

# PG-LBO: Enhancing High-Dimensional Bayesian Optimization with Pseudo-Label and Gaussian Process Guidance

Taikai Chen<sup>1</sup>, Yue Duan<sup>1</sup>, Dong Li<sup>2</sup>, Lei Qi<sup>3</sup>, Yinghuan Shi<sup>1\*</sup>, Yang Gao<sup>1</sup>

<sup>1</sup>National Key Laboratory for Novel Software Technology, Nanjing University, China

<sup>2</sup>Huawei Noah's Ark Lab, China

<sup>3</sup>School of Computer Science and Engineering, Southeast University, China

{taicaichen, yueduan}@smail.nju.edu.cn, lidong106@huawei.com, qilei@seu.edu.cn, {syh, gaoy}@nju.edu.cn

## Abstract

Variational Autoencoder based Bayesian Optimization (VAE-BO) has demonstrated its excellent performance in addressing high-dimensional structured optimization problems. However, current mainstream methods overlook the potential of utilizing a pool of unlabeled data to construct the latent space, while only concentrating on designing sophisticated models to leverage the labeled data. Despite their effective usage of labeled data, these methods often require extra network structures, additional procedure, resulting in computational inefficiency. To address this issue, we propose a novel method to effectively utilize unlabeled data with the guidance of labeled data. Specifically, we tailor the pseudo-labeling technique from semi-supervised learning to explicitly reveal the relative magnitudes of optimization objective values hidden within the unlabeled data. Based on this technique, we assign appropriate training weights to unlabeled data to enhance the construction of a discriminative latent space. Furthermore, we treat the VAE encoder and the Gaussian Process (GP) in Bayesian optimization as a unified deep kernel learning process, allowing the direct utilization of labeled data, which we term as Gaussian Process guidance. This directly and effectively integrates the goal of improving GP accuracy into the VAE training, thereby guiding the construction of the latent space. The extensive experiments demonstrate that our proposed method outperforms existing VAE-BO algorithms in various optimization scenarios. Our code will be published at <https://github.com/TaikaiChen/PG-LBO>.

## Introduction

Bayesian optimization has found widespread applications in various optimization problems and machine learning tasks, enabling the discovery of optimal solutions for complex functions (Snoek, Larochelle, and Adams 2012). Traditional BO methods based on the Gaussian process have demonstrated their remarkable performance in practical applications. However, with the increase in problem dimensions, these conventional methods face significant challenges. The computational cost of searching in high-dimensional spaces grows exponentially, leading to a drastic decrease in the efficiency of traditional BO algorithms (Kandasamy, Schneider, and Póczos 2015). Recently, Variational Autoencoder (VAE)

provides a powerful method to map high-dimensional structured data into low-dimensional continuous latent spaces, enabling efficient and effective optimization of the objective function. Consequently, VAE-BO possesses unique advantages in tackling high-dimensional structured optimization problems. Currently, the VAE-BO based methods have been widely applied in various fields, including molecular synthesis (Stanton et al. 2022; Maus et al. 2023), neural architecture search (Ru et al. 2021; Biswas et al. 2023) and anomaly detection (Zhang, Wei, and Wang 2023). The research focus of VAE-BO lies in effectively utilizing both labeled and unlabeled data to learn a discriminative latent space. Nevertheless, some existing VAE-BO methods still exhibit limitations in the utilization of labeled data and unlabeled data.

Firstly, while existing VAE-BO algorithms make use of labeled data, unfortunately, their utilization of unlabeled data is still straightforward, whereas the efficacy of unlabeled data has not been fully exploited. For example, the pioneering work of (Gómez-Bombarelli et al. 2018) was the first to introduce VAE into the BO domain. However, in their method, the VAE was trained in an unsupervised manner, resulting in a non-discriminative latent space that was ill-suited to BO processes. Subsequent VAE-BO research has largely focused on better utilizing labeled data with black-box function values to learn a discriminative latent space. For instance, (Eissman et al. 2018) introduced additional network structures that allow for the joint utilization of labeled and unlabeled data in the pre-training of the VAE. While LBO (Tripp, Daxberger, and Hernández-Lobato 2020) employed unlabeled data for the initial pre-training of the VAE and subsequently leveraged labeled data through data weighting and periodic retraining of the VAE. In addition, T-LBO (Grosnit et al. 2021) extended LBO by incorporating deep metric learning to utilize labeled data, and the training of the VAE was divided into two stages: pre-training on unlabeled data and fine-tuning on labeled data. These works proposed various ways to leverage labeled data, while the utilization of unlabeled data has only been explored in the pre-training stage of the VAE, and its potential has not been fully exploited. We notice that unlabeled data also contain valuable discriminative information regarding optimization target values. Exploiting this information can help train more discriminative models and learn a better latent space.

Furthermore, we have observed that existing methods face

\*Corresponding author.

challenges of complexity or inefficiency when it comes to utilizing labeled data. (Eissman et al. 2018) required to introduce additional network structures which predicted black-box function values through the encoded latent representations, making the method somewhat complex. And T-LBO introduced significant computational overhead due to deep metric learning, and the overall method was computationally demanding. LBO utilized the relative magnitudes of labeled data labels through data weighting, rather than directly using the actual values, potentially not fully harnessing the potential of labeled data. Hence, how to effectively utilize labeled data in a simple yet efficient manner is worth investigating.

Be aware of these two issues, we wonder if we have an effective and efficient way to fully capitalize the potential of both labeled and unlabeled data. To this end, we present a novel semi-supervised learning method that utilizes pseudo-labels to explicitly reveal the relative magnitudes of optimization target scores in unlabeled data. Subsequently, we assign appropriate training weights to unlabeled data based on these pseudo-labels. By assigning more weights to the points with higher optimization target scores, this guides VAE to focus more on high-scoring points. This contributes to the construction of a more discriminative latent space. Furthermore, we propose a more streamlined method for exploiting labeled data. Specifically, we unify the VAE encoder and the GP in BO as a cohesive deep kernel learning process, effectively integrating the goal of improving GP accuracy directly into the VAE training. This integration allows the GP to guide the training of the VAE and consequently influence the construction of the latent space. The main contributions of our work can be summarized as follows:

- We assess the limitations of current VAE-BO approaches in terms of data utilization and delve into the untapped potential of unlabeled data in enhancing the creation of a distinct latent space.
- For the comprehensive usage of unlabeled data, we propose to utilize pseudo-labels to intergrate the implicit discriminative information within unlabeled data, thereby facilitating the construction of the latent space.
- To efficiently exploit the labeled data, we introduce a straightforward method through GP guidance by integrating the VAE encoder and GP to leverage labeled data, modeled as a classical deep kernel learning paradigm.

