DENSITY RATIO ESTIMATION-BASED BAYESIAN OP-TIMIZATION WITH SEMI-SUPERVISED LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Bayesian optimization has attracted huge attention from diverse research areas in science and engineering, since it is capable of efficiently finding a global optimum of an expensive-to-evaluate black-box function. In general, a probabilistic regression model is widely used as a surrogate function to model an explicit distribution over function evaluations given an input to estimate and a training dataset. Beyond the probabilistic regression-based methods, density ratio estimation-based Bayesian optimization has been suggested in order to estimate a density ratio of the groups relatively close and relatively far to a global optimum. Developing this line of research further, supervised classifiers are employed to estimate a class probability for the two groups instead of a density ratio. However, the supervised classifiers used in this strategy are prone to be overconfident for known knowledge on global solution candidates. Supposing that we have access to unlabeled points, e.g., predefined fixed-size pools, we propose density ratio estimation-based Bayesian optimization with semi-supervised learning to solve this challenge. Finally, we show the empirical results of our methods and several baseline methods in two distinct scenarios with unlabeled point sampling and a fixed-size pool and analyze the validity of our proposed methods in diverse experiments.

004

006

008 009

010 011

012

013

014

015

016

017

018

019

020

021

024

025

1 INTRODUCTION

Bayesian optimization (Brochu et al., 2010; Garnett, 2023) has attracted immense attention from 031 various research areas such as hyperparameter optimization (Bergstra et al., 2011), battery lifetime optimization (Attia et al., 2020), chemical reaction optimization (Shields et al., 2021), and nanopho-033 tonic structure optimization (Kim et al., 2024), since it is capable of efficiently finding a global opti-034 mum of an expensive-to-evaluate black-box function. As studied in previous literature on Bayesian optimization (Snoek et al., 2012; Martinez-Cantin et al., 2018; Springenberg et al., 2016; Hutter et al., 2011), a probabilistic regression model, which can estimate a distribution of function evalu-037 ations over inputs, is widely used as a surrogate function; A Gaussian process (GP) (Rasmussen & 038 Williams, 2006) is a predominant choice for the surrogate function. An analogy between probabilistic regression models in Bayesian optimization is that they rely on an explicit function over function evaluations $p(y \mid \mathbf{x}, D)$ given an input to estimate, denoted as \mathbf{x} , and a training dataset D. 040

041 Beyond the probabilistic regression-based Bayesian optimization, density ratio estimation (DRE)-042 based Bayesian optimization has been studied recently (Bergstra et al., 2011; Tiao et al., 2021). 043 Furthermore, likelihood-free Bayesian optimization, which is equivalent to DRE-based Bayesian 044 optimization with a particular utility function, has been proposed by Song et al. (2022). Bergstra et al. (2011) attempt to model two densities $p(\mathbf{x} \mid y \leq y^{\dagger}, \mathcal{D})$ and $p(\mathbf{x} \mid y > y^{\dagger}, \mathcal{D})$, where y^{\dagger} is a threshold for dividing inputs to two groups that are relatively close and relatively far to a 046 global solution, in order to estimate ζ -relative density ratio (Yamada et al., 2011). On the other 047 hand, instead of modeling two densities separately, Tiao et al. (2021); Song et al. (2022) estimate 048 a density ratio using class-probability estimation (Qin, 1998). As discussed in the previous work, this line of research provides a new understanding of Bayesian optimization, which allows us to solve Bayesian optimization using binary classification. Moreover, it can reduce the amount of 051 computation required for building surrogate functions. 052

However, the supervised classifiers utilized in the DRE-based Bayesian optimization are prone to be overconfident for known knowledge on global solution candidates or a region that have been

054 already exploited. In this paper, supposing that we have access to unlabeled points, we propose a 055 novel DRE-based method with semi-supervised learning, which is named DRE-BO-SSL, to solve 056 this over-exploitation problem. Although our direct competitors, i.e., BORE (Tiao et al., 2021) and 057 LFBO (Song et al., 2022), show their strengths through theoretical and empirical analyses, our algo-058 rithm betters an ability to consider a wider region that satisfies $p(\mathbf{x} \mid y \leq y^{\dagger}, D) \geq p(\mathbf{x} \mid y > y^{\dagger}, D)$, than the competitors, as shown in Figure 1. By this intuitive example in Figure 1, we presume that DRE-BO-SSL appropriately balances exploration and exploitation rather than the existing meth-060 ods. Compared to a supervised classifier, e.g., random forests (Breiman, 2001), gradient boost-061 ing (Friedman, 2001), and multi-layer perceptrons, our semi-supervised classifiers, i.e., label propa-062 gation (Zhu & Ghahramani, 2002) and label spreading (Zhou et al., 2003), are less confident in terms 063 of the regions of global solution candidates by using unlabeled data points; see Figures 1 and 9 and 064 Sections 1.1 and 3 for detailed examples and analyses. 065

To make use of semi-supervised classifiers, we take into account two distinct scenarios with un-066 labeled point sampling and with a predefined fixed-size pool. For the first scenario, we randomly 067 sample unlabeled data points from the truncated multivariate normal distribution using a minimax 068 tilting method (Botev, 2017), in order that it is possible that a cluster assumption (Seeger, 2000) is 069 assumed. Finally, we demonstrate that our method shows superior performance compared to the exiting methods in diverse experiments including synthetic benchmarks, Tabular Benchmarks (Klein 071 & Hutter, 2019), NATS-Bench (Dong et al., 2021), and 64D minimum multi-digit MNIST search. 072 Note that Tabular Benchmarks, NATS-Bench, and 64D minimum multi-digit MNIST search have 073 access to fixed-size pools. Furthermore, in order to validate our methods, we provide thorough 074 analyses on the components of DRE-BO-SSL in Section 6 and the appendices.

- To sum up, our contributions are itemized as follows:
 - we identify the over-exploitation problem of supervised classifiers in DRE-based Bayesian optimization;
 - we propose DRE-based Bayesian optimization with semi-supervised learning, named DRE-BO-SSL for two distinct scenarios with unlabeled point sampling and a predefined fixed-size pool;
 - we demonstrate the effectiveness of our method in various experiments including NATS-Bench and 64D minimum multi-digit MNIST search.
 - We promise to release our implementation upon publication.
- 087 1.1 OVER-EXPLOITATION PROBLEM 088

077

078

079

081

082

084 085

As illustrated in Figures 1 and 9, the supervised classifiers used in DRE-based Bayesian optimiza-089 tion suffer from the over-exploitation problem. Interestingly, many deep learning models also share 090 a similar problem, which is called an overconfidence problem (Guo et al., 2017; Müller et al., 2019), 091 due to various reasons, but primarily due to overparameterization. It is noteworthy that the definition 092 of the over-exploitation problem in DRE-based Bayesian optimization is different from the overconfidence problem in general classification. The definition in DRE-based Bayesian optimization does 094 not imply that a single data point has a high probability for a particular class, but it implies that a few 095 data points have high probabilities for a class of interest. More precisely, our definition indicates 096 the problem of overconfidence over known knowledge on global solution candidates. For example, 097 the definition in general classification includes a case that 50% of data points are biased to one class 098 and the remainder is biased to another class. On the contrary, this definition does not include such a case and it only includes a case that a small number of data points or the small region of a search 099 space, which is colored in yellow in Figures 1 and 9, are biased to some particular class, i.e., Class 100 1 in the DRE-based Bayesian optimization problems. 101

By the aforementioned definition, at early iterations of Bayesian optimization, a supervised classifier tends to overfit to a small size of D_t due to a relatively large model capacity. This consequence makes a Bayesian optimization algorithm highly focus on exploitation. Similar to our observation, the imbalance of exploration and exploitation in the DRE-based approaches is also discussed in the recent work (Oliveira et al., 2022; Pan et al., 2024). Moreover, the consequence mentioned above is different from the characteristics of probabilistic regression-based Bayesian optimization because the regression-based methods are capable of exploring unseen regions by dealing with uncertainties.



(d) DRE-BO-SSL, Label spreading, Iterations 1 to 5

Figure 1: Comparisons of BORE and LFBO by multi-layer perceptrons, and DRE-BO-SSL with label propagation and label spreading for the Branin function. Each row shows Iterations 1 to 5 with five initial points. + (blue), \times (red), and \star (green) indicate data points of $y \leq y^{\dagger}$, data points of $y > y^{\dagger}$, and query points, respectively. Moreover, \circ (olive) stands for unlabeled points and its transparency represents the class probability predicted by a semi-supervised classifier. A query point is chosen by maximizing the output of the classifier. More results are shown in Figure 9.

142

143

132 133

134

135

136

137

138

Besides, even though a threshold ratio ζ might be able to mitigate this problem, an overconfident supervised classifier is likely to keep getting stuck in a local optimum as y^{\dagger} cannot change dramatically; more detailed analysis on ζ is provided in Sections 6 and H.

