

Neighbor Regularized Bayesian Optimization for Hyperparameter Optimization

Lei Cui¹

cuil19@mails.tsinghua.edu.cn

Yangguang Li²

liyangguang@sensetime.com

Xin Lu²

luxin@sensetime.com

Dong An³

dong.an@cripac.ia.ac.cn

Fenggang Liu²

liufgtech@163.com

¹ Tsinghua University

Beijing, China

² SenseTime Research

Beijing, China

³ Institute of Automation, Chinese

Academy of Sciences

Beijing, China

Abstract

Bayesian Optimization (BO) is a common solution to search optimal hyperparameters based on sample observations of a machine learning model. Existing BO algorithms could converge slowly even collapse when the potential observation noise misdirects the optimization. In this paper, we propose a novel BO algorithm called *Neighbor Regularized Bayesian Optimization* (NRBO) to solve the problem. We first propose a neighbor-based regularization to smooth each sample observation, which could reduce the observation noise efficiently without any extra training cost. Since the neighbor regularization highly depends on the sample density of a neighbor area, we further design a density-based acquisition function to adjust the acquisition reward and obtain more stable statistics. In addition, we design an adjustment mechanism to ensure the framework maintains a reasonable regularization strength and density reward conditioned on remaining computation resources. We conduct experiments on the bayesmark benchmark and important computer vision benchmarks such as ImageNet and COCO. Extensive experiments demonstrate the effectiveness of NRBO and it consistently outperforms other state-of-the-art methods.

1 Introduction

The performance of modern machine learning models highly depends on the good choice of hyperparameters. Thus a system that can quickly and automatically optimize the hyperparameters becomes more and more important nowadays. The hyperparameter optimization problem is often regarded as a black-box optimization problem, and the common solution is *Bayesian optimization* (BO) [5, 6, 8, 12, 24, 33, 34].

Though existing BO algorithms have shown significant improvement compared with the random search algorithm [12, 24], BO could still converge slowly even collapse when the

potential observation noise is non-negligible. Specifically, the observation (a.k.a. performance) of a model could fluctuate in a wide range even using the same hyperparameters, which might mislead the surrogate model to overfit noisy observations. A natural idea to solve the problem is to obtain reliable observations by training repetitively. However, it will cause multiplied computational cost which is unacceptable in heavy tasks. Another idea is to add noise assumptions on the sample observations, while the noise assumptions without any prior will decrease the fitting efficiency of the surrogate model.

To overcome the above difficulties, we propose a novel BO algorithm called *Neighbor Regularized Bayesian Optimization* (NRBO). NRBO applies adaptive regularization at different positions of the search space according to the neighbor sample points. Specifically, the observation of a sample will be smoothed by the observations of its neighbor samples. Then, the smoothed observations are used to optimize the surrogate model. The regularization could reduce the observation noise *which has the same spirit as the k-nearest neighbors algorithm* [14]. Thus, no repetitive training phases are needed to get extra observations. Compared with directly introducing a global noise assumption, this method can dynamically smooth observations at different regions thus regularize the optimization.

The neighbor-based regularization algorithm needs to have as many samples as possible at each desired area to obtain credible and stable statistics. Therefore, we design a new density-based acquisition function, which can adjust the acquisition reward according to the sample density of a certain area. The proposed acquisition function adopts a stronger reward for the *sparse* area where the density of observed samples is low. This design has two advantages. First, it provides a more balanced sampling strategy that the observed samples are distributed more evenly on the entire search space. Second, it prevents the surrogate model from sticking into a local sub-optimal solution prematurely.

Finally, we design a mechanism to dynamically adjust the regularization strength and acquisition reward. This mechanism ensures that the framework maintains a reasonable regularization strength and acquisition reward during the Bayesian optimization process. The regularization strength and density reward are positively correlated with the remaining computation resources. As the search progresses, the regularization strength will be weakened to help the model fit more subtle data patterns in finer spatial resolution. At the same time, the density reward will also be weakened to balance the exploration-exploitation trade-off and encourage the algorithm to converge at the high-performance area when searching is about to end.

