchmark has six input parameters corresponding to network hyperparameters.

XGBoost. A hyperparameter-optimization benchmark in 8D based on Profet. The XGBoost benchmark is generated by a generative meta-model built using a set of XGBoost regression models in 11 UCI datasets. The benchmark has eight input parameters, corresponding to XGBoost hyperparameters.

Table 4: Search space, priors, and expert configuration for the MD Grid application. The default value for each parameter is shown in bold.

Parameter name	Type	Values	Expert	Prior
loop_grid0_z	Ordinal	[1 , 2, ..., 15, 16]	1	[0.2, 0.1, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05]
loop_q	Ordinal	[1 , 2, ..., 31, 32]	8	[0.08, 0.08, 0.02, 0.1, 0.02, 0.02, 0.02, 0.1, 0.02, 0.02, 0.02, 0.02, 0.02, 0.1, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02]
par_load	Ordinal	[1 , 2, 4]	1	[0.45, 0.1, 0.45]
loop_p	Ordinal	[1 , 2, ..., 31, 32]	2	[0.1, 0.1, 0.1, 0.1, 0.05, 0.03, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02]
loop_grid0_x	Ordinal	[1 , 2, ..., 15, 16]	1	[0.2, 0.1, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05]
loop_grid1_z	Ordinal	[1 , 2, ..., 15, 16]	1	[0.2, 0.2, 0.1, 0.1, 0.07, 0.03, 0.03, 0.03, 0.03, 0.03, 0.03, 0.03, 0.03, 0.03, 0.03]
loop_grid0_y	Ordinal	[1 , 2, ..., 15, 16]	1	[0.2, 0.1, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05, 0.05]
ATOM1LOOP	Categorical	[false, true]	true	[0.1, 0.9]
ATOM2LOOP	Categorical	[false, true]	true	[0.1, 0.9]
PLOOP	Categorical	[false, true]	true	[0.1, 0.9]

and estimates the resource-usage and runtime of the application for different parameter values. These parameters can then be optimized during compilation in order to achieve the design with the fastest runtime. We fully integrate BOPrO as a pass in **Spatial**'s compiler, so that **Spatial** can automatically use BOPrO for the optimization during compilation. This enables **Spatial** to seamlessly call BOPrO during the compilation of any new application to guide the search towards the best design on an application-specific basis.

In our experiments, we introduce for the first time the automatic optimization of three **Spatial** real-world applications, namely, 7D shallow and deep CNNs, and a 10D molecular dynamics grid application. Previous work by Nardi *et al.* [31] had applied automatic optimization of **Spatial** parameters on a set of benchmarks but in our work we focus on real-world applications raising the bar of state-of-the-art automated hardware design optimization. BOPrO is used to optimize the parameters to find a design that leads to the fastest runtime. The search space for these three applications is based on ordinal and categorical parameters; to handle

these discrete parameters in the best way we implement and use a Random Forests surrogate instead of a Gaussian Process one as explained in Appendix C. These parameters are application-specific and control how much of the FPGAs’ resources we want to use to parallelize each step of the application’s computation. The goal here is to find which steps are more important to parallelize in the final design, in order to achieve the fastest runtime. Some parameters also control whether we want to enable pipeline scheduling or not, which consumes resources but accelerates runtime, and others focus on memory management. We refer to Koeplinger *et al.* [23] and Nardi *et al.* [31] for more details on **Spatial**’s parameters.

The three **Spatial** benchmarks also have feasibility constraints in the search space, meaning that some parameter configurations are infeasible. A configuration is considered infeasible if the final design requires more logic resources than what the FPGA provides, i.e., it is not possible to perform FPGA synthesis because the design does not fit in the FPGA. To handle these constraints, we use our cBO implementation (Appendix C). Our goal is thus to find the design with the fastest runtime under the constraint that the design fits the FPGA resource budget.

The priors for these **Spatial** applications take the form of a list of probabilities, containing the probability of each ordinal or categorical value being good. Each benchmark also has a default configuration, which ensures all methods start with at least one feasible configuration. The priors and the default configuration for these benchmarks were provided once by an unbiased **Spatial** developer, who is not an author of this paper, and kept unchanged during the entire project. The search space, priors, and the expert configuration used in our experiments for each application are presented in Tables 2, 3, and 4.

