

Instance Generation for Meta-Black-Box Optimization Through Latent Space Reverse Engineering

Chen Wang¹, Yue-Jiao Gong¹, Zhiguang Cao², Zeyuan Ma^{1*}

¹South China University of Technology

²Singapore Management University

Abstract

To relieve intensive human-expertise required to design optimization algorithms, recent Meta-Black-Box Optimization (MetaBBO) researches leverage generalization strength of meta-learning to train neural network-based algorithm design policies over a predefined training problem set, which automates the adaptability of the low-level optimizers on unseen problem instances. Currently, a common training problem set choice in existing MetaBBOs is well-known benchmark suites CoCo-BBOB. Although such choice facilitates the MetaBBO’s development, problem instances in CoCo-BBOB are more or less limited in diversity, raising the risk of overfitting of MetaBBOs, which might further results in poor generalization. In this paper, we propose an instance generation approach, termed as **LSRE**, which could generate diverse training problem instances for MetaBBOs to learn more generalizable policies. LSRE first trains an autoencoder which maps high-dimensional problem features into a 2-dimensional latent space. Uniform-grid sampling in this latent space leads to hidden representations of problem instances with sufficient diversity. By leveraging a genetic-programming approach to search function formulas with minimal L2-distance to these hidden representations, LSRE reverse engineers a diversified problem set, termed as **Diverse-BBO**. We validate the effectiveness of LSRE by training various MetaBBOs on Diverse-BBO and observe their generalization performances on either synthetic or realistic scenarios. Extensive experimental results underscore the superiority of Diverse-BBO to existing training set choices in MetaBBOs. Further ablation studies not only demonstrate the effectiveness of design choices in LSRE, but also reveal valuable insights on instance diversity and MetaBBO’s generalization.

Code — <https://github.com/MetaEvo/Diverse-BBO>

1 Introduction

Black-Box Optimization (BBO) problems, with neither the objective function nor corresponding derivative is accessible, are challenging to solve. Although Evolutionary Computation (EC) methods, such as Genetic Algorithm (GA), Particle Swarm Optimization (PSO), Differential Evolution (DE), as well as their modern (self-)adaptive variants are

widely recognized as effective solutions, they often require tuning or even re-design for novel scenario adaption.

To automate such manual algorithm design tasks, recent Meta-Black-Box Optimization (MetaBBO) researches propose learning generalizable design policy through meta-learning (Ma et al. 2023, 2025d). MetaBBO naturally forms a bi-level algorithmic structure. At the meta level, a learnable algorithm design policy (e.g., a neural network) outputs flexible design choices (e.g., parameter control (Ma et al. 2024a; Guo et al. 2025b) and algorithm selection (Guo et al. 2024)) for low-level optimizer. At the lower level, a BBO optimizer is deployed to optimize target problem with the given design choice. The meta-objective is to train a design policy that enhances low-level BBO process. Since such paradigm provides a promising way to learn automated optimization algorithm design, MetaBBO has been widely explored and discussed in recent two years (Shao, Tian, and Zhang 2025; Guo et al. 2025c; Shao et al. 2025; Huang et al. 2025; Chen, Feng, and Cheng 2025; Mo et al. 2025).

Given that MetaBBO essentially follows meta-learning paradigm (Thrun and Pratt 1998; Finn, Abbeel, and Levine 2017), it unavoidably requires domain generalization designs (Li et al. 2018), i.e., meta-augmentation of training data (Rajendran, Irpan, and Jang 2020), so as to ensure the adaption during the testing time. To be specific, as shown in the yellow box of Fig. 1, existing MetaBBO approaches typically train their policies over a training problem set. The diversity of this set assures that the meta-level policy generalizes across diverse optimization problem characteristics (Ma et al. 2025b). Common choices of the problem set in existing MetaBBO approaches are synthetic benchmarks such as CoCo-BBOB (Hansen et al. 2010a). Although these representative problem instances hold a long-standing history and reputation in the evaluation of BBO methods, they are primarily hand-crafted to differentiate algorithms apart from each other in terms of optimization performance (Langdon and Poli 2007). Recent studies (Ma et al. 2025d,c) reveal that such design bias may not keeps pace with the training&evaluation objectives of MetaBBO, i.e., generalization toward unseen problems. One possible solution is deriving realistic instances directly from real-world application (Kumar et al. 2020). However, collecting and formulating realistic problems require deep expertise. More importantly, the corresponding expensive simulation costs makes this so-

*Corresponding author (scut.crazy nicolas@gmail.com)
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

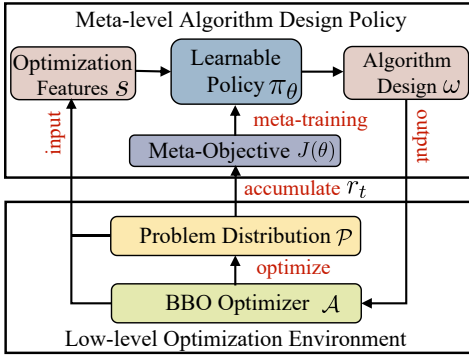


Figure 1: General workflow of MetaBBO approaches.

lution impractical (Muñoz and Smith-Miles 2020). A more promising way is generating synthetic problem under proper diversity measurement (Tian et al. 2020; Vermetten et al. 2025), however, to the best of our knowledge, related researches are still very limited.

To address the aforementioned limitations, we propose a novel generation framework that discovers diverse synthetic optimization problems through **latent space reverse engineering**, termed as **LSRE**, to advance the benchmark development in BBO and MetaBBO fields. Specifically, LSRE comprises following key designs: 1) an instance space analysis (Smith-Miles and Muñoz 2023) method that trains an autoencoder to map the Exploratory Landscape Analysis (ELA) (Mersmann et al. 2011) features of problem instances into 2-dimensional latent space; 2) an expressive symbol set including operands and operators that spans a comprehensive functional instance space; 3) a genetic programming method that effectively searches for function instance within the latent space, with an enhanced local search strategy and parallel acceleration. By grid-sampling 256 instance points from the latent space, we repeatedly call LSRE to search optimal function formula for each of them, resulting in a collection of 256 synthetic problem instances **Diverse-BBO** with adequate diversity. We validate the diversity advantage of our method against existing baseline benchmarks through large-scale MetaBBO benchmarking. The comprehensive experimental results demonstrate the effectiveness of our LSRE framework and the superiority of the generated Diverse-BBO to baseline benchmarks. We now summarize the core contributions of this paper:

- *Novelty*: our paper is the first study on how to prepare high generalization potential training data for MetaBBO.
- *Methodology*: a novel GP framework **LSRE** for generating diversified synthetic optimization problems.
- *Benchmark*: instances generate by LSRE constitute a novel benchmark **Diverse-BBO** which could be used for development of either BBO or MetaBBO approaches.