144 145 146

157

2 BACKGROUND AND RELATED WORK

Bayesian Optimization. It is a principled and efficient approach to finding a global solution of a 147 challenging objective, e.g., expensive-to-evaluate black-box functions (Brochu et al., 2010; Garnett, 148 2023). To focus on a probabilistic regression model as a surrogate function, we omit the details 149 of Bayesian optimization here; see (Brochu et al., 2010; Garnett, 2023) for details. In Bayesian 150 optimization, GP regression (Rasmussen & Williams, 2006) is widely used as a surrogate func-151 tion (Srinivas et al., 2010; Snoek et al., 2012) because of its flexibility with minimum assumptions 152 on model and smoothness. While GP is a probable choice, Bayesian optimization with diverse 153 surrogates such as Student-t process regression (Martinez-Cantin et al., 2018), Bayesian neural net-154 works (Springenberg et al., 2016), and tree-based models (Hutter et al., 2011; Kim & Choi, 2022) 155 has been proposed. An analogy between such models is that they model $p(y \mid \mathbf{x}, D)$ explicitly, so 156 that it can be used to define an acquisition function with the statistics of $p(y \mid \mathbf{x}, \mathcal{D})$.

Density-Ratio Estimation. Whereas knowing a data distribution $p(\mathbf{x})$ is important, it is difficult to directly estimate $p(\mathbf{x})$ (Sugiyama et al., 2012). For specific machine learning problems such as importance sampling (Kloek & van Dijk, 1978) and mutual information estimation (Bishop, 2006), we can bypass direct density estimation and estimate a density ratio. More recently, Rhodes et al. (2020) tackle a density-chasm problem using telescopic density-ratio estimation, which replaces an original problem with a set of logistic regression problems. Since the work (Rhodes et al., 2020) suffers from the issue on distribution shift, the recent work (Srivastava et al., 2023) proposes a density ratio estimation method with multinomial logistic regression, which is capable of mitigating the distribution shift using multi-class classification. In Bayesian optimization, Bergstra et al. (2011) have proposed a strategy with tree-structured Parzen estimators to estimate a density ratio as an alternative to probabilistic regression-based acquisition functions. In addition, the work (Tiao et al., 2021; Song et al., 2022) suggests methods with binary classifiers in order to estimate class probabilities as a density ratio; see Section 3 for details.

- 170 **Semi-Supervised Learning.** It is a learning scheme with both labeled and unlabeled data (Zhu, 171 2005; Chapelle et al., 2006; Bengio et al., 2006). To cooperate with labeled and unlabeled data, 172 this strategy generally utilizes geometry of data points or connectivity between points, and assigns 173 pseudo-labels to unlabeled data points, which referred to as transductive learning (Gammerman 174 et al., 1998). As a semi-supervised learning method on a similarity graph, Zhu & Ghahramani 175 (2002) propose a label propagation algorithm which iteratively propagates the labels of unlabeled 176 data points using labeled data. Zhou et al. (2003) compute the labels of labeled and unlabeled data 177 points by a weighted iterative algorithm with initial labels. Belkin & Niyogi (2002) predict pseudo-178 labels by finding a linear combination of a few smallest eigenvectors of the graph Laplacian.
- 179 180

181 182

183

184

185

186

187 188 189

202 203

211 212

3 DRE-BASED BAYESIAN OPTIMIZATION

Unlike probabilistic regression-based Bayesian optimization, DRE-based Bayesian optimization employs a density ratio-based acquisition function, defined with a density $p(\mathbf{x} | y \leq y^{\dagger}, \mathcal{D}_t)$, where \mathbf{x} is a *d*-dimensional vector, y is its function evaluation, y^{\dagger} is a threshold, and $\mathcal{D}_t = \{(\mathbf{x}_i, y_i)\}_{i=0}^t$ is a dataset of t + 1 pairs of data point and evaluation. In particular, the work by Bergstra et al. (2011) defines an acquisition function based on ζ -relative density ratio (Yamada et al., 2011):

$$A(\mathbf{x} \mid \zeta, \mathcal{D}_t) = \frac{p(\mathbf{x} \mid y \le y^{\dagger}, \mathcal{D}_t)}{\zeta p(\mathbf{x} \mid y \le y^{\dagger}, \mathcal{D}_t) + (1 - \zeta) p(\mathbf{x} \mid y > y^{\dagger}, \mathcal{D}_t)},$$
(1)

where $\zeta = p(y \leq y^{\dagger}) \in [0, 1)$ is a threshold ratio. We need to find a maximizer of (1), by optimizing the following composite function: $h(p(\mathbf{x} \mid y \leq y^{\dagger}, \mathcal{D}_t)/p(\mathbf{x} \mid y > y^{\dagger}, \mathcal{D}_t))$, where $h(x) = (\zeta + (1 - \zeta)x^{-1})^{-1}$. Since h is a strictly increasing function, we can directly maximize $p(\mathbf{x} \mid y \leq y^{\dagger}, \mathcal{D}_t)/p(\mathbf{x} \mid y > y^{\dagger}, \mathcal{D}_t)$. In the previous work (Bergstra et al., 2011), two treestructured Parzen estimators are used to estimate the respective densities, $p(\mathbf{x} \mid y \leq y^{\dagger}, \mathcal{D}_t)$ and $p(\mathbf{x} \mid y > y^{\dagger}, \mathcal{D}_t)$.

While the work (Bergstra et al., 2011) utilizes two distinct tree-structured Parzen estimators, Tiao et al. (2021) propose a method to directly estimate (1) using class-probability estimation (Qin, 1998; Sugiyama et al., 2012), which is called BORE. Since it can be formulated as a problem of binary classification in which Class 1 is a group of the top ζ of \mathcal{D}_t and Class 0 is a group of the bottom $1 - \zeta$ of \mathcal{D}_t in terms of function evaluations, the acquisition function defined in (1) induces the following:

$$A(\mathbf{x} \mid \zeta, \mathcal{D}_t) = \frac{p(\mathbf{x} \mid z = 1)}{\zeta p(\mathbf{x} \mid z = 1) + (1 - \zeta)p(\mathbf{x} \mid z = 0)}.$$
(2)

By the Bayes' theorem, (2) becomes $A(\mathbf{x} \mid \zeta, \mathcal{D}_t) = \zeta^{-1} p(z = 1 \mid \mathbf{x}) / (p(z = 1 \mid \mathbf{x}) + p(z = 0 \mid \mathbf{x}))$. Therefore, a class probability over \mathbf{x} for Class 1 is considered as an acquisition function; it is simply denoted as $A(\mathbf{x} \mid \zeta, \mathcal{D}_t) = \zeta^{-1} \pi(\mathbf{x})$. Eventually, the class probability is estimated by various off-the-shelf classifiers such as random forests and multi-layer perceptrons.

Song et al. (2022) have suggested a general framework of likelihood-free Bayesian optimization, called LFBO:

$$A(\mathbf{x} \mid \zeta, \mathcal{D}_t) = \underset{S:\mathcal{X} \to \mathbb{R}}{\arg\max} \mathbb{E}_{\mathcal{D}_t}[u(y; y^{\dagger})f'(S(\mathbf{x})) - f^*(f'(S(\mathbf{x})))],$$
(3)

which is versatile for any non-negative utility function $u(y; y^{\dagger})$, where f is a strictly convex function and f^* is the convex conjugate of f. By the properties of LFBO, it is equivalent to an expected utility-based acquisition function. Along with the general framework, Song et al. (2022) prove that BORE is equivalent to the probability of improvement (Kushner, 1964) and LFBO with $u(y; y^{\dagger}) =$

216 Algorithm 1 DRE-BO-SSL 217

| 218 | Input: Iteration budget T, a search space \mathcal{X} , a black-box objective f, a threshold ratio ζ , a semi- |
|-----|---|
| 210 | supervised classifier π_{C} , and unlabeled data points \mathbf{X}_{u} if available and the number of unlabeled |
| 219 | data pointe n otherwise |
| 220 | d_{u} a point h_{u} , the u |
| 001 | Output: Best candidate x ⁺ . |
| 221 | 1. Initialize $\mathcal{D} = \{(x_1, y_2)\}$ by readomly calculating x_2 from \mathcal{V} and evaluating x_2 by f |

1: Initialize $\mathcal{D}_0 = \{(\mathbf{x}_0, y_0)\}$ by randomly selecting \mathbf{x}_0 from \mathcal{X} and evaluating \mathbf{x}_0 by f.

- 2: for t = 0 to T 1 do
 - Calculate a threshold y_t^{T} using ζ . 3:
- Assign labels \mathbf{C}_t of data points in \mathcal{D}_t with y_t^{\dagger} . 4:
- 5: if \mathbf{X}_u are not available then
 - Sample n_u unlabeled data points \mathbf{X}_u from \mathcal{X} . 6:
- 7: 227 end if

222

223

224

225

226

228

229

230

232 233

235

236

237 238 239

240

241

- Estimate pseudo-labels C_t by following the procedure in Algorithm 2. 8:
- Choose the next query point \mathbf{x}_{t+1} by maximizing $\pi_{\widehat{\mathbf{C}}_t}(\mathbf{x}; \zeta, \mathcal{D}_t, \mathbf{X}_u)$ for $\mathbf{x} \in \mathcal{X}$. 9:
- Evaluate \mathbf{x}_{t+1} by f and update \mathcal{D}_{t+1} . 10:
- 11: end for 231

12: Determine the best candidate \mathbf{x}^+ , considering $y_{0:T}$.

