



Universiteit  
Leiden  
The Netherlands

## **BLADE: Benchmark suite for LLM-driven Automated Design and Evolution of iterative optimisation heuristics**

Stein, N. van; Kononova, A.V.; Yin, H.; Bäck, T.H.W.; Ochoa, G.

### **Citation**

Stein, N. van, Kononova, A. V., Yin, H., & Bäck, T. H. W. (2025). BLADE: Benchmark suite for LLM-driven Automated Design and Evolution of iterative optimisation heuristics. *Gecco '25 Companion: Proceedings Of The Genetic And Evolutionary Computation Conference Companion*, 2336-2344. doi:10.1145/3712255.3734347

Version: Publisher's Version  
License: [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)  
Downloaded from: <https://hdl.handle.net/1887/4295594>

**Note:** To cite this publication please use the final published version (if applicable).



PDF Download  
3712255.3734347.pdf  
22 January 2026  
Total Citations: 1  
Total Downloads: 527

Latest updates: <https://dl.acm.org/doi/10.1145/3712255.3734347>

RESEARCH-ARTICLE

## **BLADE: Benchmark suite for LLM-driven Automated Design and Evolution of iterative optimisation heuristics**

**NIKI VAN STEIN**, Leiden University, Leiden, Zuid-Holland, Netherlands

**ANNA V KONONOVA**, Leiden University, Leiden, Zuid-Holland, Netherlands

**HAORAN YIN**, Leiden University, Leiden, Zuid-Holland, Netherlands

**THOMAS H W BÄCK**, Leiden University, Leiden, Zuid-Holland, Netherlands

**Open Access Support** provided by:

**Leiden University**

**Published:** 14 July 2025

[Citation in BibTeX format](#)

GECCO '25 Companion: Genetic and Evolutionary Computation Conference Companion

July 14 - 18, 2025  
Malaga, Spain

**Conference Sponsors:**  
SIGEVO

# BLADE: Benchmark suite for LLM-driven Automated Design and Evolution of iterative optimisation heuristics

Niki van Stein  
n.van.stein@liacs.leidenuniv.nl  
LIACS, Leiden University  
Leiden, Netherlands

Haoran Yin  
h.yin@liacs.leidenuniv.nl  
LIACS, Leiden University  
Leiden, Netherlands

Anna V. Kononova  
a.kononova@liacs.leidenuniv.nl  
LIACS, Leiden University  
Leiden, Netherlands

Thomas Bäck  
t.h.w.baeck@liacs.leidenuniv.nl  
LIACS, Leiden University  
Leiden, Netherlands

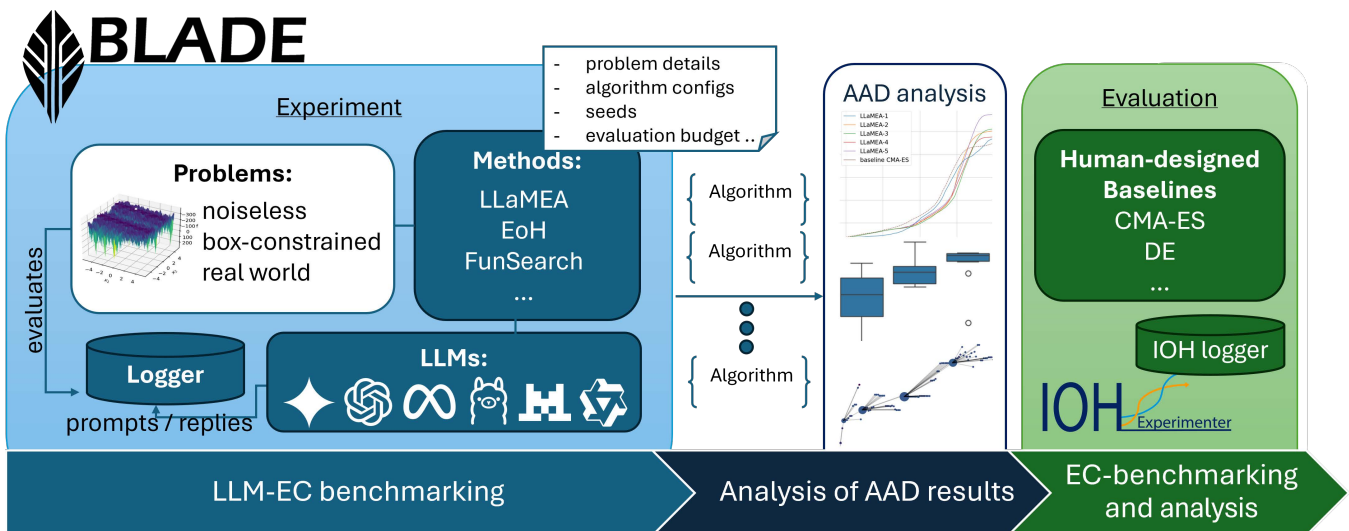


Figure 1: Abstract representation of the BLADE framework for benchmarking LLM-driven Automated Algorithm Discovery for continuous black-box optimisation.