Extensive experiments demonstrate the effectiveness of our NRBO. NRBO achieves the state-of-the-art performance on the bayesmark and six commonly used computer vision tasks. For example, NRBO achieves 0.85%, 0.22%, 2.1% and 2.36% higher accuracy on Stanfordcar, ImageNet, VOC and COCO datasets compared with random search [14]. We summarize our main contributions in this paper as follows:

- We propose a novel Bayesian optimization method based on the statistics of the neighbor observations, which can adaptively reduce the observation noise and regularize the surrogate model.
- We propose a novel density-based acquisition function to provide credible and stable statistics for Bayesian optimization process.
- We design a mechanism to dynamically adjust the regularization strength and acquisition reward according to the remaining computation resource. It helps to fit more subtle data patterns and balance the exploration-exploitation trade-off as the optimization progresses.

2 Related Works

2.1 Model-free Method

The most straightforward approach to optimize the hyperparameters is the model-free method. Grid search is one of the commonly used methods that can be easily implemented and parallelized [19, 20]. It discretizes the search space into a mesh grid and evaluates them all. The computation cost will explode exponentially as the dimension or the resolution of the hyperparameter increases. Another model-free approach is the Random search [2, 30]. Instead of traversing the discretized search space, random search selects the candidate hyperparameter randomly. Although the Random Search and the Grid Search seems to have similar efficiency, Random Search usually performs better in a limited search budget in practice. The reason is that the model performance is usually not distributed uniformly in the entire search space. The Random Search samples a fixed number of parameter combinations from the specified distribution, which improves system efficiency by reducing the probability of wasting much time on a small poorly-performing region [35].

2.2 Bayesian Optimization

In addition to the model-free approaches, the most popular method used in hyperparameter optimization is Bayesian optimization [6, 8, 12, 17, 19, 31, 33, 34]. Instead of searching the hyperparameter in a pre-defined distribution, Bayesian optimization can dynamically fit the observed data to determine the next search place. Bayesian optimization algorithm runs a search loop that iteratively fits the surrogate model and queries the acquisition function [28]. By fitting the observed data, the surrogate model can predict the performance distribution throughout the search space with prior distribution such as smoothness. The acquisition function selects the promising regions that balance the exploration-exploitation trade-off [13]. A typical Bayesian optimization algorithm is the Gaussian Process Bayesian Optimization (GPBO) [26]. GPBO utilizes the Gaussian process as the surrogate model that can predict the performance in the search space as a jointly Gaussian distribution. SMAC is another popular Bayesian optimization algorithm that uses random forest as the surrogate model and ensembled regression trees as the objective function [15]. HEBO enhances the surrogate model through input warping and output transformations and proposes a multi-objective acquisition function for candidate selection [4].

3 Method

3.1 Preliminaries

Hyperparameter optimization can be seen as a black-box optimization problem. The searching process is started with a initial observation dataset \mathcal{D}_0 . Specifically, a batch of random hyperparameters $x_{0:m-1}$ are sampled at first. Then models with these hyperparameters are trained and evaluated to get the observations $y_{0:m-1}$. The observation dataset \mathcal{D}_0 is initialized by $\{(x_{0:m-1}, y_{0:m-1})\}$. The surrogate model f_θ will be trained for N times to fit the observation dataset \mathcal{D}_0 , where θ is optimized by minimizing the negative log marginal likelihood.

After fitting the surrogate model, we need to guess the best x^* by proposing a new (batch of) sample point according to acquisition function f_{acq} . Given an input sample point x ,

f_{acq} calculates the priority of this sample point according to probability distribution predicted by the surrogate model f_{θ} . The acquisition function is usually designed to balance the exploration-exploitation trade-off. A novel sample point \hat{x} is selected by $\operatorname{argmax}_{x \in \mathcal{X}} f_{acq}(x)$, and in practice we discretize the continuous search space \mathcal{X} into a meshgrid points collection.

Once the novel sample point is selected, it will be evaluated on black-box by training the machine learning models with the hyperparameter it represents. Then the observation $\{(\hat{x}, \hat{y})\}$ will be added to datasets \mathcal{D}_i for the next iteration. When the main loop is over, the algorithm returns the best-performing data point as the result.

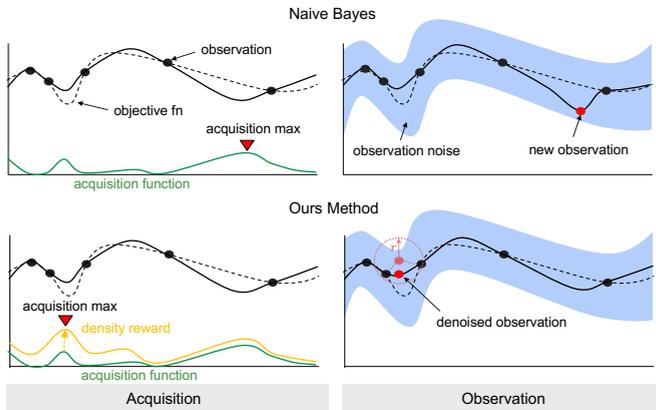


Figure 1: Illustration of the NRBO algorithm. In two key progress of acquisition and observation. Our method is different from the naive bayes algorithm. In acquisition stage, we propose a density-based acquisition function to accelerate the acquisition process, in which adjacent sample points in the neighbor are considered. In observation stage, neighbor regularized mechanism is introduced to smooth the observation noise and release the burden of repetitive observation.