 $\mathbb{I}(y \leq y^{\dagger})$ is also equivalent to the probability of improvement, where \mathbb{I} is an indicator function. Moreover, they show that their method with $u(y; y^{\dagger}) = \max(y^{\dagger} - y, 0)$, which can be implemented as weighted binary classification, is equivalent to the expected improvement (Jones et al., 1998).

4 DRE-BASED BAYESIAN OPTIMIZATION WITH SEMI-SUPERVISED LEARNING

242 We introduce DRE-based Bayesian optimization with semi-supervised learning, named DRE-BO-243 SSL, by following the previous studies in DRE-based and likelihood-free Bayesian optimiza-244 tion (Bergstra et al., 2011; Tiao et al., 2021; Song et al., 2022), which were discussed in Section 3. 245

Similar to the standard Bayesian optimization and existing DRE-based Bayesian optimization, DRE-246 BO-SSL iterates the undermentioned steps as presented in Algorithm 1. Firstly, a threshold y_t^{\dagger} is 247 calculated by considering $y_{1:t}$ with ζ . Secondly, labels \mathbf{C}_t of data points in \mathcal{D}_t are assigned to one 248 of two classes; a group of $y \le y_t^{\dagger}$ is assigned to Class 1 and a group of $y > y_t^{\dagger}$ is assigned to Class 0. 249 If we are given unlabeled data points X_u , the corresponding points X_u are used, but if not available 250 it samples n_u unlabeled data points \mathbf{X}_u from \mathcal{X} by utilizing a strategy described in Section 4.2. 251 Then, it estimates pseudo-labels $\widehat{\mathbf{C}}_t$ of a semi-supervised learning model by following the procedure 252 presented in Section 4.1 and Algorithm 2. Using $\hat{\mathbf{C}}_t$, it chooses the next query point \mathbf{x}_{t+1} : 253

$$\mathbf{x}_{t+1} = \arg\max_{\mathbf{x}\in\mathcal{X}} \pi_{\widehat{\mathbf{C}}_t}(\mathbf{x};\zeta,\mathcal{D}_t,\mathbf{X}_u),\tag{4}$$

254 255 256

257

where $\pi_{\widehat{\mathbf{C}}_t}(\mathbf{x}; \zeta, \mathcal{D}_t, \mathbf{X}_u)$ predicts a class probability over \mathbf{x} for Class 1; see (9).

258 We adopt a multi-started local optimization technique, e.g., L-BFGS-B (Byrd et al., 1995), to solve 259 (4). However, a flat landscape of $\pi_{\widehat{\mathbf{C}}_{t}}(\mathbf{x}; \zeta, \mathcal{D}_{t}, \mathbf{X}_{u})$ over \mathbf{x} may occur, so that optimization per-260 formance can be degraded. To tackle this issue, a simple heuristic of randomly selecting a query 261 point among points with identical highest class probabilities complements our method. Since the multi-started technique is utilized and the output of $\pi_{\hat{G}_t}$ is bounded in [0, 1], a flat landscape is easily 262 recognized by comparing the outcomes of the multi-started strategy. 263

264 265

266

4.1 LABEL PROPAGATION AND LABEL SPREADING

Here we describe semi-supervised learning techniques (Zhu, 2005; Chapelle et al., 2006; Bengio 267 et al., 2006) to build DRE-BO-SSL. We cover a transductive setting (Gammerman et al., 1998), 268 which is to label unlabeled data by utilizing given labeled data, and then an inductive setting, which 269 is to predict any point using pseudo-labels of unlabeled and labeled data.

Suppose that each data point is defined on a *d*-dimensional compact space $\mathcal{X} \subset \mathbb{R}^d$. We consider n_l labeled points $\mathbf{X}_l \in \mathbb{R}^{n_l \times d}$, their corresponding labels $\mathbf{C}_l \in \mathbb{R}^{n_l \times c}$, and n_u unlabeled points $\mathbf{X}_u \in \mathbb{R}^{n_u \times d}$, where *c* is the number of classes. \mathbf{X}_l and \mathbf{C}_l are query points that have already been evaluated and their class labels; we define $\mathbf{X}_l = [\mathbf{x}_0, \dots, \mathbf{x}_{n_l-1}]^\top$ for $\mathcal{D}_t = \{(\mathbf{x}_i, y_i)\}_{i=0}^{n_l-1}$. For the sake of brevity, the concatenated data points of \mathbf{X}_l and \mathbf{X}_u are defined as $\mathbf{X} = [\mathbf{X}_l; \mathbf{X}_u] \in$ $\mathbb{R}^{(n_l+n_u)\times d}$. Note that in our problem c = 2 because we address the problem with two classes.

As shown in Algorithm 2, we first initialize labels to propagate $\widehat{\mathbf{C}} \in \mathbb{R}^{(n_l+n_u)\times 2}$; it is initialized as $\widehat{\mathbf{C}} = [\mathbf{c}_0, \dots, \mathbf{c}_{n_l-1}, \mathbf{c}_{n_l}, \dots, \mathbf{c}_{n_l+n_u-1}]^\top$, where $\mathbf{c}_0, \dots, \mathbf{c}_{n_l-1}$ are one-hot representations [\mathbf{C}_l]_{1:}, ..., [\mathbf{C}_l]_{n_l}:, and $\mathbf{c}_{n_l}, \dots, \mathbf{c}_{n_l+n_u-1}$ are zero vectors. Denote that [\mathbf{C}]_i: is *i*th row of \mathbf{C} . Then, we compute a similarity between two data points $\mathbf{x}_i, \mathbf{x}_j \in \mathcal{X}$, so that we compare all pairs in \mathbf{X} . For example, one of popular similarity functions, i.e., a radial basis function kernel, can be used:

$$w_{ij} = \exp\left(-\beta \|\mathbf{x}_i - \mathbf{x}_j\|_2^2\right),\tag{5}$$

where β is a free parameter given. As discussed in the work (Zhu & Ghahramani, 2002), we can learn β in (5) by minimizing an entropy for propagated labels $\widehat{\mathbf{C}}$:

$$H(\widehat{\mathbf{C}}) = -\sum_{i=1}^{n_l+n_u} [\widehat{\mathbf{C}}]_{i:}^{\top} \log[\widehat{\mathbf{C}}]_{i:}.$$
(6)

See Figure 4 for the results of learning β and Section E for analysis on the effects of β . By (5), we compute a transition probability from \mathbf{x}_j to \mathbf{x}_i by $p_{ij} = w_{ij} / \sum_{k=1}^{n_l+n_u} w_{kj}$. Note that similarities $\mathbf{W} \in \mathbb{R}^{(n_l+n_u)\times(n_l+n_u)}$ and transition probabilities $\mathbf{P} \in \mathbb{R}^{(n_l+n_u)\times(n_l+n_u)}$ are defined, where $[\mathbf{W}]_{ij} = w_{ij}$ and $[\mathbf{P}]_{ij} = p_{ij}$. Moreover, by the definition of p_{ij} , $\mathbf{P} = \mathbf{D}^{-1}\mathbf{W}$, where \mathbf{D} is a degree matrix whose diagonal entry is defined as $[\mathbf{D}]_{ii} = \sum_{j=1}^{n_l+n_u} [\mathbf{W}]_{ij}$.

With initial $\hat{\mathbf{C}}$ and \mathbf{P} , we repeat the following steps: (i) computing the next $\hat{\mathbf{C}}$; (ii) normalizing $\hat{\mathbf{C}}$ row-wise, until a change of $\hat{\mathbf{C}}$ converges to a tolerance ε or the number of iterations propagated reaches to maximum iterations τ . In particular, label propagation (Zhu & Ghahramani, 2002) updates $\hat{\mathbf{C}}$ and constrains the labels of labeled data to \mathbf{C}_l :

$$\widehat{\mathbf{C}}_{t+1} \leftarrow \mathbf{P}^{\top} \widehat{\mathbf{C}}_t \quad \text{and} \quad [\widehat{\mathbf{C}}_{t+1}]_{i:} \leftarrow [\mathbf{C}_l]_{i:},$$
(7)

for $i \in [n_l]$ at Iteration t. On the other hand, label spreading (Zhou et al., 2003) propagates $\widehat{\mathbf{C}}$ by allowing a change of the labels of labeled data with a clamping factor $\alpha \in (0, 1)$:

$$\widehat{\mathbf{C}}_{t+1} \leftarrow \alpha \mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2} \widehat{\mathbf{C}}_t + (1-\alpha) \widehat{\mathbf{C}}_0, \tag{8}$$

where $\widehat{\mathbf{C}}_0$ is initial propagated labels, which are defined in Line 1 of Algorithm 2. Note that $\mathbf{D}^{-1/2}\mathbf{W}\mathbf{D}^{-1/2}$ can be pre-computed before the loop defined from Lines 4 to 7 of Algorithm 2.