## Abstract

The application of Large Language Models (LLMs) for Automated Algorithm Discovery (AAD), particularly for optimisation heuristics, is an emerging field of research. This emergence necessitates robust, standardised benchmarking practices to rigorously evaluate the capabilities and limitations of LLM-driven AAD methods and the resulting generated algorithms, especially given the opacity of their design process and known issues with existing benchmarks. To address this need, we introduce BLADE (Benchmark suite for LLM-driven Automated Design and Evolution), a modular and extensible framework specifically designed for benchmarking LLM-driven AAD methods in a continuous black-box optimisation

context. BLADE integrates collections of benchmark problems (including MA-BBOB and SBOX-COST among others) with instance generators and textual descriptions aimed at capability-focused testing, such as generalisation, specialisation and information exploitation. It offers flexible experimental setup options, standardised logging for reproducibility and fair comparison, incorporates methods for analysing the AAD process (e.g., Code Evolution Graphs and various visualisation approaches) and facilitates comparison against human-designed baselines through integration with established tools like IOHanalyzer and IOHexplainer. BLADE provides an ‘out-of-the-box’ solution to systematically evaluate LLM-driven AAD approaches. The framework is demonstrated through two distinct use cases exploring mutation prompt strategies and function specialisation.

## CCS Concepts

• **Theory of computation** → Design and analysis of algorithms; Optimization with randomized search heuristics; Continuous optimization; Evolutionary algorithms; • **Computing methodologies** → Heuristic function construction.



This work is licensed under a Creative Commons Attribution 4.0 International License.  
GECCO '25 Companion, July 14–18, 2025, Malaga, Spain  
© 2025 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-1464-1/2025/07  
<https://doi.org/10.1145/3712255.3734347>

## Keywords

Large Language Models, Automated Algorithm Design, Benchmarking, Evolution Strategies, Black-Box Optimization

### ACM Reference Format:

Niki van Stein, Anna V. Kononova, Haoran Yin, and Thomas Bäck. 2025. BLADE: Benchmark suite for LLM-driven Automated Design and Evolution of iterative optimisation heuristics. In *Genetic and Evolutionary Computation Conference (GECCO '25 Companion)*, July 14–18, 2025, Malaga, Spain. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3712255.3734347>

## 1 Introduction

Large Language Models (LLMs) are increasingly being used as automated algorithm designers in optimisation [15]. In this emerging paradigm, an LLM's knowledge and reasoning ability are combined with traditional optimisation frameworks to discover new algorithms or improve existing ones. Recent studies have rapidly proliferated, showing that LLMs can serve as algorithm generators for various optimisation tasks [14, 16, 23, 24, 27].

This field has become known as automatic algorithm discovery (AAD), whose goal is to have LLMs propose heuristic strategies or (snippets of) executable code for solving optimisation problems with minimal human input. It is widely accepted that the performance of optimisation algorithms depends on the optimisation problem, or, in other words, algorithms specialise on certain problem classes, following the result widely known as the No Free Lunch theorem [37]. Establishing the strengths and weaknesses of algorithms is a task of benchmarking, which has become a non-negotiable step in modern algorithm development, underpinned by established best practices [2, 18]. More specifically, *benchmarking* encompasses a systematic evaluation of algorithm performance on a curated set of problem instances, using defined metrics to assess and compare efficiency, accuracy and robustness across varying conditions or landscapes. Benchmarking becomes even more critical for automatically discovered algorithms, whose construction is opaque and often defies intuitive explanation.

It is essential to judiciously select the optimisation problems included in a benchmarking suite, as they must reflect the intended objectives of benchmarking: they should be numerous—without rendering the effort impractical—and offer *sufficient diversity* within the problem class targeted by the algorithm. A similar rationale applies to the selection of performance measures, which should align with the benchmarking objectives. For this reason, a benchmarking framework with a *modular design* is advisable, enabling components such as problem sets and performance measures to be easily swapped in or out depending on the goals of each experiment.

Unfortunately, concepts such as diversity and class of problems, which are essential for defining benchmarking setups in practice, remain elusive when it comes to rigorous formalisation. Nevertheless, researchers continue to invest considerable effort in this active area. Providing flexible tools that facilitate easy experimentation with benchmarking setups is therefore highly beneficial.