2 Related Works

2.1 Meta-Black-Box Optimization.

To automate the manual design and configuration of BBO algorithms, MetaBBO researches propose meta-learning algorithm design policy typically by a bi-level optimization framework (as illustrated in Fig. 1): 1) At the meta-level, an algorithm design policy π_θ outputs a desired design ω for low-level optimizer \mathcal{A} base on low-level optimization state s ; 2) At the lower level, the BBO optimizer \mathcal{A} is employed to optimize problem p sampled from distribution \mathcal{P} with ω . The corresponding performance gains r in low-level optimization are collected as meta-objective $J(\theta) = \mathbb{E}_{p \in \mathcal{P}} \left[\sum_{t=1}^T r_t \right]$, where T is optimization horizon. π_θ is then meta-trained to maximize the meta-objective.

Existing MetaBBO approaches can be divided into several lines according to their specific algorithm design tasks: 1) Algorithm Selection (AS), where predefined algorithms/operators are selected by the meta-level policy to solve given problems. Early researches explored selecting operators within DE frameworks (Sharma et al. 2019; Tan and Li 2021), and more recent approaches have investigated dynamic switching between multiple advanced optimizers along the optimization process (Guo et al. 2024); 2) Algorithm Configuration (AC), where given an existing optimizer, the meta-level policy dictates desired parameter values dynamically along the optimization process (Sun et al. 2021; Lange et al. 2023; Ma et al. 2024a; Guo et al. 2025b); 3) Algorithm Generation (AG), where the meta-level policy constructs both individual components and the entire workflow to synthesize novel BBO algorithms through algorithm workflow composition (Zhao et al. 2024), novel update rules generation (Chen et al. 2024) and language modelling (Ma et al. 2024b; Liu et al. 2024; van Stein and Bäck 2024). In addition, studies on MetaBBO have also been extended to automated feature learning (Ma et al. 2025b; Guo et al. 2025a), efficient offline learning (Ma et al. 2025a,e) and direct solution manipulation (Li et al. 2024, 2025).

2.2 Existing Benchmark Functions.

Existing MetaBBO methods typically adopt synthetic problems from well-established BBO benchmarks (eg., CoCoBBOB (Hansen et al. 2010a), IOHProfiler (Doerr et al. 2019), IEEE CEC Competition series (Mohamed et al. 2021)) and platforms (eg., MetaBox-v1/v2 (Ma et al. 2023, 2025c)). While these benchmark functions have been widely recognized and used for evaluating BBO methods, they are primarily hand-crafted and are collected with inherent biases (Hooker 1995), aimed at creating performance distinctions between optimization algorithms (Langdon and Poli 2007). The existing benchmarks' design bias and lack of diversity may constrain the MetaBBO to learn generalizable design policies that excel at not only the training problem distribution but also unseen problems with various landscape features (Ma et al. 2025d,b). Constructing and collecting realistic problems (Kumar et al. 2020; Hwang et al. 2010) could be a potential solution, but it requires deep expertise and is further limited by simulation costs (Muñoz and Smith-Miles 2020), ultimately rendering the solution impractical.

Generating diverse synthetic problems might be more promising, yet several challenges persist: 1) The random function generator (Tian et al. 2020) lacks control over the characteristics and diversity of generated problems; 2) A recent line of works explore the possibility of augmenting new problem instances through affine combination of existing instances in BBO benchmark (Dietrich and Mersmann 2022). A representative method is MA-BBOB (Vermetten et al. 2025), which claimed achieved efficient and fine-grained problem generation. However, a key limitation is such approaches might restrict themselves within the selected composition functions. 3) In the light of symbolic regression (Schmidt and Lipson 2009), some exploratory studies leveraged symbolic discovery tool such as Genetic Programming (GP) (Koza 1992) to generate diversified instances by searching corresponding mathematical symbolic formulas (Muñoz and Smith-Miles 2020; Long et al. 2023). However, such researches are still limited. More importantly, given the inherent expensive running of GP and the design bottlenecks in these methods, their actual generation effectiveness is still unsatisfactory. In this paper, we propose LSRE to address the limitations in the third line of work in both effectiveness and efficiency side. On the one hand, we incorporate LSRE with several novel designs: 1) Effective latent instance space analysis through autoencoder; 2) Expressive symbol set enabling both efficient generation and flexible operations; 3) Cross-dimensional local search strategy to further enhance the GP searching. On the other hand, we introduce high-performance distributed acceleration techniques to facilitate efficient GP running.

3 Methodology

3.1 Latent Instance Space Analysis

To align with practical usage of potential users, one of the most important issues in our framework is to pre-define a target problem range \mathcal{D} as the basis for constructing corresponding the latent instance space. While interests of different users may vary, in this paper we primarily select CoCo-BBOB (Hansen et al. 2010a) to serve as \mathcal{D} due to its ease of acquisition and long-standing reputation within BBO evaluation domain. We have to note that our framework can be easily extended to other user-specific problem ranges. We leave step-by-step instructions for such cases at Appendix A.1¹. Given the 24 basic functions ($f_1 \sim f_{24}$) in CoCo-BBOB, we first augment these functions with five different problem dimensions [2, 5, 10, 30, 50], resulting in 120 functions. Next, we further generate a sufficient number of instances by applying random rotation and shift transformations to these 120 functions:

$$f'(x) = f(\mathbf{R}^T(x - s)) \quad (1)$$

where f denotes the original function and f' the transformed one, x is the decision variables, \mathbf{R} and s are random rotation matrix and shift vector, respectively. We randomly generate 270 instances for each of the 120 functions following the equation above, resulting in 32,400 instances as selected \mathcal{D} , which is used in following sections. We construct

¹“Appendix.pdf” in our codebase.