By the nature of transductive setting, it only predicts the labels of particular data using the known data (Gammerman et al., 1998), which implies that it cannot estimate a categorical distribution of unseen data. To enable it, given unseen x, we define an inductive model with \hat{C} :

$$\widehat{c}_{i} = \frac{\mathbf{w}^{\top}[\widehat{\mathbf{C}}]_{:i}}{\sum_{j=1}^{c} \mathbf{w}^{\top}[\widehat{\mathbf{C}}]_{:j}},\tag{9}$$

for $i \in [2]$, where $\mathbf{w} \in \mathbb{R}^{n_l+n_u}$ is similarities of x and X by (5). $[\widehat{\mathbf{C}}]_{:i}$ denotes *i*th column of $\widehat{\mathbf{C}}$. This inductive model is better than or equivalent to other classifiers without unlabeled data in certain circumstances; its analysis is provided in Section C.

316 317

282 283

284

285

287

299 300

301

302 303

304 305

306

307

308

309 310

311 312

4.2 UNLABELED POINT SAMPLING

As described above, if unlabeled points are not available, we need to generate unlabeled points under a transductive learning scheme. However, it is a challenging problem unless we know a landscape of $\pi_{\mathbf{C}}$ adequately. Many studies (Seeger, 2000; Rigollet, 2007; Singh et al., 2008; Ben-David et al., 2008; Carmon et al., 2019; Wei et al., 2020; Zhang et al., 2022) investigate how unlabeled data can affect a model and whether unlabeled data helps improve the model or not.

In order to make use of the theoretical findings of previous literature, we define a cluster assumption:



Figure 2: Results with 20 repeated experiments on synthetic benchmark functions for a scenario with unlabeled point sampling. LP and LS stand for label propagation and label spreading, respectively.



Figure 3: Results with 20 repeated experiments on synthetic benchmark functions for a scenario with fixed-size pools. LP and LS stand for label propagation and label spreading, respectively.

Assumption 4.1 (Cluster assumption in (Seeger, 2000)). Two points $\mathbf{x}_i, \mathbf{x}_j \in \mathcal{X}$ should belong to the same label if there is a path between \mathbf{x}_i and \mathbf{x}_j which passes only through regions of relatively high P(X), where P is a distribution over a random variable $X \in \mathcal{X}$.

By Assumption 4.1, the idea of clustering on the Euclidean space or spectral clustering on a graph
can be naturally applied in semi-supervised learning (Seeger, 2000; Joachims, 2003), which is not
the scope of this work.

To build DRE-BO-SSL associated with Assumption 4.1, we sample unlabeled data points from the truncated multivariate normal distribution so that each sample is in a compact \mathcal{X} :

$$f(\mathbf{z}) = \frac{\exp(-\frac{1}{2}\mathbf{z}^{\top}\mathbf{z})\mathbb{I}(\mathbf{l} \le \mathbf{A}\mathbf{z} \le \mathbf{u})}{P(\mathbf{l} \le \mathbf{A}Z \le \mathbf{u})},\tag{10}$$

360 where $\mathbf{l}, \mathbf{u} \in \mathbb{R}^d$ are lower and upper bounds, $\Sigma = \mathbf{A}\mathbf{A}^{\top}$ is a covariance matrix, \mathbb{I} is an indicator 361 function, and $Z \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ is a random variable. It is challenging to calculate a denominator 362 of (10), $P(1 \le \mathbf{A}^{\top} Z \le \mathbf{u})$, and simulate from $f(\mathbf{z})$ because an integration of the denominator and 363 an accept-reject sampling strategy from f(z) are cumbersome in this multi-dimensional case. To 364 effectively sample from the truncated multivariate normal distribution, we adopt the minimax tilting 365 method (Botev, 2017). Compared to the method by Genz (1992), it yields a high acceptance rate 366 and accurate sampling. In this paper Σ is set as an identity matrix, and I and u are determined by a search space. We will provide more detailed discussion on point sampling in Sections 6 and F. 367

368 369 370

334

339

340 341 342

343

344 345

346

347 348

349

350

351

357 358 359

5 EXPERIMENTS

We test DRE-BO-SSL and baseline methods in the following optimization problems: synthetic benchmarks for a scenario with unlabeled point sampling, and synthetic benchmarks, Tabular Benchmarks (Klein & Hutter, 2019), NATS-Bench (Dong et al., 2021), and minimum multi-digit MNIST search for a scenario with a fixed-size pool. Note that Tabular Benchmarks, NATS-Bench, and minimum multi-digit MNIST search are defined with a fixed number of possible solution candidates, which implies that they are considered as combinatorial optimization problems. By following the previous work (Tiao et al., 2021; Song et al., 2022), we set a threshold ratio as $\zeta = 0.33$ for all experiments; the effects of a threshold ratio ζ are analyzed in Sections 6 and H. To solve (4), we



Figure 4: Results with 20 repeated experiments on learning β for label propagation, which is denoted as LP, and label spreading, which is denoted as LS. Sampling and pool indicate the experiments in Sections 5.1 and 5.2, respectively.



Figure 5: Results with 20 repeated experiments on Tabular Benchmarks for a scenario with fixedsize pools. LP and LS stand for label propagation and label spreading, respectively.

use L-BFGS-B (Byrd et al., 1995) with 1000 different initializations. All experiments are repeated
20 times with 20 fixed random seeds and the sample mean and the standard error of the sample
mean are reported. Other missing details including the details of the competitors of our methods are
presented in Section D.

410 411

412

388

389

390 391

392

393

394

395

397

398 399

400

401 402

403

404 405

5.1 A SCENARIO WITH UNLABELED POINT SAMPLING

413 Synthetic Benchmarks. We run several synthetic functions for our methods and the baseline 414 methods. For unlabeled point sampling, we sample 100 points from $\mathcal{N}(\mathbf{x}_1, \mathbf{I}_d), \ldots, \mathcal{N}(\mathbf{x}_{n_l}, \mathbf{I}_d)$, 415 where $\lfloor n_u/n_l \rfloor$ or $\lfloor n_u/n_l \rfloor + 1$ points are sampled from each truncated distribution, so that n_u 416 points are sampled in total. As shown in Figure 2, our methods outperform the baseline methods in-417 corporating labeled data with unlabeled points. Interestingly, out methods beat GP-based Bayesian 418 optimization. It implies that ours can fairly balance exploration and exploitation. Furthermore, we 419 present the results of learning β in Figure 4, where β is adaptively selected by (6) every iteration.

- 419 420 421
- 5.2 Scenarios with Fixed-Size Pools

422 **Synthetic Benchmarks.** Several synthetic benchmark functions are tested for our methods and 423 the baseline methods. To generate a fixed-size pool for each benchmark, we uniformly sample 1000 424 points from a bounded search space before an optimization round is started. As presented in Figure 3, 425 our methods perform better than the baseline methods. It implies that the use of unlabeled data helps 426 improve optimization performance. Also, the results of learning β are demonstrated in Figure 4. 427 Learned β is likely to converge to some value as iterations proceed according to the results.

428

Tabular Benchmarks. Comparisons of our methods and the existing methods are carried out in these hyperparameter optimization benchmarks (Klein & Hutter, 2019), as reported in Figure 5. We can benchmark a variety of machine learning models, which are defined with specific hyperparameters and trained on one of four datasets: naval propulsion, parkinsons telemonitoring, and slice



Figure 6: Results with 20 repeated experiments on NATS-Bench for a scenario with fixed-size pools.

localization. There exist 62,208 models, which are used as a pool in this paper, for each dataset. Our algorithms show superior performance compared to other approaches. Similar to the synthetic functions, we argue that the adoption of a predefined pool leverages its performance. In some cases, GP-based Bayesian optimization is better than our methods.

NATS-Bench. NATS-Bench (Dong et al., 2021), which is the up-to-date version of NAS-Bench-201 (Dong & Yang, 2019), is used to test our methods and the baseline methods. NATS-Bench is a neural architecture search benchmark with three popular datasets: CIFAR-10, CIFAR-100, and ImageNet-16-120, and it has 32,768 architectures, i.e., a fixed-size pool in this paper, for each dataset. Similar to the experiments mentioned above, our methods work well in three datasets, compared to the existing methods; see Figure 6 for the results.

458 64D Minimum Multi-Digit MNIST 459 Search. This task, which is proposed 460 in this work, is to find a minimum multi-461 digit number, where a fixed number of 462 multi-digit MNIST images are given. 463 As visualized in Figure 7, three random images in the MNIST dataset (LeCun et al., 464 1998) are concatenated. Eventually, "000" 465 and "999" are global minimum and global 466 maximum, respectively. Since inputs 467 are images and each concatenated image 468 consists of three different digit images, this 469 high-dimensional optimization problem is 470 challenging. The size of a fixed-size pool is 471 80,000. As shown in Figure 7, our methods 472 show satisfactory performance compared 473 to other baseline methods.