In this paper, we propose **BLADE: Benchmark suite for LLM-driven Automated Design and Evolution of iterative optimisation heuristics**, which serves the purpose. An overview of BLADE is presented in Figure 1. As demonstrated throughout the paper, BLADE includes:

- (1) several collections of single-objective continuous optimisation problems with instance generation mechanisms, which are aimed at studying capabilities of algorithms such as, e.g., generalisation over problem instances and dimensionalities, levels of specialisation and efficiency of domain knowledge integration;
- (2) flexible specification of benchmarking setup used inside AAD, such as, e.g., run budgets, number of runs and instances used for training and testing (allowing to evolve the algorithm on one set of problem instances and then test on a novel instance to truly assess generalisation);
- (3) a standardized interface to connect with many different LLMs and postprocess the generated output to get code and descriptions of solutions;
- (4) methods for the analysis of the results of the algorithm discovery process, such as, e.g., Code Evolution graphs [28] and ELO ratings [1];
- (5) methods for performance comparison of discovered algorithms against human-designed or other specified algorithms via native integration with tools like IOHanalyzer [36] and IOHexplainer [30].

Thanks to BLADE's modular design, all these components can be readily extended. At the same time, BLADE works 'out-of-the-box', using loggers that record all necessary information by default, which is then passed to the specified AAD and evolutionary computation benchmarking analysis tools. The current focus on continuous black-box optimisation stems both from our prior expertise and from a clear gap in the existing literature.

This paper is organised as follows: Section 2 summarises relevant related work, Section 3 outlines full paper methodology, Section 4 discusses two use-cases selected to demonstrate the usage of BLADE, Section 5 provides a reproducibility and transparency statement and Section 6 concludes with final remarks and future directions.

## 2 Related work

*From AS to AAD to AAD benchmarking.* Automatic algorithm discovery is rooted in algorithm selection (AS) approaches, which initially aimed to identify well-performing algorithms from a given portfolio of (parameterised) algorithms for a specific problem instance. Traditional approaches to AS relied on performance prediction models, often informed by exploratory landscape features and used tools such as irace [20] and SMAC [13]. The introduction of modular frameworks such as modular CMA-ES (modCMA) [4] and modular Differential Evolution (modDE) [31] extended these methods by enabling fine-grained configuration of algorithmic components.

While these tools significantly advanced the state of the art, they still required substantial expert input, both in defining the portfolio of candidate algorithms and in engineering suitable instance features. With the advent of powerful Large Language Models, the field has begun shifting from selecting among pre-existing algorithms to automatically generating novel algorithms—an approach now commonly referred to as Automatic Algorithm Discovery.

In this emerging paradigm, LLMs are prompted to propose algorithmic strategies, often expressed as pseudocode or directly

executable code, that are then evaluated using traditional optimisation benchmarks. AAD frameworks such as FunSearch [22], LLaMEA [27] and the Evolution of Heuristics (EoH) framework [14] exemplify this direction. These systems enable LLMs to iteratively improve their outputs by incorporating feedback from fitness evaluations. The LLaMEA framework, in particular, has seen several recent extensions, including variants that incorporate hyper-parameter optimisation [29], mutation control [38] and real-world problem specialisation [39].

A significant body of novel research on AAD [15] brought the need for a platform facilitating algorithm discovery using large language models. LLM4AD [16] has thus been proposed as a broad modular meta-framework designed to support the general development of LLM-driven algorithm discovery methods in a standardised evaluation environment. LLM4AD is primarily focused on combinatorial optimisation and is in its early development stage.

*Issues with existing benchmarking setups.* The de-facto standard for continuous optimisation is the BBOB (Black-Box optimisation Benchmark) suite [8], part of the COCO platform [7], which provides a set of 24 functions defined for any dimensionality, with varying properties (unimodal to highly multimodal, separable vs. non-separable, etc.) and allows generating multiple instances per function (by shifting or rotating the coordinate system). By using BBOB, researchers ensure generality in evaluation – an algorithm must handle rugged, smooth, separable and non-separable problems, rather than being overfitted to one test function.

While representing a definitive step forward in advancing benchmark methodology, BBOB comes with a range of problems:

- Ambiguous formulation of the optimisation problem (unconstrained vs box-constrained) and insufficient diversity of problem instances in terms of locations of optima within the search domain of BBOB problems [17]. To remedy this, SBOX-COST function suite has been proposed [33].
- Models trained on BBOB data generalise poorly to other suites such as those provided by the CEC conference [10–12] due to statistically significant differences in feature-space distributions [21]. This stems from the unintended interpretation of the generality of functions included in BBOB.
- Similarly, MA-BBOB [35] which extends BBOB by generating a wide array of benchmark functions through affine combinations of BBOB functions, the resulting problem landscapes may not exhibit sufficient diversity [5]. This limitation poses challenges for automated algorithm selection methods, as training on these generated functions does not necessarily lead to effective generalisation to novel problem instances. The authors emphasize the need for careful selection of training instances to ensure robust and generalizable algorithm selection models [5].