Figure 2: Workflow of Latent Instance Space Analysis.

latent instance space of \mathcal{D} by combining the Exploratory Landscape Analysis (ELA) (Mersmann et al. 2011) and popular dimension reduction strategy autoencoder (Rumelhart, Hinton, and Williams 1985). We illustrate the workflow of such latent instance space analysis in Fig. 2, and briefly introduce how these two techniques cooperate with each other:

ELA. It is a widely used tool for instance space analysis (Škvorc, Eftimov, and Korošec 2020) of BBO problems. Groups of analytic or statistical features, such as convexity, global structure and meta-model fitting are computed in ELA, which has been proven informative in algorithm selection (Kerschke and Trautmann 2019). Given a BBO problem instance $f \in \mathcal{D}$, we obtain its ELA features vector E_f using open-source ELA package: pflacco (Prager and Trautmann 2023). Specifically, we select a subset of representative and computationally efficient features as E_f . The selection standard and feature list are provided in Appendix A.2 .

Autoencoder. An autoencoder comprises two network components: the encoder W_θ and the decoder W_ϕ . The input to W_θ is a high-dimensional feature vector, and it then maps this feature vector into a low-dimensional hidden space. Given the hidden feature output by W_θ , the decoder W_ϕ maps it back to the original high-dimensional feature space. In our case, the high-dimensional space spans the ELA features space of the selected problem range \mathcal{D} , and the hidden space is a 2-dimensional space denoted as \mathcal{H} . The training of our autoencoder follows the normal reconstruction loss function used in existing literature. Given a BBO problem instance $f \in \mathcal{D}$, we first compute its ELA features E_f , the reconstructed feature E'_f is obtained by feeding E_f into W_θ and W_ϕ . Then the reconstruction loss is formulated as:

$$\mathcal{L}(\theta, \phi) = \frac{1}{2} \|E_f - E'_f\|_2^2, \quad f \sim \mathcal{D} \quad (2)$$

We leave the settings of the autoencoder such as its architecture and corresponding training protocol in Appendix A.3. Once the training of the autoencoder finishes, we could obtain latent representation $h \in \mathbb{R}^2$ of a BBO problem instance through the encoder. More importantly, the well-trained model allows us to measure the diversity of problem instances in a 2-dimensional space, which further facilitates uniform instance sampling from this hidden space, since directly sampling from high-dimensional ELA space is neither efficient nor practical. It is worth mentioning that principal component analysis (PCA) (Jolliffe 2011) is also a popular dimension reduction approach widely adopted in machine learning community. The reason behind our autoencoder choice is that: 1) PCA shows effectiveness in cases




Figure 3: A proximal symbolic tree example.

where features hold linear correlation assumption (Almotiri, Elleithy, and Elleithy 2017; Wang, Yao, and Zhao 2016), however, different feature groups in ELA are independently computed; 2) We have conducted a preliminary study on using PCA for ELA feature reduction, the results show that the first two principle components of certain portion of instances in \mathcal{D} can not surpass 51% importance threshold. We also provide an ablation study in Sec. 4.3.

3.2 Genetic Programming Search

Through the latent instance space analysis, we now obtain a 2-dimensional latent instance space \mathcal{H} . To generate synthetic BBO problem instance from \mathcal{H} , several challenges still remain. One of them is how we model a proximal function space to represent or formulate mathematical functions of optimization problems as computer-understandable language. To address this, we adopt a symbolic tree representation for mathematical function, which has proven to be robust and effective in symbolic regression domain (Schmidt and Lipson 2009). To make it a comprehensive proximal space, we design a flexible symbol set containing diverse operators and operands, which facilitates accurate and efficient proximal representation. Another key challenge is how we reverse engineer a proximal representation of a BBO problem from the latent space \mathcal{H} . Genetic Programming (GP) (Koza 1992) is regarded as an effective way to discover functions from symbolic function spaces (Augusto and Barbosa 2000). In this paper, we propose a novel GP framework to search for the fittest function formula given a latent instance point in \mathcal{H} . This includes how we define the searching objective and how we enhance naive GP in both effectiveness and efficiency sides. We present the pseudocode of the GP search in Alg. 1 and elaborate our designs in the following sections.

Proximal Symbolic Space We propose a comprehensive symbol set comprising flexible operators and operands that spans a proximal symbolic space in regard of single-objective continuous optimization problems. Specifically, the symbolic set comprises following symbols:

- *Operator*: 15 operators including add, sub, mul, div, neg, pow, sin, cos, tanh, exp, log, sqrt, abs, sum, and mean. Except sum and mean, the listed operators are all unary/binary operators commonly used in

symbolic regression research. We replenish them with sum and mean as additional aggregation operation, which could simplify multi-addend structure in symbolic tree, so as to enhance the searching efficiency of GP.

- *Operand*: three operands including X , $X[i:j]$ and C . X denotes the vector of decision variables. C denotes constant values or vectors. $X[i:j]$ denotes indexed decision variables, e.g., $X[2:4]$ indexes the second to fourth decision variables while $X[2:2]$ denotes the second one. It is worthy to mention that compared with symbolic set design in existing literature which only involve X and C , we add $X[i:j]$ to achieve fine-grained generation.

With the comprehensive symbol set, we can sample diverse optimization problems by constructing their equivalent symbolic trees from the corresponding symbolic space (see Fig. 3). In the proposed GP search, we represent each candidate function instance by its symbolic tree and use tree-based evolutionary operators to search for target function instance. We leave a more detailed discussion of the symbol set at Appendix A.4.