DISCUSSION



Figure 7: Examples and results on 64D minimum multi-digit MNIST search. Figure 7a shows some examples of multi-digit MNIST and Figure 7b demonstrates the results with 20 repeated experiments.

474

444 445 446

447

448

449

450 451

452

453

454

455

456

457

475 476

6

477 478

We discuss interesting topics on our methods and the properties of DRE-BO-SSL. Moreover, we provide more thorough discussion, limitations, and broader impacts in the appendices.

479 480

Effects of the Number of Points and Sampling Distributions for Unlabeled Points. We choose
 a distribution for unlabeled point sampling as the truncated multivariate normal distribution in order
 to satisfy the cluster assumption. To analyze our algorithm thoroughly, we demonstrate the effects of
 the number of sampled points and sampling distributions for unlabeled points in Section F, varying
 the number of unlabeled points and utilizing diverse sampling distributions, i.e., uniform distribu tions, Halton sequences (Halton, 1960), and Sobol' sequences (Sobol', 1967).



Figure 8: Results with 20 repeated experiments on elapsed times varying subset sizes via pool sampling. LP and LS stand for label propagation and label spreading, respectively.

Effects of Pool Sampling. Because the computational complexity of label propagation and label spreading depends on a pool size, we need to reduce a pool size appropriately in order to speed up the algorithms. Therefore, we uniformly sample a subset of the fixed-size set, which is used as unlabeled points. Figure 8 reports elapsed times over subset sizes for pool sampling. Larger subsets make the framework slower as expected. More detailed analysis on the effects of pool sampling is demonstrated in Section G.

Effects of Threshold Ratios. We study the effects of a threshold ratio ζ in order to understand how we can choose ζ . As shown in Figure 15, both small ζ and large ζ generally represent worse performance; see the results with $\zeta = 0.01$ and $\zeta = 0.8$. While it has to depend on optimization problems, $\zeta = 0.33$ and $\zeta = 0.5$ are generally appropriate choices according to this analysis. The details of this study can be found in Section H and Figure 15.

Flat Landscape of Class Probabilities over Inputs. Regardless of the use of either supervised or
 semi-supervised classifier, a flat landscape of class probabilities can occur in the framework of DRE based Bayesian optimization. To overcome the issue of optimizing a flat landscape, we add a simple
 heuristic of selecting a random query point from points with identical highest class probabilities
 if the landscape is flat, as described in Section 4. Since off-the-shelf local optimization methods
 struggle to optimize a flat landscape, this simple heuristic is effective.

General Framework of DRE-BO-SSL. As a future direction, we expect that a general framework of DRE-BO-SSL can be defined, which is similar to a likelihood-free approach by Song et al. (2022). However, it is difficult to define an utility function involved in both labeled and unlabeled data. For example, if we assume that an utility function is $u(y; y^{\dagger}) = \max(y^{\dagger} - y, 0), y$ for an unlabeled data point is unknown. Although it depends on the form of utility function, we need to define y of an unlabeled data point by utilizing the information we have if the utility function is related to y.

532 533 534

535

504

505 506

513

519

7 CONCLUSION

In this paper we have proposed a DRE-based Bayesian optimization framework with semi-supervised learning, named DRE-BO-SSL. Unlike the existing work such as BORE and LFBO, our methods make use of semi-supervised classifiers where unlabeled data points are sampled or given. The superior results by our methods and the thorough analyses on the characteristics of DRE-BO-SSL exhibit the validity of our proposed algorithms.

540 REFERENCES 541

549

550

551

552

553

554

556

558

563

568

569

573

574 575

576

577

578

579

580

581

583

585

586

587

589

- P. M. Attia, A. Grover, N. Jin, K. A. Severson, T. M. Markov, Y.-H. Liao, M. H. Chen, B. Cheong, 542 N. Perkins, Z. Yang, P. K. Herring, M. Aykol, S. J. Harris, R. D. Braatz, S. Ermon, and W. C. 543 Chueh. Closed-loop optimization of fast-charging protocols for batteries with machine learning. 544 Nature, 578(7795):397-402, 2020.
- 546 M. Belkin and P. Niyogi. Using manifold stucture for partially labeled classification. In Advances 547 in Neural Information Processing Systems (NeurIPS), volume 15, Vancouver, British Columbia, 548 Canada, 2002.
 - S. Ben-David, T. Lu, and D. Pál. Does unlabeled data provably help? worst-case analysis of the sample complexity of semi-supervised learning. In Proceedings of the Annual Conference on Learning Theory (COLT), pp. 33-44, 2008.
 - Y. Bengio, O. Delalleau, and N. Le Roux. Label propagation and quadratic criterion. In O. Chapelle, B. Schölkopf, and A. Zien (eds.), Semi-Supervised Learning, pp. 193–216. MIT Press, 2006.
 - J. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl. Algorithms for hyper-parameter optimization. In Advances in Neural Information Processing Systems (NeurIPS), volume 24, pp. 2546–2554, Granada, Spain, 2011.
- C. M. Bishop. Pattern Recognition and Machine Learning. Springer, 2006. 559
- Z. I. Botev. The normal law under linear restrictions: simulation and estimation via minimax tilting. 561 Journal of the Royal Statistical Society B, 79(1):125–148, 2017. 562
 - L. Breiman. Random forests. Machine Learning, 45(1):5-32, 2001.
- 564 E. Brochu, V. M. Cora, and N. de Freitas. A tutorial on Bayesian optimization of expensive cost 565 functions, with application to active user modeling and hierarchical reinforcement learning. arXiv 566 preprint arXiv:1012.2599, 2010. 567
 - R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu. A limited memory algorithm for bound constrained optimization. SIAM Journal on Scientific Computing, 16(5):1190–1208, 1995.
- 570 Y. Carmon, A. Raghunathan, L. Schmidt, J. C. Duchi, and P. S. Liang. Unlabeled data improves 571 adversarial robustness. In Advances in Neural Information Processing Systems (NeurIPS), vol-572 ume 32, pp. 11192–11203, Vancouver, British Columbia, Canada, 2019.
 - O. Chapelle, B. Schölkopf, and A.Zien. Semi-Supervised Learning. MIT Press, 2006.
 - T. Chen and C. Guestrin. XGBoost: A scalable tree boosting system. In Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), pp. 785–794, San Francisco, California, USA, 2016.
 - X. Dong and Y. Yang. NAS-Bench-201: Extending the scope of reproducible neural architecture search. In Proceedings of the International Conference on Learning Representations (ICLR), New Orleans, Louisiana, USA, 2019.
- X. Dong, L. Liu, K. Musial, and B. Gabrys. NATS-Bench: Benchmarking NAS algorithms for 582 architecture topology and size. IEEE Transactions on Pattern Analysis and Machine Intelligence, 44(7):3634–3646, 2021. 584
 - J. H. Friedman. Greedy function approximation: a gradient boosting machine. The Annals of Statistics, 29:1189–1232, 2001.
- A. Gammerman, V. Vovk, and V. Vapnik. Learning by transduction. In Proceedings of the Annual 588 Conference on Uncertainty in Artificial Intelligence (UAI), pp. 148–155, Madison, Wisconsin, USA, 1998.
- R. Garnett. Bayesian Optimization. Cambridge University Press, 2023. 592
- E. C. Garrido-Merchán and D. Hernández-Lobato. Dealing with categorical and integer-valued variables in Bayesian optimization with Gaussian processes. Neurocomputing, 380:20-35, 2020.