Furthermore, results in combinatorial optimisation [23] demonstrate that most heuristics discovered via LLMs also fail to generalise across diverse problem instances, performing well only in specific areas of the instance space. In contrast, traditional simple heuristics demonstrate in this case more consistent performance across a broader range of benchmarks. Thus, the authors advocate for the development and utilisation of more diverse benchmarks and applications to gain a better understanding of this emerging

paradigm in AAD [23]. Performance analysis on the results of AAD also suggests that benchmark suites for bin packing problems are simply too easy [23, 29].

### 3 Methodology

This section details the methodology of BLADE for benchmarking LLM-driven Automated Algorithm Discovery. The framework, as depicted in Figure 1, encompasses several key components: benchmark problems, AAD methods, LLMs and evaluation metrics. The BLADE framework uses the following key steps in the benchmarking pipeline:

*Experimental setup.* First, an experiment is designed, or one of the provided experimental setups is used directly for easy comparison with other works. An experiment consists of a list of search methods (such as LLaMEA, EoH, FunSearch, Random Search, ReEvo and others), a list of optimisation problems and an LLM. In addition, the experiment has several attributes such as the number of algorithms per run to evaluate (AAD budget), the number of independent runs per method/problem combination and some other settings. All experiment settings are logged for reproducibility.

*LLM-EC Benchmarking.* Next, the experiment(s) is/are carried out using the experimental setup provided. BLADE uses parallel processes to speed up the evaluations of the different runs over different CPU threads. *Loggers* are attached to the LLM and Problem objects to log all LLM queries made and all generated algorithm evaluations. The loggers operate independently of the search method and, as such, enforce a fair comparison between LLM-EC (search) algorithms. The LLM token costs, seeds and other hyperparameters are logged for transparency and reproducibility. On the optimisation problem side, each problem evaluation is logged. Problems consist of several (predefined) training and testing instances, where the training instances are used during the LLM-EC runs and the test instances are used for the final evaluation of best best-performing generated algorithms.

*Analysis of AAD results.* After the different LLM-EC runs, the convergence of the different algorithms can be analysed by visualising the best-so-far convergence curves averaged over all random seeds. In addition, we can analyse in more detail how the code evolved in one or more of the runs by looking at the Code Evolution Graphs (CEG) [28]. The CEG shows a low-dimensional embedding of various static code features against the number of algorithms generated.

*EC Benchmarking and Analysis.* The final step in the BLADE benchmarking toolbox is the evaluation of the best found algorithms on separate *validation instances* of the problem. These instances can either be larger versions of real-world problems or different instances of the same optimisation problems that the LLM-EC has used during the LLM-EC benchmarking phase. The best algorithms found by each LLM-driven AAD method are compared against each other and also against the state-of-the-art human-designed algorithms such as CMA-ES [9], Differential Evolution variants [6] and others. The final analysis of the comparison depends on the problem specification. For example, we can look at the optimal fitness reached per algorithm or at anytime performance metrics such as the Area

Over the Convergence Curve (AOCC) [19] and the Empirical Attainment Function (EAF) [19], which give better insight into the effectiveness of the different algorithms over different evaluation budgets. For the visualisation and logging of the final evaluation, BLADE is integrated with IOAnalyser [36] and IOHexplainer [30], the established benchmarking tools of the black-box optimisation field.

### 3.1 Capability-Focused Benchmarking

To move beyond generic black-box optimisation evaluations, focusing purely on benchmarking performance and addressing the limitations highlighted by the No Free Lunch theorem [37], we advocate for *capability-focused benchmarking*. This approach allows for targeted experiments designed to answer specific research questions regarding the performance of AAD methods, particularly those driven by LLMs. We aim to assess the following key capabilities:

- **Generalization:** Evaluating an AAD method's ability to generate optimisation algorithms effective across a diverse range of problems within a given dimensionality.
- **Problem Class Specialization:** Assessing an AAD method's proficiency in evolving solvers tailored for specific classes of problems, e.g., such as those characterised by multi-modality with weak global structure.
- **Information Exploitation:** Determining an AAD method's capacity to leverage problem-specific details (e.g., textual descriptions of landscape features) to construct specialised solvers for particular problem instances.