Searching Objective The objective of our proposed GP search is quite straightforward. Considering we want to reverse engineer (i.e., search for) the mathematical formula from a given instance point $h \in \mathcal{H}$, where \mathcal{H} is the 2-dimensional latent instance space we constructed. We use τ to denote any legal symbolic tree constructed from our symbol set, and use E_τ to denote the landscape features of the corresponding problem instance. Then the searching objective is to minimize a score function $obj(\tau) = \frac{1}{2} \|W_\theta(E_\tau) - h\|_2$, where W_θ is the encoder of the autoencoder in latent instance space analysis, and to find an optimal τ^* such that:

$$\tau^* = \arg \min_{\tau} obj(\tau) \quad (3)$$

Searching Strategy We aim to reverse engineer a synthetic problem instance by finding its optimal proximal representation τ^* in Eq. 3. In LSRE, we adopt *gplearn* (Stephens 2023) to achieve this goal, which is an effective and robust GP library. We initialize a *gplearn* process with our proposed symbol set and searching objective.

The pseudocode of the GP search in LSRE is provided in Alg. 1. In lines 2-3, we first randomly initialize a population of proximal tree representations $\{\tau_i\}_{i=1}^N$ by sampling from the symbolic space spanned by our proposed symbol set, and then evaluate them to attain their objective values $obj(\cdot)$. Then we begin an iterative searching process in lines 4-13. For each generation g , the solution population first undergoes a roulette_reproduction operation (line 6). It is a built-in reproduction strategy of *gplearn*, which randomly assign one of five default mutation/crossover operators for each tree individual, by a predetermined operator selection preference. After roulette_reproduction operation, an offspring population $\{\tau'_i\}_{i=1}^N$ is obtained. In line 7, we propose a novel plug-and-play cross-dimensional local search strategy which can be easily integrated into the GP process to improve the searching accuracy. The motivation behind this local search strategy is that: although we could search for optimal function formulation, the dimension of

Algorithm 1: Genetic Programming Search in LSRE

Input: population size N ; budget G ; searching target h .**Output:** best instance τ^* found ever.

```
1: Let  $g = 1$ 
2: Initialize a tree population  $\{\tau_1, \dots, \tau_N\}$ 
3: Evaluate the tree population  $\{obj(\tau_1), \dots, obj(\tau_N)\}$ 
4: while  $g \leq G$  do
5:   for each tree individual  $\tau_i$  do
6:      $\tau'_i \leftarrow \text{roulette\_reproduction}(\tau_i)$ 
7:      $\tau''_i \leftarrow \text{cross\_dimension\_local\_search}(\tau'_i)$ 
8:     if  $obj(\tau''_i) \leq obj(\tau_i)$  then
9:        $\tau_i \leftarrow \tau''_i$ .
10:    end if
11:  end for
12:   $g = g + 1$ 
13: end while
14:  $\tau^* \leftarrow \arg \min_{\tau_i} obj(\tau_i)$ 
15: return  $\tau^*$ 
```

the decision variables is not addressed. Several recent studies (Renau et al. 2020; van Stein et al. 2023) reveal that even two function instances share mathematical formulation, their ELA features may vary due to different problem dimensions. To address this, in our proposed local search strategy, for each offspring τ'_i , we instantiate it to 2D, 5D and 10D function instances and compare their objective values, then the best objective value is regarded as the objective value of this offspring. We provide effectiveness validation in following ablation studies. After a total of G generations GP searching, we obtain the best solution τ^* found so far (line 14) as well as its best-performing problem dimension.

3.3 Diverse-BBO

Overall Workflow With the latent instance space and the GP searching strategy at hand, we now elaborate on how LSRE generates uniform-diversity function instances for Diverse-BBO. We first demarcate a grid area in the 2-dimensional latent space, i.e., $x \in [-B, B]$ and $y \in [-B, B]$, which could cover the target problem range \mathcal{D} we constructed in Sec. 3.1. Then we uniformly sample $K \times K$ latent instance points (point matrix) from this grid area, where K controls the sampling granularity. At last, for each sampled instance point, we call our GP search to reverse engineer -corresponding function formulation. In this paper, we set $B = 1$ and $K = 16$ to obtain a problem set of 256 diversified instances, which we denote as Diverse-BBO. Refer to “DiverseBBO.pdf” in our codebase for their details.

Distributed Acceleration To obtain Diverse-BBO is challenging considering that: 1) we have to repeatedly call the GP process for 256 times; 2) Within each GP process, evaluating the tree population is also a tedious and time-consuming task due to repeated objective computation. The overall complexity is hence $\mathcal{O}[M * G * (N^2 + N * L^2)]$, where M is the number of functions we need to generate (256), G is the GP generations (50), N and L are population size and tree node size respectively. To accelerate LSRE, we




Figure 4: Average Success Rate Distribution

further introduce Ray (Moritz et al. 2018) to first distribute 256 GP processes on independent CPU cores and then trigger second-level parallel evaluation within each GP process. By doing this, we significantly reduce the overall complexity to $\mathcal{O}[G * L^2]$. We leave the hyper-parameters including their selection standards and the detailed design philosophy behind LSRE’s components at Appendix A.5 and Appendix A.6.

4 Experimental Analysis

We validate the effectiveness of our proposed LSRE through comparing Diverse-BBO generated by LSRE with representative baseline benchmarks. In specific, we consider following research questions. **RQ1:** Can Diverse-BBO show more diversity compared to other synthetic BBO benchmarks? **RQ2:** Can MetaBBO approaches benefit from such diversity in terms of meta-learned generalization? **RQ3:** How does our proposed novel designs contribute to LSRE? Below we first introduce all baselines and then discuss RQ1~RQ3 in the following sections. The baselines are: 1) Diverse-BBO: it includes 256 diversified synthetic problems generated by our LSRE, the problem dimensions for these instances are 2, 5, and 10; 2) CoCo-BBOB (Hansen et al. 2010a): it comprises 24 basis synthetic functions, which we used to generate 256 10-dimensional instances by transforming them via random rotation and shift; 3) MA-BBOB (Vermetten et al. 2025): it uses 24 basis functions in CoCo-BBOB as component functions and generates new instance by affine combination of them. We use MA-BBOB to generate 256 10-dimensional instances as a baseline; 4) GP-BBOB (Muñoz and Smith-Miles 2020): the first framework that uses GP for BBO instance generation. Our LSRE differs with this method in instance space analysis, symbol set construction and searching strategy. We use it to generate 256 instances as a baseline.