## Related Work

Bayesian optimization suffered from the curse of dimensionality, which referred to the exponential increase in computational complexity as the dimension of the problem increased. Although various methods had been proposed to extend Bayesian optimization to high-dimensional spaces, the focus in this paper was on the VAE-based method, also known as Latent Space Optimization (LSO) (Tripp, Daxberger, and Hernández-Lobato 2020). The VAE was used to map the high-dimensional structured space  $\mathcal{X}$  to a lower-dimensional continuous latent space  $\mathcal{Z}$ . In the latent space  $\mathcal{Z}$ , Bayesian optimization was performed to find the optimal solution  $z^* \in \mathcal{Z}$ , which was subsequently mapped back to  $\mathcal{X}$  to obtain the final solution  $x^*$ . A key challenge of VAE-BO lays

in effectively leveraging labeled and unlabeled data.

The VAE-BO was first applied to chemical design (Gómez-Bombarelli et al. 2018) and found application in various domains including automatic machine learning (Zhang et al. 2019; Biswas et al. 2023), chemical molecule synthesis (Korovina et al. 2020; Maus et al. 2023) and anomaly detection (Zhang, Wei, and Wang 2023). In those works, the VAE models were mostly trained on unlabeled data, and their parameters remained fixed during the BO process. This conventional paradigm led to a non-discriminative potential space ill-suited for BO and vulnerable to the “dead zone” issue. To address these issues, LBO suggested integrating LSO with weighted retraining of the VAE model, leveraging the black-box function values of labeled data. Assigning weights to data points according to the magnitude of the black-box function scores compelled the generative model to prioritize modeling the feasible region with high-scoring points. Furthermore, periodic VAE model retraining during Bayesian optimization iterations facilitated the incorporation of new point information. For guiding the creation of a distinctive VAE-BO latent space, T-LBO extended LBO by incorporating deep metric learning to better exploit labeled data. In addition, T-LBO adopted a pre-training and fine-tuning framework, using unlabeled data for pre-training the VAE model and then supervised fine-tuning with labeled data. This training framework effectively combined labeled and unlabeled data. Recent work, TSBO (Yin, Wang, and Li 2023), drew inspiration from Meta Pseudo Labels (Pham et al. 2021) in semi-supervised learning, introduced an additional teacher-student model structure to assign pseudo-label to the unlabeled data, which were then utilized for training the surrogate model and the VAE.

Certain VAE-BO research redirected its attention to areas beyond data utilization. For example, LOL-BO (Maus et al. 2022) concentrated on optimizing strategies within VAE-BO, introducing trust region optimization strategies that aided in searching for optimal points during the BO process. On the other hand, (Verma, Chakraborty, and Griffiths 2022) utilized Gaussian Process marginal likelihood objectives to learn invariant enhancements for determining BO query points. (Guo et al. 2022) involved using multiple VAEs to mitigate the challenges posed by confirming VAE latent space dimensions.

Our method emphasized the utilization of data in VAE-BO, specifically focusing on the utilization of unlabeled data. Compared with previous methods, we used pseudo-labels to exploit the implicit discriminative information in unlabeled data to help construct the VAE latent space, instead of simply using unlabeled data for VAE pre-training. TSBO also employed pseudo-labeling, but our method differed in pseudo-label’s application. TSBO directly utilized pseudo-labels to characterize black-box function values. While considering the potential cognitive errors introduced by pseudo-labels, we used pseudo-labels to reveal the magnitude relationship information between the black-box function values corresponding to unlabeled data and utilized this information using data weighting. Meanwhile, TSBO focused on using unlabeled data to help train the surrogate models, while we focused on using unlabeled data to aid in

constructing the latent space of VAE.

## Methodology

### Preliminaries: Bayesian Optimization

Bayesian optimization addresses the global optimization problem of finding

$$x^* = \operatorname{argmax}_{x \in \mathcal{X}} f(x), \quad (1)$$

where  $f(\cdot) : \mathcal{X} \rightarrow \mathbb{R}$  is an expensive black-box function defined on a high-dimensional structured input space  $\mathcal{X}$ . BO leverages two core components: a surrogate model to model the objective function and an acquisition function to determine the next query point for evaluation. Gaussian Process Regression (GP) is commonly used as the surrogate model, which models the posterior distribution  $f_{\Phi}(x; D_L)$  based on the available data set  $D_L = \{(x_i, y_i) | i = 1, \dots, N\}$ , where  $\Phi$  represents the hyperparameters of the GP and  $y_i = f(x_i)$ .

The posterior distribution is given as follows:

$$f_{\Phi}(x; D_L) \sim \mathcal{N}(\mu(x), \sigma^2(x)), \quad (2)$$

where

$$\mu(x) = m(x) + k(x, X)^{\top} (K + \sigma_0^2 I_N)^{-1} (y - m(x)), \quad (3)$$

$$\sigma^2(x) = k(x, x) - k(x, X)^{\top} (K + \sigma_0^2 I_N)^{-1} k(x, X), \quad (4)$$

where  $m(\cdot)$  and  $k(\cdot, \cdot)$  are the prior mean and kernel function,  $I_N$  denotes the  $N$ -dimensional identity matrix,  $k(x, X) = [k(x, x_1), k(x, x_2), \dots, k(x, x_N)]$ , and the elements of matrix  $K$  are defined as  $K_{ij} = k(x_i, x_j)$ . GP is a non-parametric model, and its hyperparameters  $\Phi$  mainly include the prior mean function, kernel function scale, and noise variance, which are typically optimized by minimizing the negative log marginal likelihood on the available data set.

The acquisition function is designed to balance the posterior mean and variance predicted by the GP model to determine the next query point. Popular acquisition functions include Expected Improvement (EI), Probability Improvement (PI), and Upper Confidence Bound (UCB), among others.