621

623

624

625

626 627

628

629

630

631

634

637

638

- 594 A. Genz. Numerical computation of multivariate normal probabilities. Journal of Computational 595 and Graphical Statistics, 1(2):141–149, 1992. 596
- C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger. On calibration of modern neural networks. In Pro-597 ceedings of the International Conference on Machine Learning (ICML), pp. 1321-1330, Sydney, 598 Australia, 2017.
- 600 J. H. Halton. On the efficiency of certain quasi-random sequences of points in evaluating multi-601 dimensional integrals. Numerische Mathematik, 2:84-90, 1960. 602
- C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, 603 J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, 604 A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, 605 W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant. Array programming with NumPy. 606 Nature, 585:357-362, 2020. 607
- 608 F. Hutter, H. H. Hoos, and K. Leyton-Brown. Sequential model-based optimization for general algo-609 rithm configuration. In Proceedings of the International Conference on Learning and Intelligent 610 Optimization (LION), pp. 507–523, Rome, Italy, 2011.
- T. Joachims. Transductive learning via spectral graph partitioning. In Proceedings of the Interna-612 tional Conference on Machine Learning (ICML), pp. 290–297, Washington, District of Columbia, 613 USA, 2003. 614
- 615 D. R. Jones, M. Schonlau, and W. J. Welch. Efficient global optimization of expensive black-box 616 functions. Journal of Global Optimization, 13:455-492, 1998.
- 617 J. Kim and S. Choi. On uncertainty estimation by tree-based surrogate models in sequential model-618 based optimization. In Proceedings of the International Conference on Artificial Intelligence and 619 Statistics (AISTATS), pp. 4359–4375, Virtual, 2022. 620
- J. Kim and S. Choi. BayesO: A Bayesian optimization framework in Python. Journal of Open 622 Source Software, 8(90):5320, 2023.
 - J. Kim, M. Li, Y. Li, A. Gómez, O. Hinder, and P. W. Leu. Multi-BOWS: Multi-fidelity multiobjective Bayesian optimization with warm starts for nanophotonic structure design. Digital Discovery, 3(2):381–391, 2024.
 - D. P. Kingma and J. L. Ba. Adam: A method for stochastic optimization. In Proceedings of the International Conference on Learning Representations (ICLR), San Diego, California, USA, 2015.
 - A. Klein and F. Hutter. Tabular benchmarks for joint architecture and hyperparameter optimization. arXiv preprint arXiv:1905.04970, 2019.
- 632 T. Kloek and H. K. van Dijk. Bayesian estimates of equation system parameters: an application of 633 integration by Monte Carlo. Econometrica: Journal of the Econometric Society, pp. 1–19, 1978.
- H. J. Kushner. A new method of locating the maximum point of an arbitrary multipeak curve in the 635 presence of noise. Journal of Basic Engineering, 86(1):97–106, 1964. 636
 - Y. LeCun, C. Cortes, and C. J. C. Burges. The MNIST database of handwritten digits. http: //yann.lecun.com/exdb/mnist/,1998.
- R. Martinez-Cantin, K. Tee, and M. McCourt. Practical Bayesian optimization in the presence of 640 outliers. In Proceedings of the International Conference on Artificial Intelligence and Statistics 641 (AISTATS), pp. 1722–1731, Lanzarote, Canary Islands, Spain, 2018. 642
- 643 R. Müller, S. Kornblith, and G. E. Hinton. When does label smoothing help? In Advances in Neural 644 Information Processing Systems (NeurIPS), volume 32, Vancouver, British Columbia, Canada, 645 2019. 646
- E. A. Nadaraya. On estimating regression. Theory of Probability and Its Applications, 9(1):141–142, 647 1964.

660

661

662

663

664

665

666

681

684

685

686 687

688

689

691

693

- 648 R. Oliveira, L. Tiao, and F. T. Ramos. Batch Bayesian optimisation via density-ratio estimation with 649 guarantees. In Advances in Neural Information Processing Systems (NeurIPS), volume 35, pp. 650 29816–29829, New Orleans, Louisiana, USA, 2022.
- J. Pan, S. Falkner, F. Berkenkamp, and J. Vanschoren. MALIBO: Meta-learning for likelihood-free 652 Bayesian optimization. In Proceedings of the International Conference on Machine Learning 653 (ICML), pp. 39102–39134, Vienna, Austria, 2024. 654
- 655 A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, 656 L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An imperative style, high-performance deep 657 learning library. In Advances in Neural Information Processing Systems (NeurIPS), volume 32, 658 pp. 8026-8037, Vancouver, British Columbia, Canada, 2019. 659
 - F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and É. Duchesnay. Scikit-learn: Machine learning in Python. Journal of Machine Learning Research, 12:2825-2830, 2011.
 - J. Qin. Inferences for case-control and semiparametric two-sample density ratio models. *Biometrika*, 85(3):619-630, 1998.
- 667 C. E. Rasmussen and C. K. I. Williams. Gaussian Processes for Machine Learning. MIT Press, 668 2006.
- 669 B. Rhodes, K. Xu, and M. U. Gutmann. Telescoping density-ratio estimation. In Advances in Neural 670 Information Processing Systems (NeurIPS), volume 33, pp. 4905–4916, Virtual, 2020. 671
- 672 P. Rigollet. Generalization error bounds in semi-supervised classification under the cluster assump-673 tion. Journal of Machine Learning Research, 8(7):1369–1392, 2007.
- 674 M. Seeger. Learning with labeled and unlabeled data. Technical Report Technical Report, University 675 of Edinburgh, 2000. 676
- 677 B. J. Shields, J. Stevens, J. Li, M. Parasram, F. Damani, J. I. M. Alvarado, J. M. Janey, R. P. Adams, 678 and A. G. Doyle. Bayesian reaction optimization as a tool for chemical synthesis. *Nature*, 590 (7844):89–96, 2021. 679
- 680 A. Singh, R. D. Nowak, and X. Zhu. Unlabeled data: now it helps, now it doesn't. In Advances in Neural Information Processing Systems (NeurIPS), volume 21, pp. 1513–1520, Vancouver, 682 British Columbia, Canada, 2008. 683
 - J. Snoek, H. Larochelle, and R. P. Adams. Practical Bayesian optimization of machine learning algorithms. In Advances in Neural Information Processing Systems (NeurIPS), volume 25, pp. 2951–2959, Lake Tahoe, Nevada, USA, 2012.
 - I. M. Sobol'. On the distribution of points in a cube and the approximate evaluation of integrals. Zhurnal Vychislitel'noi Matematiki i Matematicheskoi Fiziki, 7(4):784–802, 1967.
- J. Song, L. Yu, W. Neiswanger, and S. Ermon. A general recipe for likelihood-free Bayesian op-690 timization. In Proceedings of the International Conference on Machine Learning (ICML), pp. 20384-20404, Baltimore, Maryland, USA, 2022. 692
 - J. T. Springenberg, A. Klein, S. Falkner, and F. Hutter. Bayesian optimization with robust Bayesian neural networks. In Advances in Neural Information Processing Systems (NeurIPS), volume 29, pp. 4134–4142, Barcelona, Spain, 2016.
- 696 N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian process optimization in the bandit 697 setting: No regret and experimental design. In Proceedings of the International Conference on 698 Machine Learning (ICML), pp. 1015–1022, Haifa, Israel, 2010. 699
- A. Srivastava, S. Han, K. Xu, B. Rhodes, and M. U. Gutmann. Estimating the density ratio be-700 tween distributions with high discrepancy using multinomial logistic regression. Transactions on Machine Learning Research, 2023.

- M. Sugiyama, T. Suzuki, and T. Kanamori. Density ratio estimation in machine learning. Cambridge University Press, 2012.
- L. C. Tiao, A. Klein, M. Seeger, E. V. Bonilla, C. Archambeau, and F. Ramos. BORE: Bayesian optimization by density-ratio estimation. In Proceedings of the International Conference on Machine Learning (ICML), pp. 10289–10300, Virtual, 2021.
- P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, I. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: fundamental algorithms for scientific computing in Python. Nature Methods, 17:261-272, 2020.
- G. S. Watson. Smooth regression analysis. Sankhyā: The Indian Journal of Statistics, Series A, 26: 359-372, 1964.
- C. Wei, K. Shen, Y. Chen, and T. Ma. Theoretical analysis of self-training with deep networks on unlabeled data. In Proceedings of the International Conference on Learning Representations (ICLR), Virtual, 2020.
- M. Yamada, T. Suzuki, T. Kanamori, H. Hachiya, and M. Sugiyama. Relative density-ratio estima-tion for robust distribution comparison. In Advances in Neural Information Processing Systems (NeurIPS), volume 24, pp. 594-602, Granada, Spain, 2011.
- S. Zhang, M. Wang, S. Liu, P.-Y. Chen, and J. Xiong. How does unlabeled data improve generaliza-tion in self-training? a one-hidden-layer theoretical analysis. In Proceedings of the International Conference on Learning Representations (ICLR), Virtual, 2022.
- D. Zhou, O. Bousquet, T. N. Lal, J. Weston, and B. Schölkopf. Learning with local and global consistency. In Advances in Neural Information Processing Systems (NeurIPS), volume 16, pp. 321–328, Vancouver, British Columbia, Canada, 2003.
 - X. Zhu. Semi-supervised learning literature survey. Technical Report Computer Sciences TR-1530, University of Wisconsin–Madison, 2005.
- X. Zhu and Z. Ghahramani. Learning from labeled and unlabeled data with label propagation. Technical Report CMU-CALD-02-107, Carnegie Mellon University, 2002.

ADDITIONAL COMPARISONS OF BORE AND LFBO



(f) XGBoost, LFBO, Iterations 1 to 5

Figure 9: Comparisons of BORE and LFBO by random forests, gradient boosting, and XGBoost for the Branin function. It follows the configurations described in Figure 1.

In addition to Figure 1, we include additional comparisons of BORE and LFBO by random forests (Breiman, 2001), gradient boosting (Friedman, 2001), and XGBoost (Chen & Guestrin, 2016) for the Branin function. For Figures 1 and 9, we use $\zeta = 0.33$ and $\beta = 0.5$.