4.1 Diversity Comparison (RQ1)

We first show how we measure the diversity of a benchmark. Existing literature of BBO benchmarks barely provide explicit definition of diversity for the testing problem instances included. However, some existed papers indeed provide indirect measurements that reflect benchmark diversity through algorithm performance distribution (Hansen et al. 2010b). Following this idea, we propose a simple and intuitive average success rate distribution (ASRD) test to reflect the diversity of a benchmark through the problem difficulty

	Test Set: $\mathbb{D}_{synthetic}$									
	DEDDQN		LDE		SYMBOL		GLEET		Average	
	Mean \uparrow (\pm Std)	Rank	Mean \uparrow (\pm Std)	Rank	Mean \uparrow (\pm Std)	Rank	Mean \uparrow (\pm Std)	Rank	Mean \uparrow (\pm Std)	Rank
Diverse-BBO	8.154e-01 (\pm 5.219e-02)	1	8.315e-01 (\pm 3.136e-02)	1	7.346e-01 (\pm 9.217e-02)	1	8.106e-01 (\pm 5.983e-02)	2	7.980e-01 (\pm 5.889e-02)	1.25
CoCo-BBOB	8.106e-01 (\pm 5.217e-02)	2	8.271e-01 (\pm 3.128e-02)	2	7.319e-01 (\pm 7.363e-02)	2	8.058e-01 (\pm 4.146e-02)	4	7.939e-01 (\pm 4.963e-02)	2.5
MA-BBOB	8.006e-01 (\pm 5.133e-02)	4	8.232e-01 (\pm 3.741e-02)	3	7.021e-01 (\pm 9.149e-02)	3	8.112e-01 (\pm 4.664e-02)	1	7.843e-01 (\pm 5.672e-02)	2.75
GP-BBOB	8.045e-01 (\pm 5.378e-02)	3	8.223e-01 (\pm 3.938e-02)	4	6.907e-01 (\pm 7.625e-02)	4	8.088e-01 (\pm 4.827e-02)	3	7.816e-01 (\pm 5.442e-02)	3.5

Table 1: Generalization performance of MetaBBO when being trained on different baseline benchmarks.

distribution. In ASRD test, for a given benchmark that includes N problem instances $f_1 \sim f_N$, we first construct a BBO optimizer pool that comprises M representative optimizers $\mathcal{A}_1 \sim \mathcal{A}_M$. For any pair of problem instance f_i and optimizer \mathcal{A}_j , we use \mathcal{A}_j to optimize f_i for 51 independent runs and record the portion of successful runs (runs that achieve optimal on f_i) as success rate $SR_{i,j}$. Then the ASRD for the given benchmark is summarized as the statistical histogram of all success rate results $\{\{SR_{i,j}\}_{i=1}^N\}_{j=1}^M$.

In this experiment, we conduct ASRD test for all four baselines, the optimizer pool we construct include several representative optimizers such as DE, PSO, CMA-ES etc.. For more detailed experiment protocol, we kindly request readers refer to Appendix B.1. We visualize the test results in Fig. 4, which show that: 1) **Diverse-BBO v.s. CoCo-BBOB**: derived from the latent space of CoCo-BBOB, our Diverse-BBO achieves preferred problem diversity through fine-grained generation of LSRE; 2) **Diverse-BBO v.s. MA-BBOB**: certain diversity improvement can be observed in MA-BBOB compared to CoCo-BBOB, however, the affine combination in MA-BBOB still fails to achieve controllable uniform diversity distribution; 3) **Diverse-BBO v.s. GP-BBOB**: sharing the same GP-based generation backbone, the novel designs we proposed in LSRE result in better generation quality, which we will discuss more in the subsequent sections.

4.2 Generalization Test (RQ2)

Another key research question is whether the diversity of a BBO benchmark truly provides convenience for the generalization of MetaBBO. To this end, we conduct a large-scale validation experiment in this section.

Settings For the four baseline benchmarks, we use them for training MetaBBO approaches and then observe their generalization performance. This results in three experiment protocols we have to clarify here. First, for each baseline, we randomly split it into 75% training instances and 25% testing instances. We then collect all testing instances of the four benchmarks into one test set, termed as $\mathbb{D}_{synthetic}$. Second, we further prepare another three test sets \mathbb{D}_{hpo} , \mathbb{D}_{uav} and $\mathbb{D}_{protein}$, which denote complex realistic BBO problems: hyper-parameter optimization (Arango et al. 2021), uav planning (Shehadeh and Kudela 2025) and protein docking (Hwang et al. 2010). Third, we select four representative MetaBBO approaches from existing literature: DEDDQN (Sharma et al. 2019), LDE (Sun et al.




Figure 5: Generalization performance on realistic problems.

2021), SYMBOL (Chen et al. 2024) and GLEET (Ma et al. 2024a). These baselines cover different methodology categories such as those for operator selection (DEDDQN), for algorithm configuration (LDE & GLEET) and for algorithm generation (SYMBOL). We use their default settings for training and testing according to the original papers. More details about these protocols are provided at Appendix B.2.

Synthetic Test We train the four MetaBBO approaches on the four baseline benchmarks’ train sets respectively, with the same random seed to ensure the comparison fairness. We then test these 16 (4 MetaBBOs, 4 train sets) trained models on $\mathbb{D}_{synthetic}$, with 51 independent runs on each problem instance. We adopt the average normalized optimization performance metrics (larger is better) widely used in MetaBBO approaches (Guo et al. 2025b; Ma et al. 2025b; Li et al. 2024) in the test. The testing results are presented in Table 1, where the ranks of baseline benchmarks are annotated alongside the performance metric, and the average ranks and performance across diverse MetaBBOs are also presented in the rightmost column. The results show that generally the diversity advantage of our Diverse-BBO against the other baselines facilitates the generalization performance of various MetaBBO approaches. It only slightly underperforms MA-BBOB with large overlapping. An interesting observation is that though MA-BBOB and GP-BBOB have been claimed useful for evaluating traditional BBO optimizer, they show an opposite effect in MetaBBO domain, which anticipates further investigation in the future.