These capabilities reflect the potential of LLM-EC methods to produce solvers ranging from general-purpose tools to highly specialised algorithms for specific problem domains or individual problems (per-instance configuration). To rigorously test these capabilities, we propose utilising the following benchmark suites:

- **Many Affine Black-Box optimisation Benchmark (MA-BBOB):** This configurable suite comprises functions derived from BBOB [8], originally proposed as part of the comparing continuous optimizers (COCO) environment [7], via affine transformations [34]. We recommend using a predefined set of 20 training and 50 testing instances and provide in total 1000 predefined instances. A key feature of MA-BBOB is the random distribution of optima within the search domain, mitigating the risk of LLMs succeeding through memorisation or biased initialisation strategies [32] (a known issue in some of the original BBOB functions [17]). MA-BBOB is primarily used to assess generalisation capabilities.
- **SBOX-COST:** This benchmark suite features box-constrained black-box optimisation problems [33], also with randomly distributed optima across the full domain. It includes five distinct function groups, each equipped with an instance generator (applying rotations and shifts). We included textual descriptions of the underlying optimisation problem landscapes. SBOX-COST is employed for evaluating group-level and function-level specialization, including the ability to utilize provided problem (or problem class) information.

These synthetic benchmarking suites have been chosen because of their recognition in the field (they are based on the popular

BBOB suite, but without the center bias of the location of optima) and because they are designed to test different capabilities in optimisation algorithms. These capabilities are for example, how well the algorithms can escape local optima, how well they can exploit global structure, etc.

### 3.2 Real world applications

Next to the synthetic benchmarking functions that are used to test various capabilities in LLM-driven AAD methods, several *real-world problems* are included (and more will be added) to BLADE to verify the capabilities of AAD methods in practical domains.

In [39], LLaMEA is used to solve real-world photonic structure optimisation problems from [3], including Bragg mirror problems, inverse ellipsometry problems and photovoltaic design problems, which have practical implications for communication, semiconductors, LED displays, materials analysis and solar cells [39].

Since the evaluation of these (and many other) real-world problems is very time-consuming, the AAD methods should use smaller-scale versions of these problems (with a lower dimensionality) as *training instances* and *validate* the best found algorithms on the high-dimensional harder problems. In addition, background knowledge on the real-world problems is typically available (from physics) to be used inside the prompt by the LLM-driven AAD methods. Initial results in [39] confirmed that this knowledge usually has a positive effect on the quality of the algorithms generated, however, this likely depends on the LLM and on the real-world problem and requires deeper investigation.

## 4 Use-cases


To effectively show how the proposed BLADE toolbox works, we include two use-cases. The first use-case compares different mutation prompting strategies and uses analysis on static code features to discover the effect of different prompting techniques. In addition, it shows the BLADE benchmarking pipeline step-by-step. The second use-case shows how to test for function specialisation where the LLMs can exploit additional information about the problems to solve given in the task prompt for different types of problems. The first use-case focuses more on the EC part of the LLM-driven AAD process, while the latter use-case focuses more on LLM capabilities, showing two (of many) different approaches possible with the BLADE framework.

### 4.1 Use-case: Mutation Prompts

This use-case investigates the impact of different mutation prompt strategies on the effectiveness of LLM-driven algorithm discovery using the LLaMEA [27] framework with a (4, 12) evolution strategy (4 parents, 12 offspring).

The central research question addressed is: *Which combination of mutation prompts is most efficient for generating general-purpose continuous black-box optimizers when using a specific LLM (Gemini-2.0-flash)?*

**4.1.1 Experimental Setup.** The experiment was conducted using the 5-dimensional MA-BBOB benchmark suite [34], chosen for its diverse set of problems with randomly distributed optima. The LLM-EC process utilised 20 predefined MA-BBOB training instances,



**Figure 2: Distribution of the fitness (AOCC) of algorithms resulting from each AAD run. The AOCC is obtained by evaluating each of the resulting algorithms (10 algorithms per LLaMEA configuration) on 50 evaluation instances of MA-BBOB and 10 random seeds.**

while the final evaluation was performed on 50 distinct test instances. Each LLaMEA-EC run had a budget of generating (and evaluating) 100 candidate algorithms and each algorithm evaluation within the EC process had a budget of  $2000 \times \text{dimensionality} = 10,000$  function evaluations.