E Multivariate Prior Comparison

In this section we compare the performance of BOPrO with univariate and multivariate priors. For this, we construct synthetic univariate and multivariate priors using Kernel Density Estimation (KDE) with a Gaussian kernel. We build strong and weak versions of the KDE priors. The strong priors are computed using a KDE on the best $10D$ out of $10,000,000D$ uniformly sampled points, while the weak priors are computed using a KDE on the best $10D$ out of $1,000D$ uniformly sampled points. We use the same points for both univariate and multivariate priors. We use scipy’s Gaussian KDE implementation, but adapt its Scott’s Rule bandwidth to $100n^{-\frac{1}{d}}$, where d is the number of variables in the KDE prior, to make our priors more peaked.

Figure 8 shows a log regret comparison of BOPrO with univariate and multivariate KDE priors. We note that in all cases BOPrO achieves similar performance with univariate and multivariate priors. For the Branin and SVM benchmarks, the weak multivariate prior leads to slightly better results than the weak univariate prior. However, we note that the difference is small, in the order of 10^{-4} and 10^{-6} , respectively.

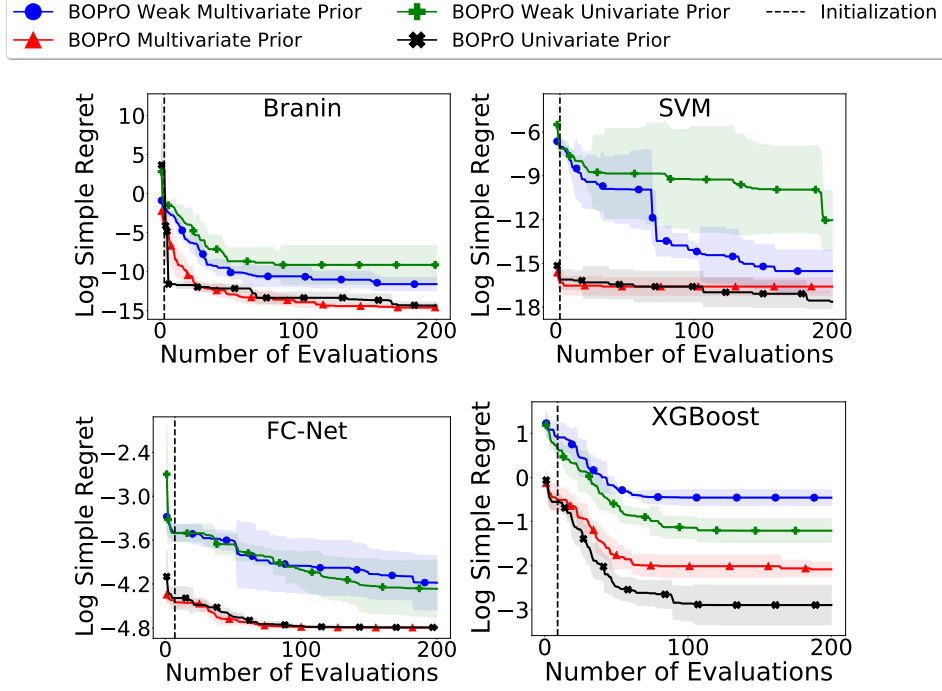


Fig. 8: Log regret comparison of BOPrO with multivariate and univariate KDE priors. The line and shaded regions show the mean and standard deviation of the log simple regret after 5 runs. All methods were initialized with $D + 1$ random samples, where D is the number of input dimensions, indicated by the vertical dashed line. We run the benchmarks for 200 iterations.

Surprisingly, for the XGBoost benchmark, the univariate version for both the weak and strong priors lead to better results than their respective multivariate counterparts, though, once again, the difference in performance is small, around 0.2 and 0.03 for the weak and strong prior, respectively, whereas the XGBoost benchmark can reach values as high as $f(\mathbf{x}) = 700$. Our hypothesis is that this difference comes from the bandwidth estimator ($100n^{-\frac{1}{d}}$), which leads to larger bandwidths, consequently, smoother priors, when a multivariate prior is constructed.