 А

B DETAILS OF DRE-BO-SSL

Algorithm 2 Labeling Unlabeled Data

- 814 Input: Labeled data points X_l , their labels C_l , unlabeled data points X_u , maximum iterations τ , 815 and a tolerance ε . Additionally, a clamping factor α for label spreading.
- 816 **Output:** Propagated labels $\widehat{\mathbf{C}}$.
 - 1: Initialize propagated labels $\widehat{\mathbf{C}}$ of \mathbf{X} .
 - 2: Compute similarities W and a degree matrix D.
- 3: Compute transition probabilities **P** with **W** and **D**.
 - 4: repeat

812

813

817

818

820

821

822

823 824

826

827 828

829 830

835 836 837

- 5: Propagate $\widehat{\mathbf{C}}$ with \mathbf{P} and the previous $\widehat{\mathbf{C}}$, and additionally α for label spreading.
- 6: Normalize $\widehat{\mathbf{C}}$ row-wise.
 - 7: **until** a change of C converging to ε or reaching τ .

Algorithm 2 describes a procedure to label unlabeled data points; see the main article for the details of DRE-BO-SSL.

C ANALYSIS OF DRE-BO-SSL

Under the cluster assumption, i.e., Assumption 4.1, a margin γ is defined as a minimum distance between two decision boundaries.

Definition C.1. Let a compact connected decision set be $C_i \subseteq \mathcal{X}$ for $i \in \{0, 1\}$ and a boundary subset, i.e., a set of boundary points, of a compact connected set S be ∂S . A margin γ is defined as

$$\gamma = (2\mathbb{I}(\mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset) - 1) \min_{\mathbf{x}_1 \in \partial \mathcal{C}_1 \setminus \partial \mathcal{X}, \mathbf{x}_2 \in \partial \mathcal{C}_2 \setminus \partial \mathcal{X}} \|\mathbf{x}_1 - \mathbf{x}_2\|_2.$$
(11)

Using Definition C.1, we claim that a semi-supervised classifier in DRE-BO-SSL can mitigate the 838 over-exploitation problem presented in Section 1.1, because it can expand a decision set C_1 for Class 839 1 by reducing γ with unlabeled data. However, we need to verify if a large decision set is derived 840 from the characteristics of non-parametric classifiers, since a semi-supervised classifier we use is 841 converted to the Nadaraya-Watson non-parametric model (Nadaraya, 1964; Watson, 1964) without 842 unlabeled data. As shown in Figure 12, there is no strong relationship between the performance of 843 semi-supervised classifiers without unlabeled data, i.e., the Nadaraya-Watson estimator, and one of 844 semi-supervised classifiers with unlabeled data. It implies that configuration selection for a semi-845 supervised classifier is dependent on a class of objective function, and the presence of unlabeled 846 data is likely to be effective for alleviating the over-exploitation problem.

847 Singh et al. (2008) provide a sample error bound of supervised and semi-supervised learners, re-848 lated to γ , n_l , n_u , and d. This work proves that a semi-supervised learner can be better than any 849 supervised learners, by assuming that $n_u \gg n_l$ and access to perfect knowledge of decision sets. 850 However, these theoretical results cannot be directly applied in our sequential problem because this 851 work assumes that both labeled and unlabeled data points are independent and identically distributed. 852 Nevertheless, these results can hint a theoretical guarantee on better performance of semi-supervised classifiers with unlabeled data points. Further analysis for the circumstances of Bayesian optimiza-853 tion is left for future work. 854

- 855
- 856
- 00/ 858
- 859
- 200
- 861
- 862
- 863

864 D **EXPERIMENT DETAILS** 865

866 Here we present the missing details of the experiments shown in the main part. 867

To carry out the experiments in our work, we use dozens of commercial Intel and AMD CPUs 868 such as Intel Xeon Gold 6126 and AMD EPYC 7302. For the experiments on minimum multi-digit MNIST search, the NVIDIA GeForce RTX 3090 GPU is used. 870

To minimize (6) for finding an adequate β of label propagation and label spreading, we use L-871 872 BFGS-B with a single initialization, which is implemented in SciPy (Virtanen et al., 2020), for all the experiments in Figures 2, 3, 5, 6, and 7. For the other empirical analyses, we set β as a specific 873 fixed value; see the corresponding sections for the details. 874

875 To reduce the computational complexity of DRE-BO-SSL for a scenario with a fixed-size pool, we 876 randomly sample 2,000 unlabeled points from the predefined pool for all experiments excluding synthetic benchmark functions. More thorough analysis can be found in Section G. 877

878 To compare baseline methods with our methods, we assess optimization algorithms using a simple 879 regret: 880

simple regret
$$(f(\mathbf{x}_1), \dots, f(\mathbf{x}_t), f(\mathbf{x}^*)) = \min_{i=1}^t f(\mathbf{x}_i) - f(\mathbf{x}^*),$$
 (12)

882 where \mathbf{x}^* is a global optimum.

883

884

885

889 890

891

892 893

894

895

896

897

899

900

901

902

903

904

905

906

909

910

915

917

Our proposed methods and baseline methods are implemented with scikit-learn (Pedregosa et al., 2011), PyTorch (Paszke et al., 2019), NumPy (Harris et al., 2020), SciPy (Virtanen et al., 2020), XGBoost (Chen & Guestrin, 2016), and BayesO (Kim & Choi, 2023). Scikit-learn and SciPy are under the BSD 3-Clause license, NumPy is under the modified BSD license, XGBoost is under the 887 Apache 2.0 license, and BayesO is under the MIT license. On the other hand, PyTorch is under its own license; refer to its repository for the license.

D.1 DETAILS OF BASELINE METHODS

As the competitors of our method, we test the following baseline methods:

- Gaussian process, EI and UCB: It is a Bayesian optimization strategy, which is defined with GP regression with the Matérn 5/2 kernel (Rasmussen & Williams, 2006). Expected improvement (Jones et al., 1998) and GP-UCB (Srinivas et al., 2010) are used as acquisition functions.
- Random forest, BORE and LFBO: These are BORE and LFBO that employ random forests (Breiman, 2001) with 1000 decision trees. Minimum samples to split are set to 2 for these baselines.
- Gradient boosting, BORE and LFBO: These methods are BORE and LFBO with gradient boosting classifiers (Friedman, 2001) with 100 decision trees. Learning rate for the classifier is set to 0.3.
 - XGBoost, BORE and LFBO: Similar to gradient boosting, BORE and LFBO with XGBoost have 100 decision trees as base learners with learning rate of 0.3.
 - MLP, BORE and LFBO: These methods are built with two-layer fully-connected networks. The detail of the multi-layer perceptron is described as follows.
- 907 The architecture of a multi-layer perceptron is 908
 - First layer: fully-connected, input dimensionality d, output dimensionality 32, ReLU; Second layer: fully-connected, input dimensionality 32, output dimensionality 1, Logistic,
- 911 where d is the dimensionality of the problem we solve. 912

913 Note that most configurations for the baseline methods follow the configurations described in the 914 work (Song et al., 2022).

916 D.2 DETAILS OF SYNTHETIC BENCHMARKS

Here we describe the details of synthetic benchmarks.

918 Beale Function. This function is defined as follows: 919 $f(\mathbf{x}) = (1.5 - x_1 + x_1 x_2)^2 + (2.25 - x_1 + x_1 x_2^2)^2 + (2.625 - x_1 + x_1 x_2^3)^2,$ (13)920 where $\mathbf{x} = [x_1, x_2] \in [-4.5, 4.5]^2$. 921 922 **Branin Function.** It is defined as follows: 923 924 $f(\mathbf{x}) = \left(x_2 - (5.1/4\pi^2)x_1^2 + (5/\pi)x_1 - 6\right)^2 + 10\left(1 - (1/8\pi)\right)\cos(x_1) + 10,$ (14)925 where $\mathbf{x} = [x_1, x_2] \in [[-5, 10], [0, 15]].$ 926 927 Bukin6 Function. This benchmark is given by the following: 928 929 $f(\mathbf{x}) = 100\sqrt{|x_2 - 0.01x_1^2|} + 0.01|x_1 + 10|,$ (15)930 931 where $\mathbf{x} = [x_1, x_2] \in [[-15, -5], [-3, 3]].$ 932 933 **Six-Hump Camel Function.** It is given by the following: 934 $f(\mathbf{x}) = (4 - 2.1x_1^2 + x_1^4/3)x_1^2 + x_1x_2 + (-4 + 4x_2^2)x_2^2.$ (16)935 936 where $\mathbf{x} = [x_1, x_2] \in [[-3, 3], [-2, 2]].$ 937 938 D.3 DETAILS OF TABULAR BENCHMARKS 939 Table 1: Search space for Tabular Benchmarks. "tanh" and "relu" represent hyperbolic tangent and 940 941 ReLU, respectively. 942 Hyperparameter **Possible Values** 943 944 The number of units for 1st layer $\{16, 32, 64, 128, 256, 512\}$ 945 {16, 32, 64, 128, 256, 512} The number of units for 2nd layer Dropout rate for 1st layer $\{0.0, 0.3, 0.6\}$ 946 $\{0.0, 0.3, 0.6\}$ Dropout rate for 2nd layer 947 $\{ \text{``tanh'', ``relu''} \\ \{ \text{``tanh'', ``relu''} \\ \{ \text{``tanh'', ``relu''} \\ \{ 5 \times 10^{-4}, 1 \times 10^{-3}, 5 \times 10^{-3}, 1 \times 10^{-2}, 5 \times 10^{-2}, 1 \times 10^{-1} \\ \{ \text{``cosine'', ``constant''} \}$ Activation function for 1st layer 948

Batch size $\{8, 16, 32, 64\}$ The search space of Tabular Benchmarks (Klein & Hutter, 2019) is described in Table 1. To handle categorical and discrete variables, we treat them as integer variables by following the previous

D.4 DETAILS OF NATS-BENCH

Activation function for 2nd layer

Initial learning rate

Learning rate scheduling

949

950

951 952 953

954

955

956 957

958 959

960 961 962

963 964

965 966

967





We describe the search space for NATS-Bench (Dong et al., 2021) in Figure 10 and Table 2.