3.2 Neighbor Regularized Bayesian Optimization

Figure 1 shows a brief illustration of our NRBO. The main innovations of NRBO can be summarized as: a novel regularization to reduce the observation noise (section 3.2.1), a density-based acquisition function to adjust the acquisition reward (section 3.2.2) and a dynamic adjustment of the regularization strength and acquisition reward according to the remaining computation resource (section 3.2.3).

3.2.1 Neighbor-based regularization

In the context of Bayesian optimization, the algorithm maintains an observation dataset \mathcal{D}_i at any iteration i , which consists of all sampled hyperparameters as well as their corresponding observation received from the black box. To regularize the surrogate model, we define

$$\bar{y}_j = \frac{\sum_k^{N_D} P_j(x_k) \cdot y_k}{\sum_k^{N_D} P_j(x_k)}, \quad P_j(x_k) = \begin{cases} 1 & \|x_j - x_k\| \leq \sigma_1 \\ 0 & \|x_j - x_k\| > \sigma_1 \end{cases} \quad (1)$$

where N_D is the number of samples in \mathcal{D}_i , $P_j(x_k)$ is a filter function that keeps the points in the neighbor of x_j with σ_1 as radius. \bar{y}_j is the observation smoothed by its neighbor sample observations, and we have a smoothed observation dataset $\bar{\mathcal{D}}_i = \{(x_l, \bar{y}_l)\}_{l=0:N-1}$. Different from naive Bayes, the surrogate model is regularized by fitting the $\bar{\mathcal{D}}_i$ instead of \mathcal{D}_i .

Here we give a brief analysis on neighbor-based regularization. Firstly, NRBO smoothes each observation by taking the neighbor observations into account. In this way, NRBO can efficiently regularize the model without repetitive training. In practice, the computation burden of repetitive training is unacceptable because each observation needs to re-train the entire model. Secondly, the strength of regularization is determined by the statistical properties of all sample points lie in the neighbor area. As the variance of observations increase, NRBO will also impose stronger regularization. Compared with simply modeling noise with $y = f(x) + \eta$ that regularize the model evenly throughout the entire search space, the NRBO regularization mechanism is more adaptive and effective. Note that in practice, σ_1 moves according to the searching progress, we will discuss it in section 3.2.3.

3.2.2 Density-based acquisition function

As demonstrated in section 3.2.1, NRBO imposes adaptive regularization by adopting neighbor sample observations. Obviously, to make the NRBO works more efficiently, we need to increase the density of the sample points for a more reliable and stable statistics. If there exists no other observed sample point lies in the neighbor area of a certain point, NRBO degenerates into a normal Bayesian optimization without any regularization in this area.

Based on the above analysis, we propose a density-based strategy to adjust the acquisition function. Our baseline acquisition function is the multi-objective acquisition ensemble (MACE) proposed in HEBO [14]. The proposed density-based acquisition function can be formally denoted as:

$$\min(-f_{EI} - g_d(x)S_{EI}; -f_{PI} - g_d(x)S_{PI}; f_{UCB} - g_d(x)S_{UCB}), \quad g_d(x) = e^{-f_n(x, \sigma_2)} \quad (2)$$

where f_{EI} , f_{PI} and f_{UCB} are three widely-used myopic acquisition functions. S_{EI} , S_{PI} and S_{UCB} are norm items that represent the standard deviation of acquisition values of all candidate points in search space. We follow HEBO [14] to define and use these symbols, please refer to HEBO for more details. Different from HEBO, we introduce new adjustment items g_d . In equation 2, function f_n returns the number of sample points in D_i that lies in the neighbor area of point x . σ_2 is the radius that defines the neighbor range.

With adjustment item g_d , the acquisition value of the candidate point x can be adjusted by the density of its observed adjacent sample points in D_i , thus g_d can be regarded as density reward. It will highly pump the acquisition value in the area that has fewer observed points in dataset D_i , and encourage the solver to search this sparse area.