F Misleading Prior Comparison

Figure 9 shows the effect of injecting a misleading prior in BOPrO. We compare BOPrO with a misleading prior, no prior, a weak prior, and a strong prior. For our misleading prior, we use a Gaussian centered at the worst point out of 10,000,000 D uniform random samples. Namely, for each parameter, we inject a

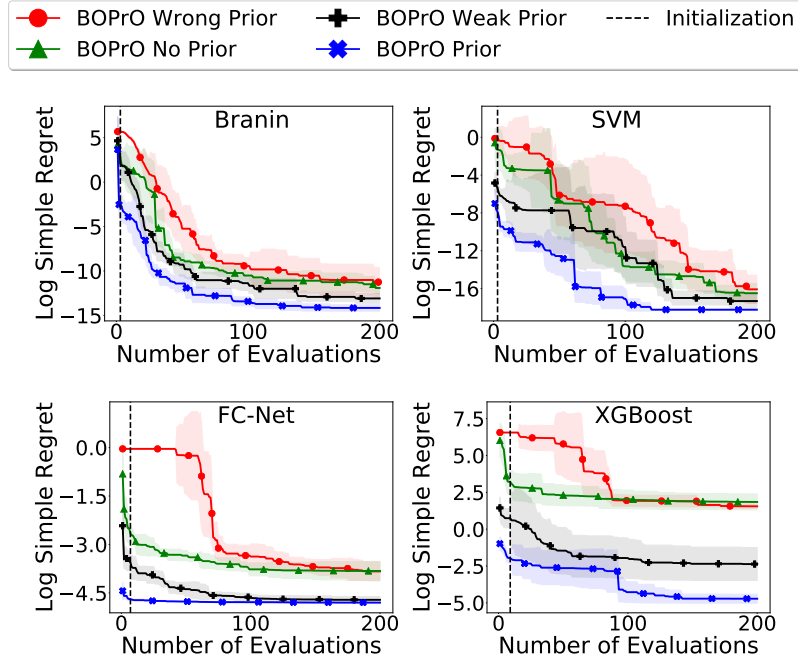


Fig. 9: Log regret comparison of BOPrO with varying prior quality. The line and shaded regions show the mean and standard deviation of the log simple regret after 5 runs. All methods were initialized with $D + 1$ random samples, where D is the number of input dimensions, indicated by the vertical dashed line. We run the benchmarks for 200 iterations.

prior of the form $\mathcal{N}(x_w, \sigma_w^2)$, where x_w is the value of the parameter at the point with highest function value out of 10,000,000 D uniform random samples and $\sigma_w = 0.01$. For all benchmarks, we note that the misleading prior slows down convergence, as expected, since it pushes the optimization away from the optima in the initial phase. However, BOPrO is still able to forget the misleading prior and achieve similar regret to BOPrO without a prior.

G Comparison to Other Baselines

We compare BOPrO to SMAC [19], TuRBO [9], and TPE [3] on our four synthetic benchmarks. We use Hyperopt’s implementation¹³ of TPE, the public implementation of TuRBO¹⁴, and the SMAC3 Python implementation of SMAC¹⁵. Hyperopt defines priors as one of a list of supported distributions, including