D.5 DETAILS OF MINIMUM MULTI-DIGIT MNIST SEARCH

literature (Garrido-Merchán & Hernández-Lobato, 2020).

Since multi-digit MNIST, which is composed of images of size (28, 84), is high-dimensional, some 968 969 of the methods used in this work, e.g., methods with random forests, gradient boosting, and XG-Boost, struggle to process such data. Therefore, we embed an original image to a lower-dimensional 970 vector using an auxiliary convolutional neural network. The convolutional neural network is trained 971 to classify a three-digit image to one of labels from "000" to "999," with the following architecture:

Table 2: Search space for NATS-Bench. There exist $8^5 = 32,768$ models.

| 973 | | | | |
|------------|---|---|---|--|
| 974 | | Hyperparameter | Possible Values | |
| 975 | | Output channels of 1st convolutional laver | {8, 16, 24, 32, 40, 48, 56, 64} | |
| 976 | | Output channels of 1st cell stage | {8, 16, 24, 32, 40, 48, 56, 64} | |
| 977 | | Output channels of 1st residual block | $\{8, 16, 24, 32, 40, 48, 56, 64\}$ | |
| 978 | | Output channels of 2nd cell stage | $\{8, 16, 24, 32, 40, 48, 56, 64\}$ | |
| 979 | | Output channels of 2nd residual block | {8, 16, 24, 32, 40, 48, 56, 64} | |
| 980 | | | | |
| 981 | First law | er: convolutional input channel 1 output | at channel 8 kernel size 3×3 modding 1 | |
| 982 983 | ReLU, n | hax-pooling 2×2 ; | at channel 6, kenter size 5×5 , padding 1, | |
| 984 985 | Second layer: convolutional, input channel 8, output channel 16, kernel size 3×3 , padding 1, ReLU, max-pooling 2×2 ; | | | |
| 986 987 | Third layer: convolutional, input channel 16, output channel 32, kernel size 3×3 , padding 1, ReLU, max-pooling 2×2 ; | | | |
| 988 | Fourth la | iver: fully-connected, input dimensionalit | v 960. output dimensionality 128. ReLU: | |
| 989 | Fifth layer fully connected input dimensionality 100, output dimensionality 64, Del U. | | | |
| 990 | Fifth layer: fully-connected, input dimensionality 128, output dimensionality 64, ReLU; | | | |
| 991 | Sixth lay | er: fully-connected, input dimensionality | 64, output dimensionality 1000, Softmax. | |
| 992 | The Adam on | timizer (Kingma & Ba_2015) with learning | ng rate 1×10^{-3} is used to train the network | |
| 993 | for 100 epoch | s. To train and test the model fairly, we cr | reate a training dataset of 440,000 three-digit | |
| 994 | images, a val | idation dataset of 40,000 three-digit image | ges, and a test dataset of 80,000 three-digit | |
| 995 | images using | a training dataset of 55,000 single-digit i | mages, a validation dataset of 5,000 single- | |
| 996 | digit images, | and a test dataset of 10,000 single-digit in | nages in the original MNIST dataset (LeCun | |
| 997 | et al., 1998). | For example, supposing that a test dataset | has 1,000 single-digit images per class – it is | |
| 990 | three-digit im | age "753" 10^9 combinations for "753" ca | n be created. We therefore randomly sample | |
| 1000 | a fixed numb | er of three-digit images from a vast num | ber of possible combinations. In addition, | |
| 1001 | an early stop | bing technique is utilized by comparing t | the current validation loss to the average of | |
| 1001 | validation los | ses for the recent five epochs. Eventually, | our network achieves 99.6% in the training | |
| 1002 | dataset, 97.0% | % in the validation dataset, and $96.9%$ in t | he test dataset. | |
| 1004 | To construct a | a fixed-size pool, we use 80,000 embeddi | ngs of dimensionality 64, which are derived | |
| 1005 | from the outp | outs of the fifth layer without ReLU, by p | assing the test dataset of three-digit images | |
| 1006 | through the n | etwork. | | |
| 1007 | | | | |
| 1008 | | | | |
| 1009 | | | | |
| 1010 | | | | |
| 1011 | | | | |
| 1012 | | | | |
| 1013 | | | | |
| 1014 | | | | |
| 1015 | | | | |
| 1016 | | | | |
| 1017 | | | | |
| 1018 | | | | |
| 1019 | | | | |
| 1020 | | | | |
| 1021 | | | | |
| 1022 | | | | |
| 1023 | | | | |
| 1024 | | | | |
| 1025 | | | | |



1026EDISCUSSION ON A FREE PARAMETER IN LABEL PROPAGATION AND1027LABEL SPREADING

Figure 11: Effects of a free parameter β in label propagation, denoted as LP, and label spreading, denoted as LS. All experiments are repeated 20 times.

In the Bayesian optimization process of DRE-BO-SSL, a free parameter β in label propagation and label spreading is learned every iteration by minimizing (6); see Figure 4 for the results on learned β . Furthermore, to show the effects of β , we empirically analyze β as depicted in Figure 11. We sample 1,000 unlabeled points and use all of them as unlabeled points without pool sampling. For the cases of four benchmark functions, higher β tends to show better performance than lower β . These results considerably correspond with the results in Figure 4.

F DISCUSSION ON UNLABELED POINT SAMPLING







Figure 13: Effects of sampling strategies for unlabeled point sampling. LP and LS stand for label propagation and label spreading, respectively. We repeat all experiments 20 times.

1101 We design two studies to analyze the effects of the number of unlabeled points n_u and sampling 1102 strategies in a process of unlabeled point sampling, where unlabeled points are not provided and 1103 $\beta = 0.5$ is given.

For the first study, we conduct five settings, no unlabeled data, which implies that transductive learning is not used, and $n_u = 10,100,1000$. Interestingly, the tendency of the number of unlabeled points are unclear as presented in Figure 12. It implies that a setting for the number of unlabeled data points depend on the characteristics of benchmark functions, which is common in Bayesian optimization and black-box optimization. Besides, γ is different across benchmarks and iterations and it lets optimization results sensitive to n_u . Therefore, we cannot determine a suitable setting without access to a black-box function of interest.

1111 As another elaborate study, we test the effects of sampling strategies. Four strategies, the truncated multivariate normal distributions, uniform distributions, Halton sequences (Halton, 1960), and Sobol' sequences (Sobol', 1967), are compared. As depicted in Figure 13, the normal distribution is better than the other sampling methods in four cases and shows robust performance in most of the cases, but it is not always the best. Similar to the previous study on the effects of n_u , we presume that it is also affected by γ , which is hard to define in practice.

1117

1100

1118 G DISCUSSION ON POOL SAMPLING

To see the impact of an additional hyperparameter, i.e., the size of a subset of the original pool, which is introduced to speed up semi-supervised learning algorithms, we demonstrate numerical analysis on pool sampling where the size of a predefined pool is 4,000 and $\beta = 0.5$ is given. Based on Figures 8 and 14, we can accelerate our framework without significant performance loss.

1124 1125 1126

H DISCUSSION ON A THRESHOLD RATIO

Figure 15 demonstrates the effects of a threshold ratio ζ , where we use $\beta = 0.5$. It follows the setting described in Section E. As presented in Figure 15, a smaller ζ , for example, $\zeta = 0.01$ or $\zeta = 0.1$ tends to show worse performance than a larger ζ , which implies that the over-exploitation is not due to conservative y^{\dagger} . Interestingly, the results with $\zeta = 0.8$ also generally under-perform. We presume that it is basically due to over-exploration.



Figure 15: Effects of a threshold ratio ζ with $\beta = 0.5$. We repeat all experiments 20 times, and LP and LS stand for label propagation and label spreading, respectively.

¹¹⁸⁸ I LIMITATIONS

As discussed above, our algorithms slow down if a pool size is significantly large. As presented in Figure 8, elapsed times are certainly dependent on subset sizes. To tackle this issue, we suggest a method to randomly select a subset of the pool, but a more sophisticated subsampling method can be devised for our framework. In particular, we can leverage the impacts of the subset of the pool by utilizing the geometric information of unlabeled data points.

1196 J IMPACT STATEMENT

Our work does not have any direct negative societal impacts because it proposes a new Bayesian optimization method to optimize a black-box function. However, in the sense of optimization, this line of research can be used in optimizing any objectives including unethical optimization tasks.