Equipped with the density-based adjustment strategy, observed sample points in D_i will distribute more evenly to enhance the efficiency of the neighbor-based regularization. Note that in practice, σ_2 also moves according to the searching progress and we will discuss it in section 3.2.3.

3.2.3 Dynamic regularization strength and density reward

In Eq 1, we adopt a fixed σ_1 to smooth the observations and regularize the surrogate model. In practice, we use a dynamic σ_1 to weaken the strength of the regularization as the search progresses.

At the beginning of the search process, we adopt a larger σ_1 to strengthen the smoothness. This helps the model to fit the observation dataset with lower resolution and concentrate more on long-range trends in the entire search space. As the search progresses, the σ_1 will gradually decrease to weaken the regularization, encouraging the model to fit more details when the solver is more close to the optimal solution. Specifically, we can rewrite the filter function in equation 1 as follows:

$$P_j(x_k) = \begin{cases} 1 & \|x_j - x_k\| \leq \sigma_1(i) \\ 0 & \|x_j - x_k\| > \sigma_1(i) \end{cases} \quad \sigma_1(i) = \sigma_1^0 + (1 - i/N) * \sigma_1^1 \quad (3)$$

where σ_1^0 is the base regularization strength, i is the current optimization iteration and N is the total number of the optimization iteration. The neighbor radius will starts at $\sigma_1^0 + \sigma_1^1$ and end up in σ_1^0 .

Similarly, we also adopt a dynamic σ_2 in Eq 2, the difference is that it moves in an opposite direction. At the beginning of the search process, we use a relatively smaller σ_2 to encourage the algorithm to search the entire search space especially the sparse area that contains fewer observed points. Then the σ_2 will gradually decrease the density reward and rebalance the exploration-exploitation trade-off. Formally, we re-write the equation 2 as

$$g_d(x) = e^{-f_n(x, \sigma_2(i))}, \quad \sigma_2(i) = \sigma_2^0 + i/N * \sigma_2^1 \quad (4)$$

It helps the search process to get rid of the local optimal at very beginning and finally rebalance the exploration and exploitation to the normal state because the adjustment item value g_d will finally tend to be zero throughout the entire search space, making the searching process easier to converge on the optimal point.

4 Experiments

We conduct extensive experiments on both conventional BO benchmark bayesmark and some commonly used computer vision benchmarks such as ImageNet [18], COCO [2]. We first compare NRBO against several state-of-the-arts BO methods on both bayesmark (section 4.1) and computer vision benchmarks (section 4.2). Then we conduct a more detailed ablation study to demonstrate the impact of each component of NRBO in section 4.4. In addition, we present some qualitative visualization of NRBO in section 4.5.

4.1 Results on Bayesmark

Bayesmark¹ benchmark contains 6 standard datasets (breast, digits, iris, wine, boston, diabetes) and 9 commonly used machine learning models (DT, MLP-ADAM, MLP-SGD, RF, SVM, ADA, KNN, Lasso, Linear). Each model has two variants for classification and regression. Combining these datasets and models, there are 108 tasks to evaluate the hyperparameter optimization algorithms.

We adopt two official metrics to measure the performance on Bayesmark. The loss score metric for each task is calculated by $100 * (1 - loss)$. The normalized mean score first calculates the performance gap between observations and the global optimal point, then normalize it by the gap between random search results and the optimal. Table 1 shows the performance

¹<https://github.com/uber/bayesmark>

Method	MLP_digits	RF_breast	SVM_digits	SVM_wine	ada_breast	ada_digits	linear_breast	Avg._score	Avg._norm
Random	98.63	92.07	94.95	81.19	93.75	69.26	89.27	90.93	1.07
Hyperopt	100.29	96.50	96.87	88.72	98.96	73.87	92.85	96.53	0.411
Opentuner	100.76	92.07	98.43	92.61	98.96	100.21	94.39	93.16	0.81
Nevergrad	100.11	98.83	100.07	70.30	98.96	76.09	94.55	93.95	0.72
Pysot	101.13	95.34	97.56	99.87	95.83	96.55	97.80	98.34	0.20
Skopt	100.08	98.60	96.86	96.24	93.75	76.79	89.88	96.23	0.45
Turbo	101.10	96.74	98.53	100	95.83	91.45	97.56	98.28	0.20
HEBO	101.38	97.67	101.70	96.24	98.96	112.88	95.82	99.94	0.01
NRBO	101.49	100.00	103.53	103.63	100.00	113.01	101.01	100.34	-0.04

Table 1: The performance of 9 optimizers on bayesmark benchmark. We list the average loss score(denoted as Avg.score) and normalized mean score(denoted as Avg.norm) for all 108 tasks in bayesmark benchmark.