### Overview of the PG-LBO

**Weighted retraining module.** Weighted retraining, proposed by LBO, employs data weighting to enhance the modeling of high-scoring points in the VAE training objective, allocating a larger proportion of the feasible region and effectively utilizing all known data points in order to learn informative representations and avoid overfitting. Let  $q_{\phi}^{enc}$  be the pre-trained VAE encoder,  $g_{\theta}^{dec}$  be the decoder. The initial available labeled data set is denoted as  $D_L = \{(x_i, f(x_i)) | i = 1, 2, \dots, N\}$ .

The widely used weighting method is based on a rank-based weighting function:

$$w(x; D_L, k) \propto \frac{1}{kN + \operatorname{rank}_{f, D_L}(x)}, \quad (5)$$

$$\operatorname{rank}_{f, D_L}(x) = |\{x_i : f(x_i) > f(x), x_i \in D_L\}|. \quad (6)$$

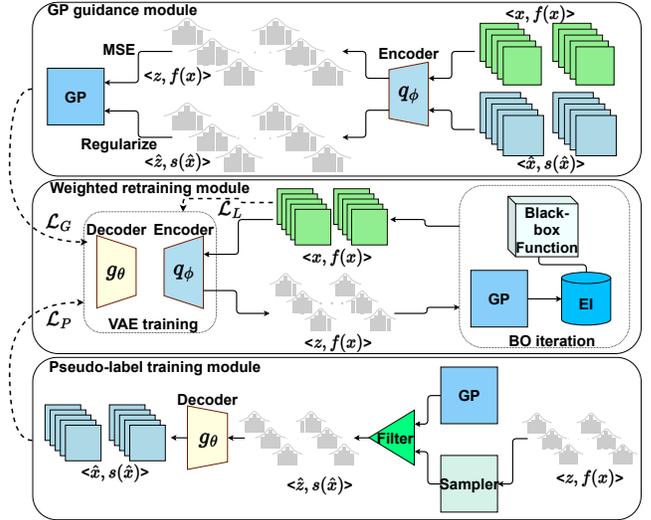


Figure 1: A overview of the PG-LBO.

This function assigns weights approximately proportional to the inverse rank (starting from zero) of each data point. A tunable hyperparameter  $k$  controls the degree of weighting, where  $k = \infty$  corresponds to uniform weighting ( $w_i = \frac{1}{N}$  for all  $i$ ), and  $k = 0$  assigns all weight to the single point with the highest objective function value. The VAE is trained on the weighted data set, and training loss can be seen as the weighted average over the data points:

$$\mathcal{L}_L(\theta, \phi) = w(x_i) [E_{q_{\phi}(z_i | x_i)} [\log g_{\theta}(x_i | z_i)] - \text{KL}(q_{\phi}(z_i | x_i) || p(z))], \quad (7)$$

where  $w(x_i)$  represents the data weight calculated according to the weighting function.

In addition to data weighting, periodic retraining of the VAE is also necessary to propagate the newly generated points by the BO process. Specifically, assuming that retraining is performed every  $r$  BO iterations, before each retraining, the labeled data set is updated as follows:  $D_L^{(k)} = D_L^{(k-1)} \cup \{(x_i^*, f(x_i^*))\}_{i=0, \dots, q}$ , and where  $D_L^{(0)} = D_L$ . New points are added to the labeled data set and then the updated data set is used to retrain the VAE.

**Pseudo-label training module.** Existing VAE-BO methods primarily focus on utilizing the black-box function labels of labeled data to guide the construction of the latent space, while overlooking the potential role of implicit discriminative information regarding optimization target values in unlabeled data. In response, we introduce pseudo-label training, a novel approach leveraging pseudo-labels to uncover latent discriminative information within unlabeled data. This knowledge is effectively employed via weighted retraining to guide latent space construction. The training objective of VAE-BO can be seen as training a Gaussian process latent variable model (Siivola et al. 2021), resembling a deep regression task. However, unlike extensively studied classification tasks in semi-supervised learning that have explored settings such as mismatch (Duan et al. 2023),

open-set (Li et al. 2023; Yang et al. 2022) and barely-supervised learning (Gui et al. 2022), the exploration of semi-supervised methodologies for deep regression tasks is relatively limited. Regression tasks directly use real-valued targets as pseudo-labels, being sensitive to prediction quality. However, Weighted retraining mitigates this sensitivity by indirectly incorporating relative objective function ranking through data weighting, avoiding direct utilization of actual values. This method reduces training sensitivity to black-box function values, enabling the effective use of unlabeled data via pseudo-labels.

Specifically, we sample unlabeled data points in the latent space  $\mathcal{Z}$  to obtain a set of points  $\{\hat{z}_j\}_{j=0,\dots,N_P}$ , then assign pseudo-labels to these unlabeled data points to reflect the relative magnitudes of the black-box function values, resulting in a pseudo-labeled data set  $D_P = \{\hat{x}_j, s(\hat{x}_j) | j = 0, \dots, N_P\}$ , where  $\hat{x}_j = g_\theta(\hat{z}_j)$ . We then use the pseudo-labels to weigh the data during training. By introducing an additional pseudo-label loss  $\mathcal{L}_P$  during training, it becomes integrated into the weighted retraining process of VAE:

$$\mathcal{L}_P(\theta, \phi) = \hat{w}(\hat{x}_j) [E_{q_\phi(\hat{z}_j|\hat{x}_j)}[\log g_\theta(\hat{x}_j|\hat{z}_j)] - \text{KL}(q_\phi(\hat{z}_j|\hat{x}_j)||p(z))], \quad (8)$$

$$\hat{w}(\hat{x}; D_P, k) \propto \frac{1}{kN_P + \text{rank}_{s, D_P}(\hat{x})}, \quad (9)$$

$$\text{rank}_{s, D_P}(\hat{x}) = |\{\hat{x}_j : s(\hat{x}_j) > s(\hat{x}), \hat{x}_j \in D_P\}|, \quad (10)$$

where  $s(\hat{x})$  represents the pseudo-label of the unlabeled data sample  $\hat{x}$ , which only needs to reflect the relative magnitudes of the black-box function values among unlabeled data, rather than accurately reflecting the actual values. For simplicity, we directly use the posterior mean provided by the GP model as the pseudo-label, *i.e.*,  $s(\hat{x}_j) = \mu_{f_\Phi}(\hat{z}_j)$ , where  $\hat{x}_j = g_\theta(\hat{z}_j)$ .