Realistic Test According to the insights obtained from recent MetaBBO-related researches, out-of-distribution generalization, especially MetaBBO’s zero-shot performance on unseen realistic problems should be thoroughly evaluated (Ma et al. 2025c). To this end, in this section we set a very high bar for the four baseline benchmarks. We test

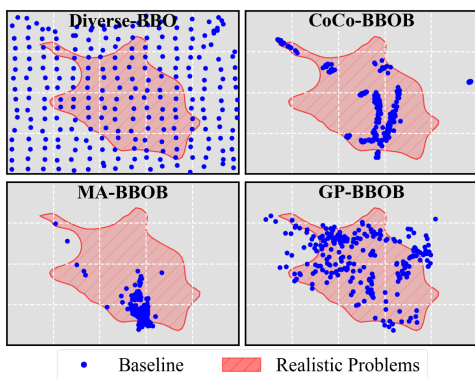


Figure 6: Generalization interpretation of baselines.

the four MetaBBO approaches trained previously on three realistic problem sets \mathbb{D}_{hpo} , \mathbb{D}_{uav} and $\mathbb{D}_{protein}$. We illustrate in Fig. 5 the average performance gains of baseline benchmarks across the four MetaBBO approaches, which is computed as $\frac{1}{rank}$. The results further demonstrate the diversity advantage of our Diverse-BBO, indicating the effectiveness of LSRE in generating high-quality training instances for MetaBBO. To further interpret the observed superiority, we provide an intuitive visualization in Fig. 6 where the shaded red area denotes the coverage of all instances of the three realistic sets in the latent space \mathcal{H} we constructed in Sec. 3.1. The blue points denotes the positions of the 256 instances of each baseline benchmark in \mathcal{H} . We believe results there could clearly explain the relationship between benchmark diversity and corresponding generalization potential. Through our LSRE, we achieve controllable, uniform and fine-grained instance generation within the instance space, which hence benefits training of MetaBBO.

4.3 Ablation Study (RQ3)

We provide following ablation baselines to validate our proposed novel designs in LSRE: 1) LSRE-PCA: where we use PCA for the latent instance space analysis, selecting the most important two features for dimension reduction; 2) LSRE-Simple: where we use simpler symbol set design of existing symbolic regression research (Tian et al. 2020), without vectorized operations; 3) LSRE-NO_LS: where the cross-dimensional local search strategy is ablated from the GP search process. Besides, we additionally include a background baseline GP-BBOB (Muñoz and Smith-Miles 2020), which can be regarded as a baseline that ablates all novel designs in LSRE. Specifically, GP-BBOB share the same GP backbone as our LSRE, while it uses the simpler symbol set, PCA-based instance space analysis and naive GP search.

We present in Fig. 7 the optimization curves of these ablation baselines and our LSRE when searching for τ^* , which are averaged across all 256 target instance points. The results not only demonstrate significant improvement LSRE achieved against GP-BBOB, but also validate that each novel design in LSRE contributes to the final performance. Specifically, while the augmented symbol set and refined autoencoder-based instance analysis provide comparably

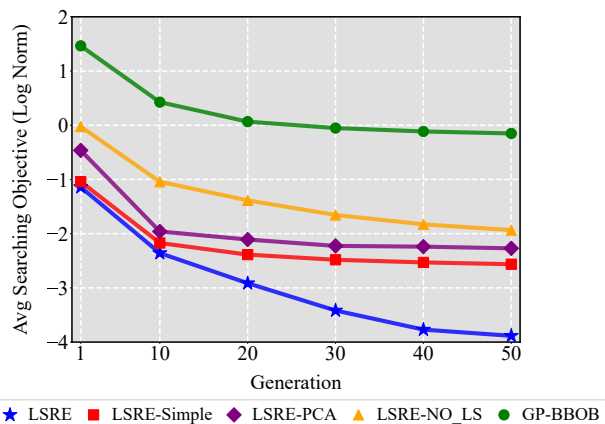


Figure 7: LSRE optimization curves when setting vary.

equal contribution, we observe that the cross-dimensional local search plays a more important role in LSRE. Compared to existing works that fix the problem dimension during their searching process (Muñoz and Smith-Miles 2020; Long et al. 2023; Vermetten et al. 2025), our finding reveals that the ELA feature space might not be smooth enough if we do not co-optimize the problem dimension. There are also several insightful and interesting findings such as the visualized comparison between PCA and autoencoder in instance space analysis, the complexity of the function instance generated by LSRE etc.. We leave them at Appendix B.3 for curious readers.

5 Conclusions

This paper serves as the first study to emphasize the importance of training data for Meta-Black-Box Optimization (MetaBBO). We have identified design bias and limitations of existing BBO benchmark and related instance generation approach through empirical evaluation. To address such design bias, we propose LSRE framework that first leverages an effective autoencoder-based latent instance space analysis to establish a comprehensive searching space. LSRE then deploys GP-based instance generation (searching) method to discover diversified synthetic problem instances. Several novel designs facilitate this process such as a replenished symbol set to enhance expressiveness for mathematical formulas and a cross-dimensional local search strategy to significantly improve the generation accuracy. We obtain several valuable conclusions from the comprehensive experiment results: 1) While existing BBO benchmarks hold long-standing reputation on evaluating traditional BBO optimizers, they may not be a proper choice for training MetaBBO; 2) The diversity of training problem set significantly impacts the generalization potential of MetaBBO; 3) The Diverse-BBO problem set generated by our LSRE framework, is at least so far, a better choice for future MetaBBO researches. We appeal more and more researchers for related studies on algorithm performance evaluation, benchmark problem diversity and generalization potential within both BBO and MetaBBO fields.

Acknowledgments

This work was supported in part by National Natural Science Foundation of China (Grant No. 62276100), in part by the Guangdong Provincial Natural Science Foundation for Outstanding Youth Team Project (Grant No. 2024B1515040010), in part by Guangzhou Science and Technology Elite Talent Leading Program for Basic and Applied Basic Research (Grant No. SL2024A04J01361), and in part by the Fundamental Research Funds for the Central Universities (Grant No. 2025ZYGXZR027).