Method	ImageNet	VOC	CIFAR10	CIFAR100	Stanford Car	COCO
Random	62.00	75.02	95.28	81.96	87.09	31.56
HEBO	62.07(+0.07)	74.88(-0.14)	95.35(+0.07)	81.91(-0.05)	87.56(+0.47)	32.34(+0.78)
NRBO	62.22(+0.22)	77.12(+2.10)	95.43(+0.15)	82.16(+0.20)	87.94(+0.85)	33.92(+1.36)

Table 2: Experiment results on computer vision tasks. The positive numbers in parentheses represent absolute improvements relative to Random Search.

of NRBO and other 8 commonly used hyperparameter optimizers, including random search, hyperopt [3], opentuner [4], pysot [5], skopt [32], turbo [8], nevergrad [25] and HEBO [9]. We also select 7 tasks in bayesmark and plot the optimization process in detail at figure 2. NRBO starts to surpass other optimizers at the 9-th iteration and stays ahead till the end of the optimization. The results indicate that NRBO converges faster and reaches better final score.

4.2 Results on Computer Vision Tasks

In this section, we show the experiment results about different hyperparameter optimization algorithms used on common computer vision tasks. We only choose three optimization algorithms including random search, HEBO and our NRBO due to the high cost of computer vision tasks on both time and resource.

Our computer vision experiments include classification task and detection task. For the classification task, we use ResNet-18 [14] and train it on ImageNet [18], CIFAR10 [17], CIFAR100 [17] and Stanford Cars [16] datasets. Four hyperparameters including momentum, weight decay, label smooth and learning rate compose a 4-dimension searching space. For the detection task, we use RetinaNet [23] and train it on Pascal VOC [9] and MS-COCO2017 [22] datasets. The searching space is also composed of 4 hyperparameters including momentum, weight decay, positive IoU thresh and negative IoU thresh.

The best detection results with different hyperparameter optimization algorithms are shown in table 2, where HEBO performs better than random search on COCO, and NRBO performs best on both VOC and COCO datasets. These results are consistent with those in bayesmark benchmark, demonstrating the effectiveness of our proposed method.

The classification results are shown in table 2. It is seen that the performance of HEBO is close to the naive random search. This may be caused by a relatively flat performance

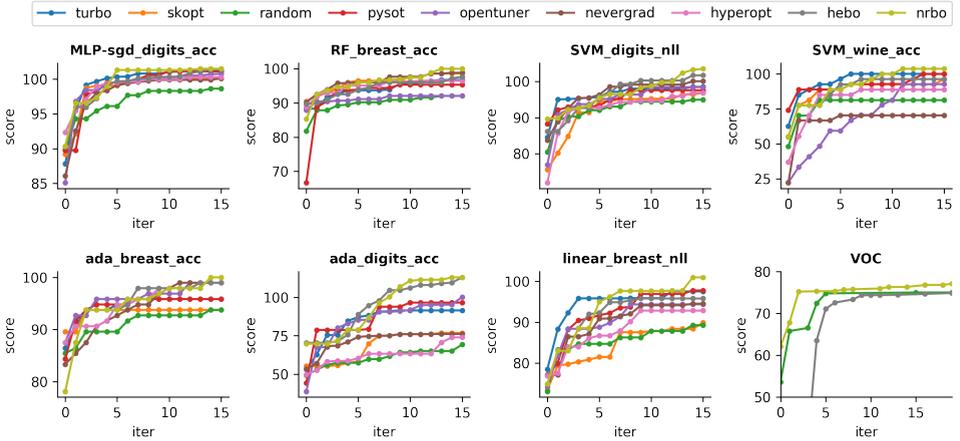


Figure 2: The convergence progress of different methods. The results indicate that NRBO converges faster and reaches better score on both bayesmarks and CV tasks. For VOC task, score represents the AP50 metric.

landscape near the optimal solution in search space, and even a random search could have a high probability of getting a good result. Compared to bayesmark and COCO, the output of the ImageNet experiments with a specific hyperparameters is more stable. Therefore, the NRBO could hardly benefit from the regularization. Even in this case, NRBO still surpasses the HEBO on all four classification tasks.