**GP guidance module.** The core purpose of utilizing labeled data is to construct a more discriminative latent space tailored for the Bayesian optimization process. Drawing inspiration from the principles of deep kernel learning, we present a more direct approach called GP guidance to harness the black-box function values of labeled data. We treat the VAE encoder and the GP model as integral components of a unified deep kernel learning process. This technique introduces a loss term that aims to minimize the mean squared error between the GP model’s predictions and the actual black-box function values of labeled data. This loss term is seamlessly integrated into the VAE’s training objective, driving the VAE to generate latent representations that closely align with the GP model’s expectations. Notably, the GP model offers a predictive posterior distribution that gauges the uncertainty of predictions. To enhance stability, we regularize the posterior by incorporating an unsupervised loss term designed to minimize the prediction variance for unlabeled data points (Jean, Xie, and Ermon 2018). The ultimate formulation of this loss term can be expressed as follows:

$$\mathcal{L}_G(\phi) = w(x_i) E_{q_\phi(z_i|x_i)} [(f(x_i) - \mu_{f_\Phi}(z_i))^2] + \hat{w}(\hat{x}_j) E_{q_\phi(\hat{z}_j|\hat{x}_j)} [\sigma_{f_\Phi}^2(\hat{z}_j)]. \quad (11)$$

Since the GP model is a non-parametric model and lacks learnable parameters like neural networks, it primarily relies on the hyperparameters  $\Phi$  to fit the training data. Therefore, during VAE training, we keep the GP model fixed and only use it for guidance. After VAE training, we re-encode the labeled data using VAE to obtain new data latent representations, then fit the GP model on these new representations, and update the hyperparameters of the GP model. This way, the VAE and the GP model are continually optimized through this approximate alternating update process.

**Data sampling and pseudo-label selection.** Unlabeled data is acquired through sampling. Weighted retraining encourages the VAE to use larger feasible regions for modeling high-scoring points. However, if the sampled unlabeled data mainly comprises low-scoring points, it might not contribute effectively to the model and could even worsen its performance. Hence, the data sampling strategy must be carefully designed to prioritize higher-scoring points. In this study, we utilize two straightforward sampling approaches: noisy sampling and heuristic sampling. Noisy sampling introduces Gaussian noise to existing labeled data, effectively sampling around high-scoring points. Heuristic sampling treats sampling as an optimization task, utilizing a simple heuristic optimization algorithm to iteratively maximize data point pseudo-labels. These approaches enhance the overall quality of sampled points. A detailed introduction of the sampling methods are in Appendix A.

While data weighting effectively mitigates potential impacts from discrepancies between pseudo-labels and true labels, pseudo-labels carrying significant errors can still introduce cognitive bias into the model. Therefore, we filter the sampled data during the sampling process to improve the accuracy of the pseudo-label data as much as possible. BO can be viewed as an optimization for regression tasks. Different from classification tasks where straightforward thresholding can sift confident class predictions, regression tasks employ real-valued targets as pseudo-labels, rendering the sole reliance on model outputs for confidence assessment challenging. Inspired by the work of (Rizve et al. 2021), we use uncertainty to select robust pseudo-label data. This choice presents distinct benefits when coupled with GP models, as GP models can jointly express both mean and variance. The variance conveniently doubles as a measure of prediction uncertainty, negating the need for extra computation. We provide a more detailed discussion on the threshold selection in Appendix B. Drawing inspiration from the work of FreeMatch (Wang et al. 2023) in semi-supervised learning, we set the uncertainty threshold  $\tau_t$  for filtering as the model’s average prediction variance on unlabeled data and estimate the uncertainty threshold as the exponential moving average (EMA) of threshold values at each training step:

$$\tau_t = \begin{cases} \frac{1}{N'_P} \sum_{j=1}^{N'_P} \sigma_{f_\Phi}^2(z'_j), & \text{if } t=0, \\ \lambda\tau_{t-1} + (1-\lambda)\frac{1}{N_P} \sum_{j=1}^{N_P} \sigma_{f_\Phi}^2(z_j), & \text{otherwise,} \end{cases} \quad (12)$$

where  $\lambda \in (0, 1)$  is the momentum decay for EMA. When  $t = 0$ , we pre-sample a subset of points  $\{z'_j\}_{j=1,\dots,N'_P}$ , where  $N'_P = N_P/10$ , to determine the initial threshold.

**Algorithm 1:** Pseudo code of PG-LBO

---

**Input:** Data  $D_L$ , query budget  $M$ , object function  $f(x)$ , VAE encoder/decoder  $q_\phi(x)/g_\theta(z)$ , retrain frequency  $r$ , weight function  $w(x)$

- 1 **for**  $1, \dots, M/r$  **do**
- 2      $\mathcal{W} \leftarrow \{w(x_i)\}_{x_i \in D_L}$  ;
- 3      $\hat{\mathcal{W}} \leftarrow \{w(\hat{x}_j)\}_{\hat{x}_j \in D_P}$  ;
- 4     Weight  $D_L$  by  $\mathcal{W}$ , weight  $D_P$  by  $\hat{\mathcal{W}}$  ;
- 5     Solve  $\phi^*, \theta^* \leftarrow \operatorname{argmin}_{\phi, \theta} \mathcal{L}(\phi, \theta)$  on  $D_L, D_P$  ;
- 6      $\{z_i \leftarrow q_\phi(x_i)\}_{x_i \in D_L}$  ;
- 7      $D_Z \leftarrow \{(z_i, f(x_i))\}_{x_i \in D_L}$  ;
- 8     **for**  $1, \dots, r$  **do**
- 9         Fit GP on  $D_Z$  to obtain  $f_\Phi(z; D_Z)$  ;
- 10         Optimize EI to obtain new latent point  $\tilde{z}$  ;
- 11          $\tilde{x} \leftarrow g_\theta(\tilde{z})$  ;
- 12          $D_L \leftarrow D_L \cup \{\tilde{x}, f(\tilde{x})\}$  ;
- 13          $D_Z \leftarrow D_Z \cup \{\tilde{z}, f(\tilde{x})\}$  ;
- 14     **end**
- 15     Sample latent data point  $\hat{z}$  ;
- 16     Filter  $\hat{z}$  by thresholds  $\tau$  ;
- 17      $\hat{x} \leftarrow g_\theta(\hat{z})$  ;
- 18      $D_P \leftarrow \{\hat{x}, \mu_{f_\Phi}(\hat{z})\}$  ;
- 19     Update  $\tau$  ;
- 20 **end**
- 21  $x^* \leftarrow \operatorname{argmax}_{x \in D_L} f(x)$  ;
- 22 **return**  $x^*$