¹³ <https://github.com/hyperopt/hyperopt>

¹⁴ <https://github.com/uber-research/TuRBO>

¹⁵ <https://github.com/automl/SMAC3>

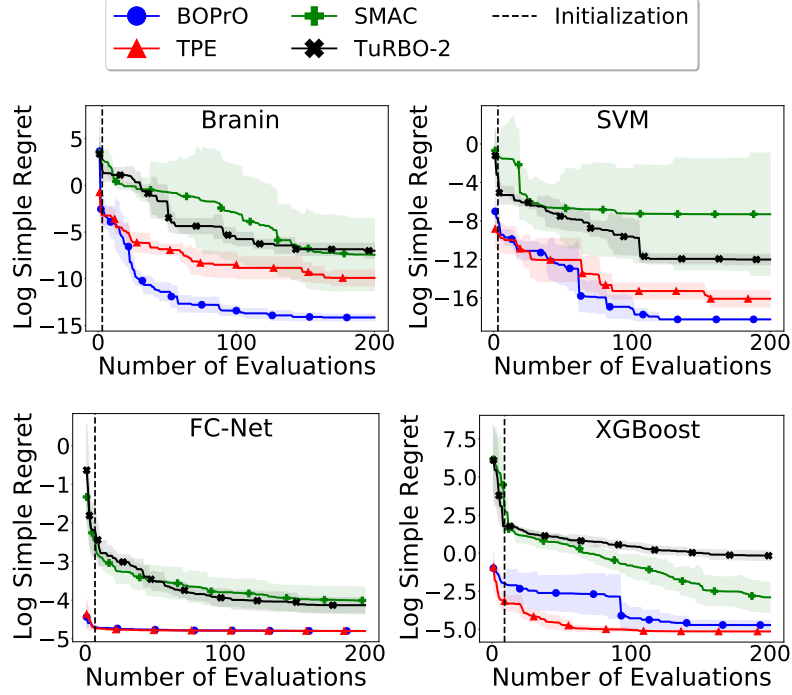


Fig. 10: Log regret comparison of BOPrO, SMAC, and TPE. The line and shaded regions show the mean and standard deviation of the log simple regret after 5 runs. BOPrO was initialized with $D + 1$ random samples, where D is the number of input dimensions, indicated by the vertical dashed line. We run the benchmarks for 200 iterations.

Uniform, Normal, and Lognormal distributions, while SMAC and TuRBO do not support priors on the locality of an optimum under the form of probability distributions.

For the three Profet benchmarks (SVM, FCNet, and XGBoost), we inject the strong priors defined in Section 4.2 into both Hyperopt and BOPrO. For Branin, we also inject the strong prior defined in Section 4.2 into BOPrO, however, we cannot inject this prior into Hyperopt. Our strong prior for Branin takes the form of a Gaussian mixture prior peaked at all three optima and Hyperopt does not support Gaussian mixture priors. Instead, for Hyperopt, we arbitrarily choose one of the optima $((\pi, 2.275))$ and use a Gaussian prior centered near that optimum. We note that since we compare all approaches based on the log simple regret, both priors are comparable in terms of prior strength, since finding one optimum or all three would lead to the same log regret. Also, we note that using Hyperopt’s Gaussian priors leads to an unbounded search space, which sometimes leads TPE to suggest parameter configurations outside the allowed parameter range. To prevent these values from being evaluated, we convert values outside

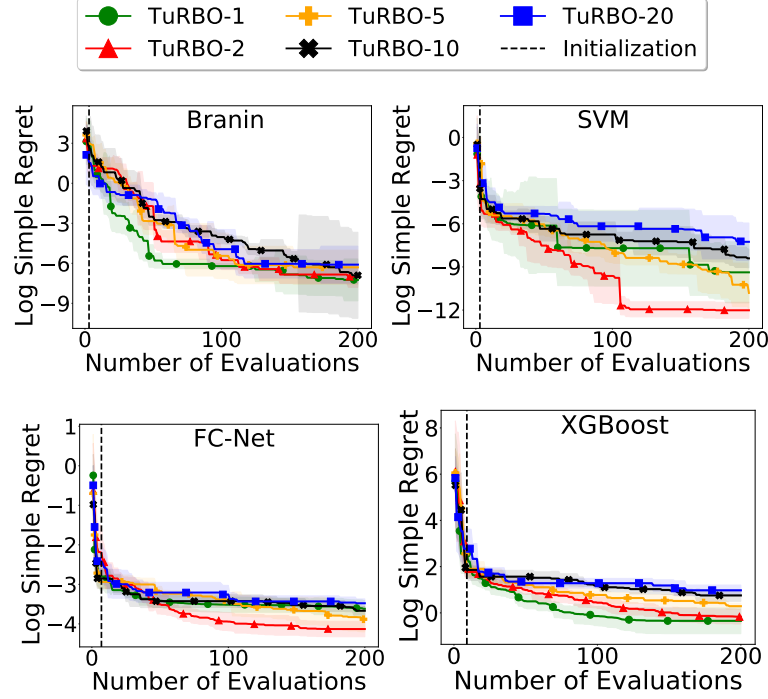


Fig. 11: Log regret comparison of TuRBO with different number of trust regions. TuRBO- M denotes TuRBO with M trust regions. The line and shaded regions show the mean and standard deviation of the log simple regret after 5 runs. TuRBO was initialized with $D + 1$ uniform random samples, where D is the number of input dimensions, indicated by the vertical dashed line. We run the benchmarks for 200 iterations.

the parameter range to be equal to the upper or lower range limit, depending on which limit was exceeded. We do not inject any priors into SMAC and TuRBO, since these methods do not support priors about the locality of an optimum.