4.3 Full-training Results of Computer Vision Tasks

Since the hyperparameter optimization needs repetitive trials, the computation cost is unaffordable on large-scale datasets. To overcome the computation resource limitation, in large computer vision datasets ImageNet and COCO, we optimize the hyperparameter on proxy tasks with 60% of total iterations for only 20 trials. Then we use the hyperparameters optimized on the *proxy* task for a *full* training to verify whether the NRBO is still in the lead. We repeat each optimizer 4 times for stable and convincing results.

Classification The configuration is same as ImageNet experiments in section 4.2, except that the learning rate of proxy task decreases at 25000th, 50000th, and 75000th iteration and ends at 80000th iteration. In full-training task, the learning rate schedule time is set as 37500, 75000, and 112500. Training ends at the 125000th iteration. As shown in table 3, the average accuracy of NRBO surpasses HEBO and random search by 0.37 and 0.33 respectively in the full-training setting.

Detection The configuration is same as COCO experiments in section 4.2. In full-training task, the learning rate drops at the 9th and 12th epoch. Training ends at the 14th epoch. As shown in table 3, the mAP of NRBO surpasses HEBO and random search by 2.29 and 1.37 respectively in the full-training setting.

Experiments above demonstrate that the hyperparameters searched by the small proxy task can be transferred to the full-training task.

Exp. Index	Task	Random	HEBO	NRBO
1	proxy	67.268	65.808	66.490
	full	68.324	67.660	68.362
2	proxy	65.644	67.258	66.972
	full	67.858	68.642	68.562
3	proxy	66.290	66.908	67.058
	full	67.668	68.494	68.446
4	proxy	66.916	66.116	67.020
	full	68.638	67.848	68.606
Avg.	proxy	66.530	66.523	66.885
	full	68.122	68.161	68.494

Exp. Index	Task	Random	HEBO	NRBO
1	proxy	32.283	32.070	34.853
	full	34.549	34.438	36.952
2	proxy	32.085	32.334	33.231
	full	33.773	34.609	35.634
3	proxy	31.705	33.571	34.410
	full	34.123	35.997	36.199
4	proxy	30.147	31.395	33.222
	full	32.487	33.585	35.307
Avg.	proxy	31.555	32.343	33.929
	full	33.733	34.657	36.023

Table 3: (Left) Proxy task and full-training task results on ImageNet. (Right) Proxy task and full-training task results on COCO.

4.4 Ablation Study

In this section, we analyze the efficiency of each component proposed in NRBO. We first conduct experiments with a variant NRBO that cancels the density-based acquisition function. If we further cancel the regularization mechanism proposed in section 3.2.1, the NRBO will degenerate to a normal HEBO optimizer.

Figure 3 left shows the performance of HEBO, variant NRBO that without density-based acquisition function and a full NRBO. We can find that the variant NRBO still outperforms the basic HEBO within a wide range of computation resource settings. The variant NRBO significantly surpasses the HEBO from the tenth to the twelfth iteration. As the optimization progresses, the variant NRBO and HEBO finally reach the tie in the end. Note that the full NRBO still outperforms the HEBO at the end of the optimization. It suggests that the density-based acquisition function helps the optimizer jump out of the local optimal and leads to a better solution.

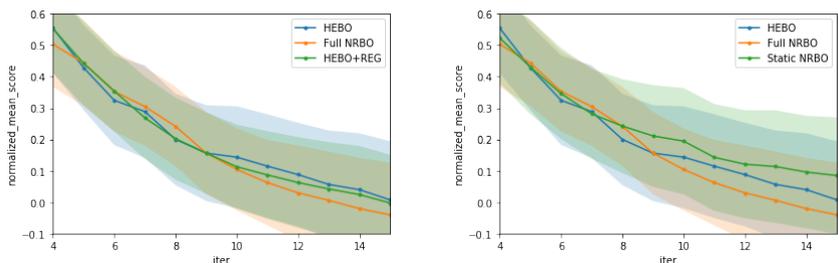


Figure 3: (Left) Ablation study for regularization and acquisition function. HEBO denotes the NRBO w/o regularization and density-based acquisition reward. HEBO+REG denotes the NRBO w/o density-based acquisition reward. (Right) Ablation study for dynamic regularization and acquisition function. Static NRBO denotes the NRBO w/o dynamic regularization and acquisition function.