---

**Training loss.** By combining the aforementioned components, we have developed a novel high-dimensional Bayesian optimization algorithm, named PG-LBO. Figure 1 provides an overview of the PG-LBO. The training loss of the algorithm can be represented as follows:

$$\mathcal{L} = \mathcal{L}_L + \lambda_P \mathcal{L}_P + \lambda_G \mathcal{L}_G, \quad (13)$$

where  $\lambda_P$  and  $\lambda_G$  represent the weights for the Pseudo-label training loss and the GP guidance loss, respectively. Algorithm 1 presents the pseudo-code for PG-LBO.

## Experiments

In this section, we apply PG-LBO to three high-dimensional structured optimization tasks and compare it with several VAE-BO algorithms. For each task, we first pre-train the VAE using unlabeled data, and all algorithms start with the same pre-trained VAE as the backbone network.

### High-Dimensional BO Tasks and Baselines

**Topology shape fitting task:** The goal of the task is to generate a  $40 \times 40$  binary image  $x$ , and **maximize** the cosine similarity  $\cos(\mathbf{x}, \mathbf{x}') = \mathbf{x} \cdot \mathbf{x}'^T / \|\mathbf{x}\| \|\mathbf{x}'\|$  between  $\mathbf{x}$  and a predefined target image  $\mathbf{x}'$ . The task involves using 10000 topology images from the dataset (Sosnovik and Oseledets 2019) and a VAE with the latent space dimension of 20.

**Expression reconstruction task:** The expression reconstruction task aims to generate single-variable expressions  $x(v)$  and **minimize** the distance to the target equation

$x'(v) = 1/3 + v * \sin(v * v)$ . The objective function is a distance metric  $f(x) = \max\{-7, -\int_{-10}^{10} \log(1 + (x(v) - x'(v))^2 dv)\}$ . Task access to 40,000 data points and use the grammar VAE from (Kusner, Paige, and Hernández-Lobato 2017) with the latent space dimension of 25.

**Chemical design task:** The task uses the ZINC250K dataset (Sterling and Irwin 2015) to synthesize chemical molecules with the objective of **maximizing** the penalized water-octanol distribution coefficient (PlogP) of molecules. A Junction Tree Variational Autoencoder (JT-VAE) (Jin, Barzilay, and Jaakkola 2018) with a latent space dimension of 56 encodes and generates efficient molecules.

**Baselines:** We compare our proposed method with four VAE-BO baselines: LSBO, LBO, T-LBO and LOL-BO. LSBO (Gómez-Bombarelli et al. 2018) performs BO in the latent space with a fixed pre-trained VAE. LBO trains VAE on the labeled data through data weighting and periodically fine-tuning VAE. T-LBO introduces deep metric learning to LBO by additionally minimizing the triplet loss of the labeled data. LOL-BO introduces the trust region optimization strategies that aided in searching for optimal points during the BO process. We follow the setups in the existing literature and utilize the same GP surrogate and acquisition function across all baselines. The surrogate is a sparse GP (Titsias 2009) with the radial basis function (RBF) kernel. The acquisition is the EI function.

## Experimental Setup and Results

**Experimental setup:** PG-LBO builds upon the foundation of LBO and uses the same VAE updating strategy and data weighting scheme during the BO process. In the training of pseudo-label data, the size of the pseudo-label dataset is maintained at half of the labeled dataset size, *i.e.*,  $N_P = N_L/2$ . As the BO iterations progress, the accuracy of pseudo-labels improves. Therefore, we linearly increase the weight of the pseudo-label loss during VAE retraining rounds. For topology shape fitting task,  $\lambda_P = \text{LinearIncrease}(0.5, 0.75)$ . For expression reconstruction task and chemical design task,  $\lambda_P = \text{LinearIncrease}(0.1, 0.75)$ . Regarding the GP guidance loss weight, we consider the varying difficulty levels of different tasks, and accordingly, the loss weight varies. For topology shape fitting task and chemical design task, the weight  $\lambda_G = 1$ , while for expression reconstruction task,  $\lambda_G = 0.1$ . The momentum decay of the pseudo-label selection threshold,  $\lambda = 0.9$ . The data sampling method employs noisy sampling, with Gaussian noise  $\mathcal{N}(0, 0.1)$ .

**Results:** As shown in Figure 2, PG-LBO consistently outperforms all baselines by the end of optimization. Table 1 shows more details. We notice that the improvement of PG-LBO over LBO is not substantial in the topology task. This might be attributed to the fact that LBO already performs well in this task, and due to the inherent difficulty of the problem, achieving significant enhancements becomes challenging. In expression task and molecule task, PG-LBO shows a notable improvement over LBO. The experimental results of T-LBO are obtained by replicating the code