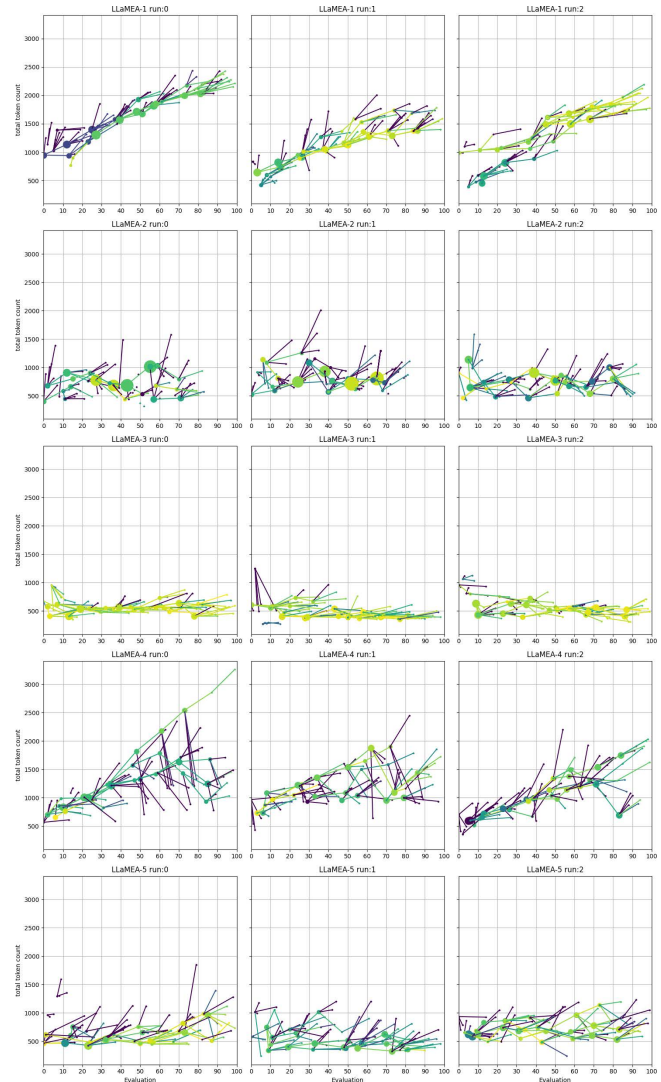
Three distinct mutation prompts were defined:

- (1) "Refine the strategy of the selected algorithm to improve it"
- (2) "Generate a new algorithm that is different from the algorithms you have tried before"
- (3) "Refine and simplify the selected algorithm to improve it"

Five LLaMEA configurations were tested, each employing different combinations of these mutation prompts:

- LLaMEA-1: Uses only prompt 1
- LLaMEA-2: Uses only prompt 2
- LLaMEA-3: Uses only prompt 3
- LLaMEA-4: Uses prompts 1 and 2 (chosen randomly)
- LLaMEA-5: Uses all three prompts (chosen randomly)


**4.1.2 AAD Evaluation Results.** Figure 2 illustrates the distribution of fitness values, measured by the Area Over the Convergence Curve (AOCC), for the algorithms produced by each LLaMEA configuration during the automated design phase. Each boxplot represents the AOCC values obtained by evaluating the 10 algorithms generated (by 10 different runs) of each respective LLaMEA variant on the 50 MA-BBOB evaluation instances, aggregated over 10 independent seeds (500 evaluation runs per algorithm). Visual inspection suggests variability in the effectiveness of different prompt strategies, with configurations using prompt 3 ("simplify") (LLaMEA-3 and LLaMEA-5) showing potential for generating higher-performing algorithms compared to the default prompt 1 "refine" option provided by vanilla LLaMEA. In addition, when we look at the Code Evolution Graphs in Figure 3, we can clearly observe that including prompt 3 (LLaMEA-3 and 5) keeps the size of the code



**Figure 3: Code Evolution Graphs of 3 runs (columns) per LLaMEA variant (rows) showing how the number of Python tokens (y-axis) changes per generated algorithm (x-axis). The colour of the nodes denotes the AOCC score (yellow is higher and better, blue is lower and worse), while connections between nodes denote parent-offspring relations and the size of nodes encodes the number of connections produced from this node. Runs are sometimes showing fewer than 100 evaluations due to runtime errors inside the generated algorithms (resulting in a negative fitness which is not displayed).**

small. It can also be observed that adding prompt 2 (LLaMEA-2, 4 and 5) increases the diversity in code substantially.

**4.1.3 Validation Against Human-Designed Baseline.** To assess the quality of the generated algorithms relative to established methods, the best-performing algorithm discovered by each of the five LLaMEA configurations (selected based on performance across 10 independent AAD runs and their training instance performance)



**Figure 4: Empirical Attainment Function (EAF) curves of the best solutions found per LLaMEA configuration (best of 10 runs) and a CMA-ES baseline. Each algorithm is tested on 50 different MA-BBOB instances and 10 independent runs.**


was further evaluated. These selected algorithms were compared against CMA-ES, a state-of-the-art human-designed evolutionary algorithm, serving as a strong baseline [9]. The evaluation involved running each selected algorithm and the baseline on the 50 MA-BBOB test instances over 10 independent runs.

The performance comparison is visualised using Empirical Attainment Function (EAF) curves and ELO ratings [1]. Figure 4 displays the EAF curves, illustrating the probability of reaching a certain target function value within a given budget of function evaluations [19]. This provides insight into the anytime performance characteristics of the algorithms. The EAF curves show that the best algorithms generated by several LLaMEA configurations, particularly LLaMEA-5, are competitive with or outperform the CMA-ES baseline, especially in the later stages of the search corresponding to higher evaluation budgets.