We also compare the NRBO with its variant that cancels the dynamic regularization strength and density reward described in section 3.2.3. Figure 3 right shows that, without the dynamic regularization strength and density reward, the performance of NRBO deteriorates

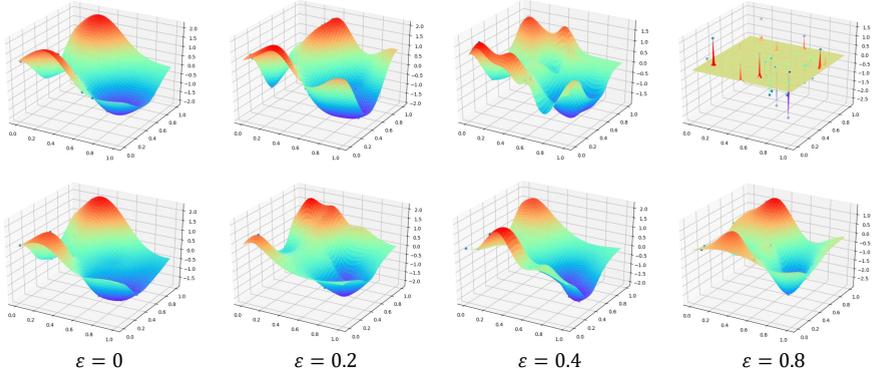


Figure 4: **Fitting surrogate model on the same dataset with different noise level.** From left to right: Gradually increased noise level ε , from 0 to 0.8. Top row: A typical Bayesian optimization with Gaussian process and Matern Kernel. Bottom row: Neighbor regularized Bayesian optimization.

severely as the optimization progresses. This is reasonable because always keeping a strong regularization will prevent the surrogate model from fitting the subtle patterns at the high-performance area where the search process is close to optimal. In addition, a fixed density-based acquisition reward encourages the optimizer to explore rather than exploit. It will also prevent the optimizer from a detailed search in the high-performance area.

4.5 Qualitative Visualization

In this section we visualize the effectiveness of regularization on an example dataset. We compare the standard Gaussian process Bayesian optimizer with a neighbor-based regularized optimizer on different noise levels. The search space is normalized to $[0, 1]$ and the neighbor radius for regularization is set to 0.1. The example dataset is generated with $\sin(2\pi x) + \cos(2\pi y) + \varepsilon\sigma$, where ε is the noise level and $\sigma \sim N(0, 1)$.

As shown in Figure 4, in the first row, both of them fit the dataset well when the noise level is 0. When the noise level is set to 0.4, standard Gaussian process model starts to overfit the noise data and its output tends to be sharp and unstable. If we further increase the noise level to 0.8, the standard Gaussian process model completely collapsed. Meanwhile, the regularized model still outputs a smooth prediction that is very close to the ground truth.

5 Conclusion

We propose a novel hyperparameter optimization algorithm NRBO in this work. We first propose a neighbor-based regularization to smooth sample observations by the neighbor statistics. To further improve the stability of the neighbor statistics, we propose a density-based acquisition function. In addition, an adjustment mechanism is adopted to adjust the regularization strength and acquisition reward based on the remaining computation resources. Extensive experiments demonstrate NRBO could accelerate the convergence of hyperparameter optimization and reduce the risk of collapse.