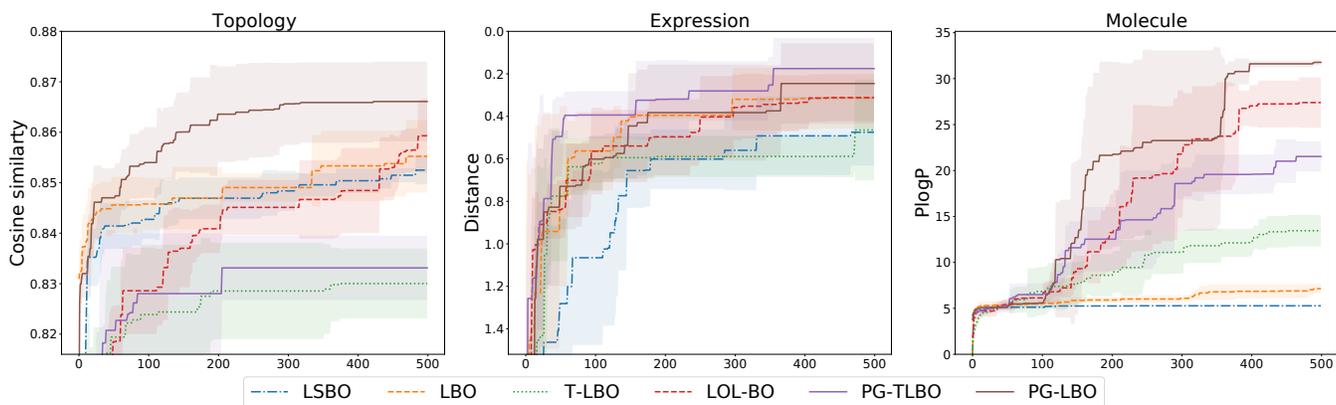


Figure 2: Comparison between PG-LBO and other VAE-BO methods in three tasks, plot the mean and the standard deviation of the best value over 5 seed over 500 evaluations. For topology task, the y-axis represents the cosine similarity between generated images and target images. For expression task, the y-axis represents the distance between generated expressions and target expressions, and for molecule task, the y-axis represents the PlogP score of generated molecules.

Method \ Task	Topology ( $\uparrow$ )	Expression ( $\downarrow$ )	Molecule ( $\uparrow$ )
LSBO	$0.852 \pm 0.003$	$0.474 \pm 0.157$	$5.45 \pm 0.252$
LBO	$0.855 \pm 0.007$	$0.314 \pm 0.144$	$7.01 \pm 0.517$
T-LBO	$0.830 \pm 0.007$	$0.464 \pm 0.238$	$13.44 \pm 1.711$
LOL-BO	$0.859 \pm 0.008$	$0.311 \pm 0.111$	$27.53 \pm 2.393$
<b>PG-TLBO (ours)</b>	$0.833 \pm 0.006$	<b><math>0.175 \pm 0.144</math></b>	$21.53 \pm 1.650$
<b>PG-LBO (ours)</b>	<b><math>0.866 \pm 0.007</math></b>	$0.227 \pm 0.186$	<b><math>31.77 \pm 0.329</math></b>

Table 1: Mean and standard deviation of best values found by baselines and PG-LBO. For Topology task, values represent cosine similarity, for Expression task, values represent distance, and for Molecule task, values represent PlogP score.  $\uparrow$  indicates higher is better,  $\downarrow$  indicates lower is better.

and parameters provided in T-LBO paper. However, the results in all three tasks significantly deviate from those reported in T-LBO paper. To address this, we conducted additional experiments by applying our proposed method to T-LBO, named PG-TLBO. Notably, the performance of model is improved after incorporating our method. In expression task, PG-TLBO ( $0.175 \pm 0.1443$ ) even outperforms PG-LBO ( $0.227 \pm 0.1864$ ), demonstrating the effectiveness of our proposed method. We also conducted comparative experiments on a smaller scale labeled dataset, our method consistently maintained superior performance over all baseline methods. The detailed discussion and comparison are in Appendix C.

## Ablation Studies

Unless otherwise specified, we conduct ablation studies on the topology shape fitting task. In the paper, we primarily present ablation experiments on the modules and key designs of the proposed methods. Ablation experiments for the hyperparameters are deferred to Appendix D.

**Individual effectiveness of Pseudo-labeled training and GP guidance.** To investigate the impact of the two main components of our method, pseudo-label training and GP

guidance, on model performance, we conduct ablation experiments and attempt to visualize the VAE latent space. **Experimental setup:** We compare three scenarios: using only pseudo-label training, using only GP guidance, and combining both techniques, denoted as P-LBO, G-LBO, and PG-LBO respectively. Other settings remain consistent. **Result:** As shown in Figure 3a, Pseudo-label training (P-LBO) and GP guidance (G-LBO) both contribute to performance improvements in the model, and for the topology task, Pseudo-label training provides more significant improvement to the model’s performance. This might be attributed to the fact that for topology tasks, the GP model has already fitted quite well, thus the improvement from GP guidance is limited. We conduct a visual analysis to further investigate the impact of pseudo-label training and GP guidance on the VAE latent space. We randomly sample 10,000 points from the trained VAE latent space and compute their corresponding true label values. We then perform PCA dimensionality reduction for visualization. The results are shown in Figure 4, the two axes are the principal components selected from the PCA analysis. The color bar represents the values of the selected property (cosine similarity). LSBO represents an unsupervised trained VAE. We can observe that the distribution of points in the latent space of P-LBO, G-LBO, and PG-LBO exhibits greater discriminative characteristics.

## The impact of data sampling on model performance.

The quality of sampling unlabeled data significantly impacts the model’s performance. We conduct ablation experiments to compare different sampling methods. **Experimental setup:** We compare three sampling methods: random sampling, heuristic sampling, and noisy sampling, denoted as PG-LBO (Random), PG-LBO (CMAES), and PG-LBO (Noise) respectively. In the Methodology section, we provide a detailed description of the implementation process for heuristic sampling and noisy sampling. We employ the Covariance Matrix Adaptation Evolution Strategy (CMAES) algorithm for heuristic sampling as the optimization algorithm. We select 100 high-scoring points from the exist-

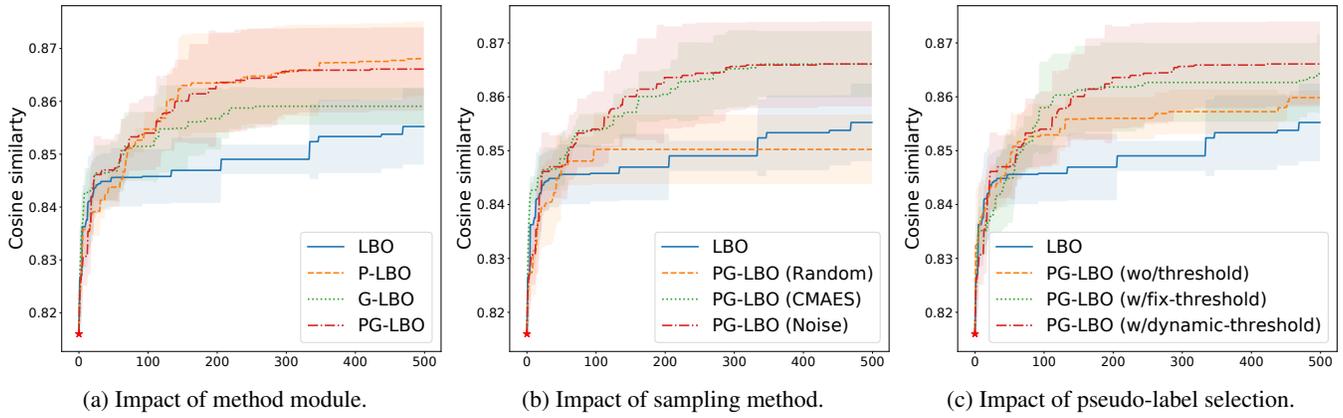


Figure 3: Ablation studies result, plot the mean and the standard deviation of the best value over 5 seed over 500 evaluations. The y-axis represents the cosine similarity.