Figure 5 presents the ELO ratings calculated from pairwise comparisons between the best algorithm from each LLaMEA configuration and the CMA-ES baseline. The ELO rating provides a single metric summarising the relative strength of each algorithm based on tournament-style comparisons across the 50 test instances and 10 runs using a tournament size of 100 000 pair-wise comparisons. The results indicate that LLaMEA-5 achieved the highest ELO rating, suggesting that utilising a combination of all three mutation prompts yielded the most robust algorithm overall within this experimental setup, outperforming other configurations and the CMA-ES baseline. LLaMEA-2 and LLaMEA-4 also demonstrate strong performance relative to the baseline according to this metric.

## 4.2 Use-case: LLM Comparison

This use-case aims to evaluate and compare the effectiveness of different LLMs in the context of AAD for black-box optimisation. Specifically, it seeks to determine *which LLM is most proficient in*



**Figure 5: ELO rating (higher is better) of the best solutions found per LLaMEA configuration (best of 10 runs) and a CMA-ES baseline. Each algorithm is tested on 50 different MA-BBOB instances and 10 independent runs. ELO rating is calculated using a tournament of 100 000 one against one comparisons.**

*generating optimizers tailored to a specific type of function or problem class, focusing on specialisation in landscape function properties.*

The experiments involve using LLaMEA with the following LLMs: local models codestral (22b) and qwen2.5-coder (14b) and closed-source models gemini-1.5-flash [25] and the more recently introduced gemini-2.0-flash <sup>1</sup>.

The experiment is comprised of two parts. The per-function specialisation and the per-function-group specialisation. In the *per-function* specialisation part, the AAD methods are given 5 instances of a particular optimisation problem including a textual description of the problem, such as “Separable Ellipsoidal Function” in the case of function 2 (*fid2*). The best resulting algorithm per run is then evaluated using 10 (different) evaluation instances of the same problem. This is repeated for different functions (with different characteristics), in this case we choose *fid2* (Separable Ellipsoidal Function), *fid5* (Linear Slope), *fid13* (Sharp Ridge Function), *fid15* (Rastrigin Function) and *fid21* (Gallagher’s Gaussian 101-me Peaks Function) as these functions form a very diverse set. In the *per-function-group* part of the experiment, the LLM-driven AAD method is given 5 instances of all problems in a particular function group (4 or 5 problems per group). The groups are: group1 “Separable Functions”, group2 “Functions with low or moderate conditioning”, group3 “Functions with high conditioning and unimodal”, group4 “Multi-modal functions with adequate global structure” and group5 “Multi-modal functions with weak global structure”. The final algorithms are validated on 10 different instances for each problem in the same group.

For both parts of the experiment, each LLM-EC run had a budget of generating 50 candidate algorithms, and each algorithm evaluation within the EC process had a budget of  $2000 \times d(5) = 10,000$  function evaluations. The performance of these LLMs in combination with the LLaMEA framework is assessed based on the fitness (AOCC) of the final solutions they generate, as illustrated in Figure 6 and Table 1 (higher is better). These results are derived from

<sup>1</sup>Specifically codestral <https://ollama.com/library/codestral:22b> and qwen2.5-coder <https://ollama.com/library/qwen2.5-coder:14b> from the ollama repository.