References

- [1] Jason Ansel, Shoaib Kamil, Kalyan Veeramachaneni, Jonathan Ragan-Kelley, Jeffrey Bosboom, Una-May O’Reilly, and Saman Amarasinghe. Opentuner: An extensible framework for program autotuning. In *Proceedings of the 23rd international conference on Parallel architectures and compilation*, pages 303–316, 2014.
- [2] James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *Journal of machine learning research*, 13(2), 2012.
- [3] James Bergstra, Daniel Yamins, and David Cox. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In *International conference on machine learning*, pages 115–123. PMLR, 2013.
- [4] Alexander I Cowen-Rivers, Wenlong Lyu, Zhi Wang, Rasul Tutunov, Hao Jianye, Jun Wang, and Haitham Bou Ammar. Hebo: Heteroscedastic evolutionary bayesian optimisation. *arXiv e-prints*, pages arXiv–2012, 2020.
- [5] Samuel Daulton, David Eriksson, Maximilian Balandat, and Eytan Bakshy. Multi-objective bayesian optimization over high-dimensional search spaces. *CoRR*, abs/2109.10964, 2021. URL <https://arxiv.org/abs/2109.10964>.
- [6] Katharina Eggensperger, Philipp Müller, Neeratyoy Mallik, Matthias Feurer, René Sass, Aaron Klein, Noor H. Awad, Marius Lindauer, and Frank Hutter. Hpobench: A collection of reproducible multi-fidelity benchmark problems for HPO. In Joaquin Vanschoren and Sai-Kit Yeung, editors, *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*, 2021.
- [7] David Eriksson, David Bindel, and Christine A Shoemaker. pysot and poap: An event-driven asynchronous framework for surrogate optimization. *arXiv preprint arXiv:1908.00420*, 2019.
- [8] David Eriksson, Michael Pearce, Jacob Gardner, Ryan D Turner, and Matthias Poloczek. Scalable global optimization via local bayesian optimization. *Advances in Neural Information Processing Systems*, 32:5496–5507, 2019.
- [9] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2): 303–338, June 2010.
- [10] Stefan Falkner, Aaron Klein, and Frank Hutter. Bohb: Robust and efficient hyperparameter optimization at scale. In *International Conference on Machine Learning*, pages 1437–1446. PMLR, 2018.
- [11] Evelyn Fix and Joseph Lawson Hodges. Discriminatory analysis. nonparametric discrimination: Consistency properties. *International Statistical Review/Revue Internationale de Statistique*, 57(3):238–247, 1989.
- [12] Daniel Golovin, Benjamin Solnik, Subhdeep Moitra, Greg Kochanski, John Karro, and David Sculley. Google vizier: A service for black-box optimization. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1487–1495, 2017.

- [13] Elad Hazan, Adam Klivans, and Yang Yuan. Hyperparameter optimization: A spectral approach. *arXiv preprint arXiv:1706.00764*, 2017.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [15] Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. Sequential model-based optimization for general algorithm configuration. In *International conference on learning and intelligent optimization*, pages 507–523. Springer, 2011.
- [16] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, Sydney, Australia, 2013.
- [17] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [18] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012.
- [19] PM Lerman. Fitting segmented regression models by grid search. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 29(1):77–84, 1980.
- [20] Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. *The Journal of Machine Learning Research*, 18(1):6765–6816, 2017.
- [21] Petro Liashchynskyi and Pavlo Liashchynskyi. Grid search, random search, genetic algorithm: A big comparison for nas. *arXiv preprint arXiv:1912.06059*, 2019.
- [22] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [23] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [24] Marius Lindauer, Katharina Eggensperger, Matthias Feurer, André Biedenkapp, Difan Deng, Carolin Benjamins, Tim Ruhkopf, René Sass, and Frank Hutter. SMAC3: A versatile bayesian optimization package for hyperparameter optimization. *J. Mach. Learn. Res.*, 23:54:1–54:9, 2022.
- [25] J. Rapin and O. Teytaud. Nevergrad - A gradient-free optimization platform. <https://GitHub.com/FacebookResearch/Nevergrad>, 2018.
- [26] Carl Edward Rasmussen. Gaussian processes in machine learning. In *Summer school on machine learning*, pages 63–71. Springer, 2003.

- [27] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2015.
- [28] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical bayesian optimization of machine learning algorithms. *Advances in neural information processing systems*, 25, 2012.
- [29] Jasper Snoek, Oren Rippel, Kevin Swersky, Ryan Kiros, Nadathur Satish, Narayanan Sundaram, Mostofa Patwary, Mr Prabhat, and Ryan Adams. Scalable bayesian optimization using deep neural networks. In *International conference on machine learning*, pages 2171–2180. PMLR, 2015.
- [30] Francisco J Solis and Roger J-B Wets. Minimization by random search techniques. *Mathematics of operations research*, 6(1):19–30, 1981.
- [31] Kevin Swersky, Jasper Snoek, and Ryan P Adams. Multi-task bayesian optimization. *Advances in neural information processing systems*, 26, 2013.
- [32] Timgates42. scikit-optimize, 2020.
- [33] A Helen Victoria and G Maragatham. Automatic tuning of hyperparameters using bayesian optimization. *Evolving Systems*, 12(1):217–223, 2021.
- [34] Jia Wu, Xiu-Yun Chen, Hao Zhang, Li-Dong Xiong, Hang Lei, and Si-Hao Deng. Hyperparameter optimization for machine learning models based on bayesian optimization. *Journal of Electronic Science and Technology*, 17(1):26–40, 2019.
- [35] Li Yang and Abdallah Shami. On hyperparameter optimization of machine learning algorithms: Theory and practice. *Neurocomputing*, 415:295–316, 2020.