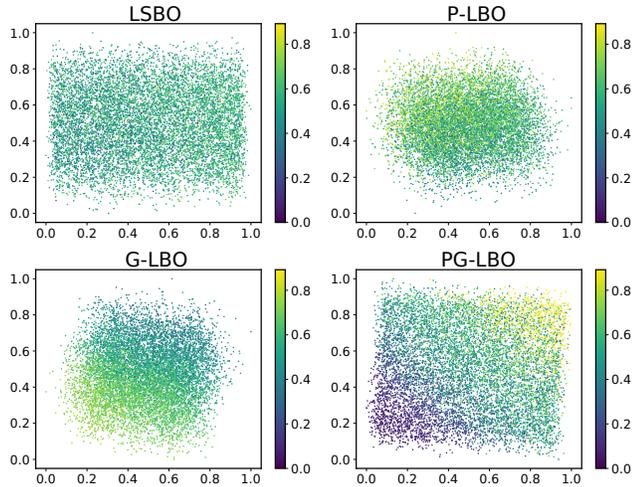


Figure 4: Two-dimensional PCA analysis on the VAE latent space for the topology shape fitting task.

ing labeled dataset as the initial means for the CMA-ES algorithm. The initial standard deviation is set to 0.25, and we perform 100 iterations, using the points generated in each iteration as sampling points. For noisy sampling, we sample noise from a Gaussian distribution  $\mathcal{N}(0, 0.1)$ . Similar to heuristic sampling, we select 100 high-scoring points from the existing labeled dataset as seed points, add random Gaussian noise to these seed points, and use the newly obtained points as sampling points. Aside from the sampling method, other parameter settings remain consistent. **Result:** As shown in Figure 3b, the performance of random sampling is the worst, even lower than the baseline (LBO). This indicates that the effectiveness of the pseudo-label training is highly dependent on the quality of the unlabeled data points sampled. If the sampled points are of poor quality, pseudo-label training might even degrade the performance compared to the baseline. Heuristic sampling and noisy sampling exhibit comparable performance. However, since heuristic

sampling involves executing an additional optimization algorithm, it introduces additional computational overhead. As a result, we ultimately choose simple and effective noisy sampling as the primary sampling method.

**The role of pseudo-label selection.** To improve the quality of pseudo-labels, we select pseudo-label data based on a dynamic threshold derived from uncertainty. We perform ablation experiments to investigate the impact of pseudo-label thresholds on model performance. **Experimental setup:** We compare three scenarios: using a fixed threshold, using a dynamic threshold, and using no threshold, denoted as PG-LBO (w/fix-threshold), PG-LBO (w/dynamic-threshold), and PG-LBO (wo/threshold) respectively. The fixed threshold is set to 0.0015, and the momentum decay  $\lambda$  for the dynamic threshold scheme is set to 0.9. Other experimental settings remain consistent. **Results:** As shown in Figure 3c, after threshold filtering, there is an improvement in model performance, indicating that enhancing the accuracy of pseudo-labels contributes to better model performance. The dynamic threshold filtering scheme performs better than the fixed threshold scheme, suggesting that dynamically adjusting the threshold based on the model’s learning state is beneficial. Moreover, the fixed threshold scheme requires task-specific prior knowledge to determine the threshold, whereas the dynamic threshold scheme does not.

## Conclusion

In this paper, we propose a novel method to enhance the construction of the VAE latent space by leveraging unlabeled data through pseudo-labeling techniques in semi-supervised learning. Furthermore, drawing inspiration from deep kernel learning, we directly incorporate the goal of improving GP prediction accuracy into the training of the VAE, providing a new way of exploiting labeled data. By combining our method with the weighted retraining proposed by LBO, we obtain a novel high-dimensional BO algorithm named PG-LBO. Extensive experiments on multiple task datasets demonstrate the effectiveness of our method in enhancing the performance of existing VAE-BO methods.

## Acknowledgments

This work is supported by the National Key R&D Program of China (2022ZD0116408), the Science and Technology Innovation 2030 New Generation Artificial Intelligence Major Projects (SQ2023AAA010051), NSFC Program (62222604, 62206052, 62192783), Jiangsu Natural Science Foundation Project (BK20210224), Huawei-Nanjing University Technical Cooperation Project (20221108058).