References

- Almotiri, J.; Elleithy, K.; and Elleithy, A. 2017. Comparison of autoencoder and principal component analysis followed by neural network for e-learning using handwritten recognition. In *IEEE LISAT 2017*.
- Arango, S. P.; Jomaa, H. S.; Wistuba, M.; and Grabocka, J. 2021. Hpo-b: A large-scale reproducible benchmark for black-box hpo based on openml. *arXiv preprint arXiv:2106.06257*.
- Augusto, D. A.; and Barbosa, H. J. 2000. Symbolic regression via genetic programming. In *Proceedings. Vol. 1. Sixth Brazilian symposium on neural networks*.
- Chen, J.; Ma, Z.; Guo, H.; Ma, Y.; Zhang, J.; and Gong, Y.-j. 2024. Symbol: Generating Flexible Black-Box Optimizers through Symbolic Equation Learning. In *ICLR*.
- Chen, M.; Feng, C.; and Cheng, R. 2025. Metade: Evolving differential evolution by differential evolution. *IEEE TEVC*.
- Dietrich, K.; and Mersmann, O. 2022. Increasing the diversity of benchmark function sets through affine recombination. In *International Conference on Parallel Problem Solving from Nature*.
- Doerr, C.; Ye, F.; Horesh, N.; Wang, H.; Shir, O. M.; and Bäck, T. 2019. Benchmarking discrete optimization heuristics with IOHprofiler. In *GECCO*.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*.
- Guo, H.; Ma, S.; Huang, Z.; Hu, Y.; Ma, Z.; Zhang, X.; and Gong, Y.-J. 2025a. Reinforcement learning-based self-adaptive differential evolution through automated landscape feature learning. *arXiv preprint arXiv:2503.18061*.
- Guo, H.; Ma, Y.; Ma, Z.; Chen, J.; Zhang, X.; Cao, Z.; Zhang, J.; and Gong, Y.-J. 2024. Deep Reinforcement Learning for Dynamic Algorithm Selection: A Proof-of-Principle Study on Differential Evolution. *TSMC*.
- Guo, H.; Ma, Z.; Chen, J.; Ma, Y.; Cao, Z.; Zhang, X.; and Gong, Y.-J. 2025b. Configx: Modular configuration for evolutionary algorithms via multitask reinforcement learning. In *AAAI*.
- Guo, H.; Ma, Z.; Ma, Y.; Zhang, X.; Chen, W.-N.; and Gong, Y.-J. 2025c. DesignX: Human-Competitive Algorithm Designer for Black-Box Optimization. *arXiv preprint arXiv:2505.17866*.
- Hansen, N.; Auger, A.; Finck, S.; and Ros, R. 2010a. *Real-parameter black-box optimization benchmarking 2010: Experimental setup*. Ph.D. thesis, INRIA.
- Hansen, N.; Auger, A.; Ros, R.; Finck, S.; and Pošfik, P. 2010b. Comparing results of 31 algorithms from the black-box optimization benchmarking BBOB-2009. In *GECCO*.
- Hooker, J. N. 1995. Testing heuristics: We have it all wrong. *Journal of heuristics*.
- Huang, B.; Wu, X.; Zhou, Y.; Wu, J.; Feng, L.; Cheng, R.; and Tan, K. C. 2025. Evaluation of Large Language Models as Solution Generators in Complex Optimization. *IEEE Computational Intelligence Magazine*.
- Hwang, H.; Vreven, T.; Janin, J.; and Weng, Z. 2010. Protein-protein docking benchmark version 4.0. *Proteins: Structure, Function, and Bioinformatics*.
- Jolliffe, I. 2011. Principal component analysis. In *International encyclopedia of statistical science*.
- Kerschke, P.; and Trautmann, H. 2019. Automated algorithm selection on continuous black-box problems by combining exploratory landscape analysis and machine learning. *Evolutionary computation*.
- Koza, J. 1992. On the programming of computers by means of natural selection. *Genetic programming*.
- Kumar, A.; Wu, G.; Ali, M. Z.; Mallipeddi, R.; Suganthan, P. N.; and Das, S. 2020. A test-suite of non-convex constrained optimization problems from the real-world and some baseline results. *Swarm Evol. Comput.*
- Langdon, W. B.; and Poli, R. 2007. Evolving problems to learn about particle swarm optimizers and other search algorithms. *IEEE TEVC*.
- Lange, R.; Schaul, T.; Chen, Y.; Zahavy, T.; Dalibard, V.; Lu, C.; Singh, S.; and Flennerhag, S. 2023. Discovering evolution strategies via meta-black-box optimization. In *GECCO*.
- Li, D.; Yang, Y.; Song, Y.-Z.; and Hospedales, T. 2018. Learning to generalize: Meta-learning for domain generalization. In *AAAI*.
- Li, X.; Wu, K.; Zhang, X.; and Wang, H. 2025. B2opt: Learning to optimize black-box optimization with little budget. In *AAAI*.
- Li, X.; Wu, K.; Zhang, X.; Wang, H.; Liu, J.; et al. 2024. Pretrained optimization model for zero-shot black box optimization. *Advances in Neural Information Processing Systems*.
- Liu, F.; Tong, X.; Yuan, M.; Lin, X.; Luo, F.; Wang, Z.; Lu, Z.; and Zhang, Q. 2024. Evolution of heuristics: Towards efficient automatic algorithm design using large language model. *arXiv preprint arXiv:2401.02051*.
- Long, F. X.; Vermetten, D.; Kononova, A. V.; Kalkreuth, R.; Yang, K.; Bäck, T.; and van Stein, N. 2023. Challenges of ELA-Guided Function Evolution Using Genetic Programming. In *IJCCI*.
- Ma, Z.; Cao, Z.; Jiang, Z.; Guo, H.; and Gong, Y.-J. 2025a. Meta-Black-Box-Optimization through Offline Q-function Learning. In *ICML*.