Figure 10 shows a log regret comparison between BOPrO, SMAC, TuRBO-2 and TPE on our four synthetic benchmarks. We use TuRBO-2 since it led to better performance overall compared to other variants, see Fig. 11. BOPrO achieves better performance than SMAC and TuRBO on all four benchmarks. Compared to TPE, BOPrO achieves similar or better performance on three of the four synthetic benchmarks, namely Branin, SVM, and FCNet, and slightly worse performance on XGBoost. We note, however, that the good performance of TPE on XGBoost may be an artifact of the approach of clipping values to its maximal or minimal values as mentioned above. In fact, the clipping nudges TPE towards promising configurations in this case, since XGBoost has low function values near the edges of the search space. Overall, the better performance of

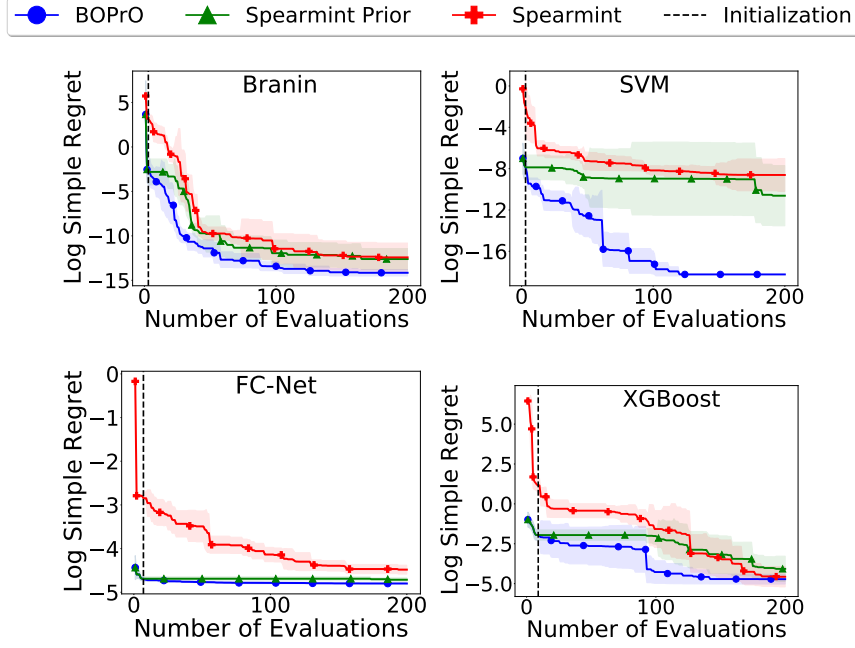


Fig. 12: Log regret comparison of BOPrO, Spearmint with prior initialization, and Spearmint with default initialization. The line and shaded regions show the mean and standard deviation of the log simple regret after 5 runs. BOPrO and Spearmint Prior were initialized with $D + 1$ random samples from the prior, where D is the number of input dimensions, indicated by the vertical dashed line. We run the benchmarks for 200 iterations.

BOPrO is expected, since BOPrO is able to combine prior knowledge with more sample-efficient surrogates, which leads to better performance.

H Prior Baselines Comparison

We show that simply initializing a BO method in the DoE phase by sampling from a prior on the locality of an optimum doesn't necessarily lead to better performance. Instead in BOPrO, it is the pseudo-posterior in Eq. (4) that drives its stronger performance by combining prior and new observations. To show that, we compare BOPrO with Spearmint and HyperMapper such as in section Sections 4.2 and 4.3, respectively, but we initialize all three methods using the same approach. Namely, we initialize all methods with $D + 1$ samples from the prior. Our goal is to show that simply initializing Spearmint and HyperMapper with the prior will not lead to the same performance as BOPrO, because, unlike BOPrO, these baselines do not leverage the prior after the DoE initialization phase. We report results on both our synthetic and real-world benchmarks




Fig. 13: Log regret comparison of BOPrO, HyperMapper with prior initialization, and HyperMapper with default initialization. The line and shaded regions show the mean and standard deviation of the log simple regret after 5 runs. BOPrO and HyperMapper Prior were initialized with $D + 1$ random samples from the prior, where D is the number of input dimensions, indicated by the vertical dashed line.

Figure 12 shows the comparison between BOPrO and Spearmint Prior. In most benchmarks, the prior initialization leads to similar final performance. In particular, for XGBoost, the prior leads to improvement in early iterations, but to worse final performance. We note that for FCNet, Spearmint Prior achieves better performance, however, we note that the improved performance is given almost solely from sampling from the prior. There is no improvement for Spearmint Prior until around iteration 190. In contrast, in all cases, BOPrO is able to leverage the prior both during initialization and its Bayesian Optimization phase, leading to improved performance. BOPrO still achieves similar or better performance than Spearmint Prior in all benchmarks.