## References

- Biswas, A.; Vasudevan, R.; Ziatdinov, M.; and Kalinin, S. V. 2023. Optimizing training trajectories in variational autoencoders via latent Bayesian optimization approach. *Machine Learning: Science and Technology*, 4(1): 015011.
- Duan, Y.; Zhao, Z.; Qi, L.; Zhou, L.; Wang, L.; and Shi, Y. 2023. Towards Semi-supervised Learning with Non-random Missing Labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 16121–16131.
- Eissman, S.; Levy, D.; Shu, R.; Bartzsch, S.; and Ermon, S. 2018. Bayesian optimization and attribute adjustment. In *Proc. 34th Conference on Uncertainty in Artificial Intelligence*.
- Gómez-Bombarelli, R.; Wei, J. N.; Duvenaud, D.; Hernández-Lobato, J. M.; Sánchez-Lengeling, B.; Sheberla, D.; Aguilera-Iparraguirre, J.; Hirzel, T. D.; Adams, R. P.; and Aspuru-Guzik, A. 2018. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science*, 4(2): 268–276.
- Grosnit, A.; Tutunov, R.; Maraval, A. M.; Griffiths, R.-R.; Cowen-Rivers, A. I.; Yang, L.; Zhu, L.; Lyu, W.; Chen, Z.; Wang, J.; et al. 2021. High-dimensional Bayesian optimisation with variational autoencoders and deep metric learning. *arXiv preprint arXiv:2106.03609*.
- Gui, G.; Zhao, Z.; Qi, L.; Zhou, L.; Wang, L.; and Shi, Y. 2022. Improving Barely Supervised Learning by Discriminating Unlabeled Samples with Super-Class. In Oh, A. H.; Agarwal, A.; Belgrave, D.; and Cho, K., eds., *Advances in Neural Information Processing Systems*.
- Guo, Z.; Liu, H.; Ong, Y.-S.; Qu, X.; Zhang, Y.; and Zheng, J. 2022. Generative multimodal Bayesian optimization. *IEEE Transactions on Cybernetics*.
- Jean, N.; Xie, S. M.; and Ermon, S. 2018. Semi-supervised deep kernel learning: Regression with unlabeled data by minimizing predictive variance. *Advances in Neural Information Processing Systems*, 31.
- Jin, W.; Barzilay, R.; and Jaakkola, T. 2018. Junction tree variational autoencoder for molecular graph generation. In *International Conference on Machine Learning*, 2323–2332. PMLR.
- Kandasamy, K.; Schneider, J.; and Póczos, B. 2015. High dimensional Bayesian optimisation and bandits via additive models. In *International Conference on Machine Learning*, 295–304. PMLR.
- Korovina, K.; Xu, S.; Kandasamy, K.; Neiswanger, W.; Póczos, B.; Schneider, J.; and Xing, E. 2020. Chembo: Bayesian optimization of small organic molecules with synthesizable recommendations. In *International Conference on Artificial Intelligence and Statistics*, 3393–3403. PMLR.
- Kusner, M. J.; Paige, B.; and Hernández-Lobato, J. M. 2017. Grammar variational autoencoder. In *International Conference on Machine Learning*, 1945–1954. PMLR.
- Li, Z.; Qi, L.; Shi, Y.; and Gao, Y. 2023. IOMatch: Simplifying open-set semi-supervised learning with joint inliers and outliers utilization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15870–15879.
- Maus, N.; Jones, H.; Moore, J.; Kusner, M. J.; Bradshaw, J.; and Gardner, J. 2022. Local latent space bayesian optimization over structured inputs. *Advances in Neural Information Processing Systems*, 35: 34505–34518.
- Maus, N.; Zeng, Y.; Anderson, D. A.; Maffettone, P.; Solomon, A.; Greenside, P.; Bastani, O.; and Gardner, J. R. 2023. Inverse Protein Folding Using Deep Bayesian Optimization. *arXiv preprint arXiv:2305.18089*.
- Pham, H.; Dai, Z.; Xie, Q.; and Le, Q. V. 2021. Meta pseudo labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11557–11568.
- Rizve, M. N.; Duarte, K.; Rawat, Y. S.; and Shah, M. 2021. In Defense of Pseudo-Labeling: An Uncertainty-Aware Pseudo-label Selection Framework for Semi-Supervised Learning. In *International Conference on Learning Representations*.
- Ru, B.; Wan, X.; Dong, X.; and Osborne, M. 2021. Interpretable Neural Architecture Search via Bayesian Optimisation with Weisfeiler-Lehman Kernels. In *International Conference on Learning Representations*.
- Siivola, E.; Paleyes, A.; González, J.; and Vehtari, A. 2021. Good practices for Bayesian optimization of high dimensional structured spaces. *Applied AI Letters*, 2(2): e24.
- Snoek, J.; Larochelle, H.; and Adams, R. P. 2012. Practical bayesian optimization of machine learning algorithms. *Advances in Neural Information Processing Systems*, 25.
- Sosnovik, I.; and Oseledets, I. 2019. Neural networks for topology optimization. *Russian Journal of Numerical Analysis and Mathematical Modelling*, 34(4): 215–223.
- Stanton, S.; Maddox, W.; Gruver, N.; Maffettone, P.; Delaney, E.; Greenside, P.; and Wilson, A. G. 2022. Accelerating bayesian optimization for biological sequence design with denoising autoencoders. In *International Conference on Machine Learning*, 20459–20478. PMLR.
- Sterling, T.; and Irwin, J. J. 2015. ZINC 15—ligand discovery for everyone. *Journal of Chemical Information and Modeling*, 55(11): 2324–2337.
- Titsias, M. 2009. Variational learning of inducing variables in sparse Gaussian processes. In *Artificial Intelligence and Statistics*, 567–574. PMLR.
- Tripp, A.; Daxberger, E.; and Hernández-Lobato, J. M. 2020. Sample-efficient optimization in the latent space of deep generative models via weighted retraining. *Advances in Neural Information Processing Systems*, 33: 11259–11272.
- Verma, E.; Chakraborty, S.; and Griffiths, R.-R. 2022. High-dimensional Bayesian optimization with invariance. In

*ICML Workshop on Adaptive Experimental Design and Active Learning.*

Wang, Y.; Chen, H.; Heng, Q.; Hou, W.; Fan, Y.; Wu, Z.; Wang, J.; Savvides, M.; Shinozaki, T.; Raj, B.; Schiele, B.; and Xie, X. 2023. FreeMatch: Self-adaptive Thresholding for Semi-supervised Learning. In *The Eleventh International Conference on Learning Representations*.

Yang, Y.; Wei, H.; Sun, Z.; Li, G.; Zhou, Y.; Xiong, H.; and Yang, J. 2022. S2OSC: A Holistic Semi-Supervised Approach for Open Set Classification. *ACM Trans. Knowl. Discov. Data*, 16(2): 34:1–34:27.

Yin, Y.; Wang, Y.; and Li, P. 2023. High-dimensional Bayesian Optimization via Semi-supervised Learning with Optimized Unlabeled Data Sampling. *arXiv preprint arXiv:2305.02614*.

Zhang, M.; Jiang, S.; Cui, Z.; Garnett, R.; and Chen, Y. 2019. D-vae: A variational autoencoder for directed acyclic graphs. *Advances in Neural Information Processing Systems*, 32.

Zhang, X.; Wei, P.; and Wang, Q. 2023. A hybrid anomaly detection method for high dimensional data. *PeerJ Computer Science*, 9: e1199.