- Ma, Z.; Chen, J.; Guo, H.; and Gong, Y.-J. 2025b. Neural Exploratory Landscape Analysis for Meta-Black-Box Optimization. In *ICLR*.
- Ma, Z.; Chen, J.; Guo, H.; Ma, Y.; and Gong, Y.-J. 2024a. Auto-configuring Exploration-Exploitation Tradeoff in Evolutionary Computation via Deep Reinforcement Learning. In *GECCO*.
- Ma, Z.; Gong, Y.-J.; Guo, H.; Qiu, W.; Ma, S.; Lian, H.; Zhan, J.; Chen, K.; Wang, C.; Huang, Z.; et al. 2025c. MetaBox-v2: A Unified Benchmark Platform for Meta-Black-Box Optimization. *arXiv preprint arXiv:2505.17745*.
- Ma, Z.; Guo, H.; Chen, J.; Li, Z.; Peng, G.; Gong, Y.-J.; Ma, Y.; and Cao, Z. 2023. MetaBox: A Benchmark Platform for Meta-Black-Box Optimization with Reinforcement Learning. In *NeurIPS*.
- Ma, Z.; Guo, H.; Chen, J.; Peng, G.; Cao, Z.; Ma, Y.; and Gong, Y.-J. 2024b. Llamoco: Instruction tuning of large language models for optimization code generation. *arXiv preprint arXiv:2403.01131*.
- Ma, Z.; Guo, H.; Gong, Y.-J.; Zhang, J.; and Tan, K. C. 2025d. Toward automated algorithm design: A survey and practical guide to meta-black-box-optimization. *IEEE TEVC*.
- Ma, Z.; Huang, Z.; Chen, J.; Cao, Z.; and Gong, Y.-J. 2025e. Surrogate learning in meta-black-box optimization: A preliminary study. *arXiv preprint arXiv:2503.18060*.
- Mersmann, O.; Bischl, B.; Trautmann, H.; Preuss, M.; Weihs, C.; and Rudolph, G. 2011. Exploratory landscape analysis. In *GECCO*.
- Mo, S.; Wu, K.; Gao, Q.; Teng, X.; and Liu, J. 2025. AutoSGNN: automatic propagation mechanism discovery for spectral graph neural networks. In *AAAI*.
- Mohamed, A. W.; Hadi, A. A.; Mohamed, A. K.; Agrawal, P.; Kumar, A.; and Suganthan, P. N. 2021. Problem definitions and evaluation criteria for the CEC 2021 Special Session and Competition on Single Objective Bound Constrained Numerical Optimization. Technical report.
- Moritz, P.; Nishihara, R.; Wang, S.; Tumanov, A.; Liaw, R.; Liang, E.; Elibol, M.; Yang, Z.; Paul, W.; Jordan, M. I.; et al. 2018. Ray: A distributed framework for emerging {AI} applications. In *OSDI 18*.
- Muñoz, M. A.; and Smith-Miles, K. 2020. Generating new space-filling test instances for continuous black-box optimization. *Evolutionary computation*.
- Prager, R. P.; and Trautmann, H. 2023. Pflacco: Feature-Based Landscape Analysis of Continuous and Constrained Optimization Problems in Python. *Evolutionary Computation*.
- Rajendran, J.; Irpan, A.; and Jang, E. 2020. Meta-learning requires meta-augmentation. *Advances in Neural Information Processing Systems*.
- Renau, Q.; Doerr, C.; Dreo, J.; and Doerr, B. 2020. Exploratory landscape analysis is strongly sensitive to the sampling strategy. In *International Conference on Parallel Problem Solving from Nature*.
- Rumelhart, D. E.; Hinton, G. E.; and Williams, R. J. 1985. Learning internal representations by error propagation. Technical report.
- Schmidt, M.; and Lipson, H. 2009. Distilling free-form natural laws from experimental data. *science*.
- Shao, S.; Tian, Y.; and Zhang, Y. 2025. Deep reinforcement learning assisted surrogate model management for expensive constrained multi-objective optimization. *Swarm and Evolutionary Computation*.
- Shao, S.; Tian, Y.; Zhang, Y.; and Zhang, X. 2025. Knowledge learning-based dimensionality reduction for solving large-scale sparse multiobjective optimization problems. *IEEE TC*.
- Sharma, M.; Komninos, A.; López-Ibáñez, M.; and Kazakov, D. 2019. Deep reinforcement learning based parameter control in differential evolution. In *GECCO*.
- Shehadeh, M. A.; and Kudela, J. 2025. Benchmarking global optimization techniques for unmanned aerial vehicle path planning. *arXiv:2501.14503*.
- Škvorc, U.; Eftimov, T.; and Korošec, P. 2020. Understanding the problem space in single-objective numerical optimization using exploratory landscape analysis. *Applied Soft Computing*.
- Smith-Miles, K.; and Muñoz, M. A. 2023. Instance space analysis for algorithm testing: Methodology and software tools. *ACM Computing Surveys*.
- Stephens, T. 2023. gplearn Documentation. *Release 0.4*.
- Sun, J.; Liu, X.; Bäck, T.; and Xu, Z. 2021. Learning adaptive differential evolution algorithm from optimization experiences by policy gradient. *IEEE TEVC*.
- Tan, Z.; and Li, K. 2021. Differential evolution with mixed mutation strategy based on deep reinforcement learning. *Appl. Soft Comput.*
- Thrun, S.; and Pratt, L. 1998. Learning to learn: Introduction and overview. In *Learning to learn*.
- Tian, Y.; Peng, S.; Zhang, X.; Rodemann, T.; Tan, K. C.; and Jin, Y. 2020. A recommender system for metaheuristic algorithms for continuous optimization based on deep recurrent neural networks. *IEEE TAI*.
- van Stein, B.; Long, F. X.; Frenzel, M.; Krause, P.; Gitterle, M.; and Bäck, T. 2023. Doe2vec: Deep-learning based features for exploratory landscape analysis. In *GECCO*.
- van Stein, N.; and Bäck, T. 2024. Llamea: A large language model evolutionary algorithm for automatically generating metaheuristics. *IEEE TEVC*.
- Vermetten, D.; Ye, F.; Bäck, T.; and Doerr, C. 2025. MABBOB: A problem generator for black-box optimization using affine combinations and shifts. *ACM Transactions on Evolutionary Learning*.
- Wang, Y.; Yao, H.; and Zhao, S. 2016. Auto-encoder based dimensionality reduction. *Neurocomputing*.
- Zhao, Q.; Liu, T.; Yan, B.; Duan, Q.; Yang, J.; and Shi, Y. 2024. Automated metaheuristic algorithm design with autoregressive learning. *IEEE TEVC*.