Figure 13 shows similar results for our Spatial benchmarks. The prior does not lead HyperMapper to improved final performance. For the Shallow CNN benchmark, the prior leads HyperMapper to improved performance in early iterations, compared to HyperMapper with default initialization, but HyperMapper Prior is still outperformed by BOPrO. Additionally, the prior leads to degraded performance in the Deep CNN benchmark. These results confirm that BOPrO is able to leverage the prior in its pseudo-posterior during optimization, leading to improved performance in almost all benchmarks compared to state-of-the-art BO baselines.

I γ -Sensitivity Study

We show the effect of the γ hyperparameter introduced in Section 3.2 for the quantile identifying the points considered to be good. To show this, we compare the performance of BOPrO with our strong prior and different γ values. For all experiments, we initialize BOPrO with $D + 1$ random samples and then run BOPrO until it reaches $10D$ function evaluations. For each γ value, we run BOPrO five times and report mean and standard deviation.




Fig. 14: Comparison of BOPrO with the strong prior and different values for the γ hyperparameter on our four synthetic benchmarks. We run BOPrO with a budget of $10D$ function evaluations, including $D + 1$ randomly sampled DoE configurations.

Figure 14 shows the results of our comparison. We first note that values near the lower and higher extremes lead to degraded performance, this is expected, since these values will lead to an excess of either exploitation or exploration. Further, we note that BOPrO achieves similar performance for all values of γ , however, values around $\gamma = 0.05$ consistently lead to better performance.

J β -Sensitivity Study

We show the effect of the β hyperparameter introduced in Section 3.3 for controlling the influence of the prior over time. To show the effects of β , we compare the performance of BOPrO with our strong prior and different β values on our four synthetic benchmarks. For all experiments, we initialize BOPrO with $D + 1$ random samples and then run BOPrO until it reaches $10D$ function evaluations. For each β value, we run BOPrO five times and report mean and standard deviation.

Figure 15 shows the results of our comparison. We note that values of β that are too low (near 0.01) or too high (near 1000) often lead to lower performance. This shows that putting too much emphasis on the model or the prior will lead to degraded performance, as expected. Further, we note that $\beta = 10$ lead to the

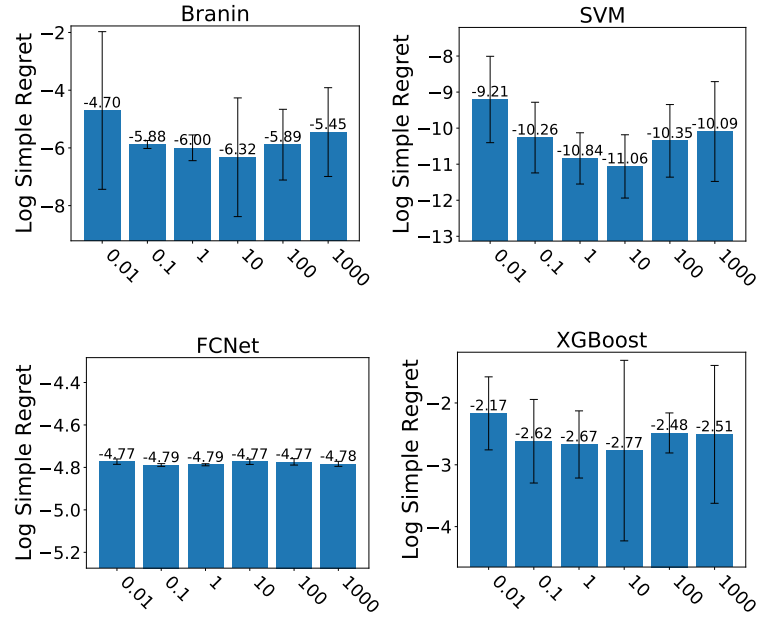


Fig. 15: Comparison of BOPrO with the strong prior and different values for the β hyperparameter on our four synthetic benchmarks. We run BOPrO with a budget of $10D$ function evaluations, including $D + 1$ randomly sampled DoE configurations.

best performance in three out of our four benchmarks. This result is reasonable, as $\beta = 10$ means BOPrO will put more emphasis on the prior in early iterations, when the model is still not accurate, and slowly shift towards putting more emphasis on the model as the model sees more data and becomes more accurate.