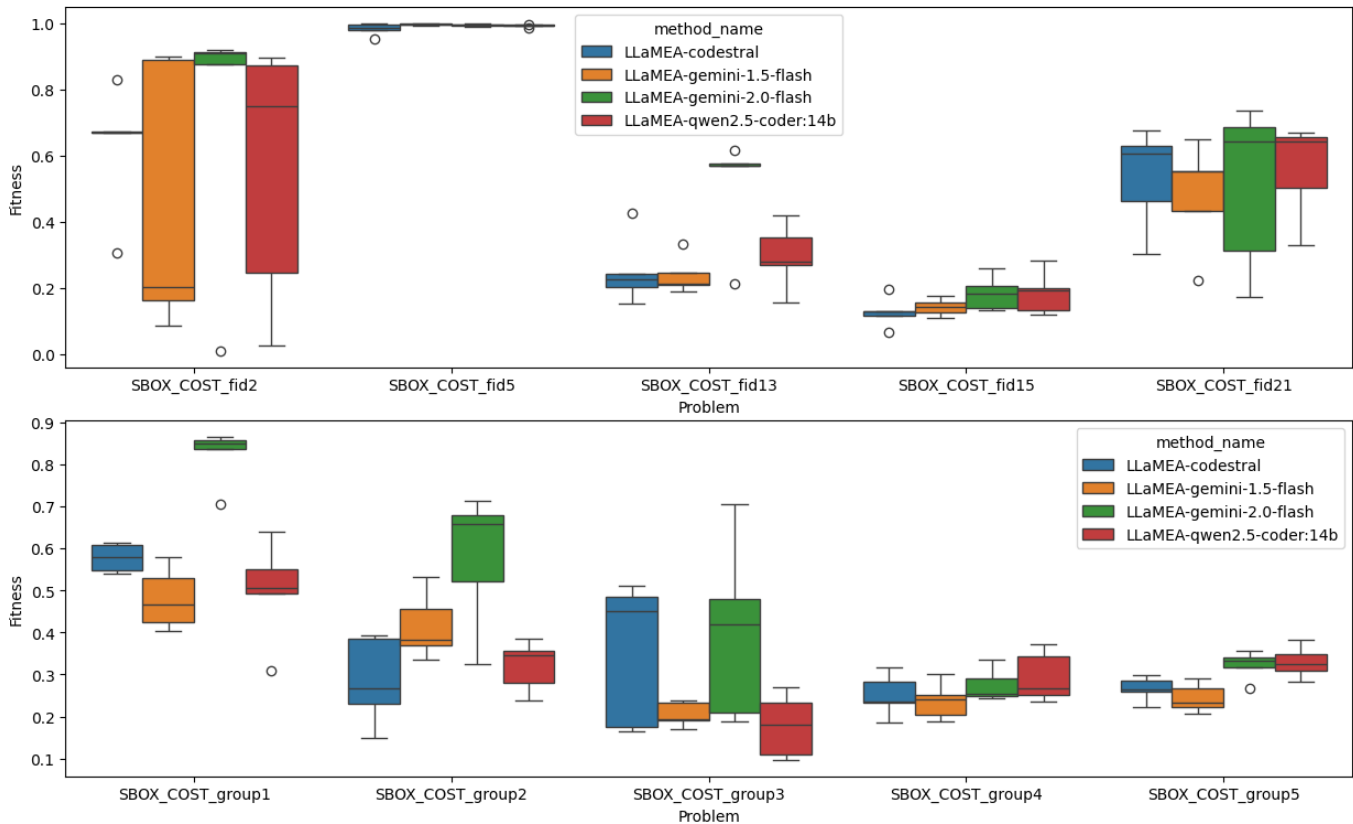


Figure 6: Fitness (AOCC) distribution of all final solutions generated using LLaMEA with different LLMs.

Table 1: Mean AOCC for each LLM-EC over the final solutions of all runs per problem. Results in boldface are significantly better with a threshold of 0.05, p-values below this threshold are reported.

	$f_2$	$f_5$	$f_{13}$	$f_{15}$	$f_{21}$
codestral	0.63 ± 0.17	0.98 ± 0.02	0.25 ± 0.09	0.13 ± 0.04	0.53 ± 0.14
gemini-1.5-flash	0.45 ± 0.37	1.00 ± 0.00	0.24 ± 0.05	0.14 ± 0.02	0.48 ± 0.15
gemini-2.0-flash	0.72 ± 0.36	0.99 ± 0.00	<b>0.51 ± 0.15</b> (p=0.047)	0.18 ± 0.05	0.51 ± 0.22
qwen2.5-coder:14b	0.56 ± 0.36	0.99 ± 0.00	0.29 ± 0.09	0.18 ± 0.06	0.56 ± 0.13

	Group 1	Group 2	Group 3	Group 4	Group 5
codestral	0.58 ± 0.03	0.29 ± 0.09	0.36 ± 0.15	0.25 ± 0.05	0.27 ± 0.03
gemini-1.5-flash	0.48 ± 0.07	0.41 ± 0.07	0.21 ± 0.03	0.24 ± 0.04	0.24 ± 0.03
gemini-2.0-flash	<b>0.82 ± 0.06</b> (p=0.002)	0.58 ± 0.14	0.40 ± 0.19	0.27 ± 0.03	0.32 ± 0.03
qwen2.5-coder:14b	0.50 ± 0.11	0.32 ± 0.05	0.18 ± 0.07	0.29 ± 0.05	0.33 ± 0.03

the validation instances, providing insight into the generalisation capabilities of the algorithms produced by each LLM.

In Figure 6 we can observe that gemini-2.0-flash is outperforming other LLMs mostly for unimodal and low conditioning functions (fid2, group1, fid13, group2). Other models show relatively similar performance and it is interesting to note that the *smaller open-source models do not underperform significantly* in most cases. From Table 1 we can further see that gemini-2.0-flash is only significantly better

(using a pairwise independent t-test) in two of the ten problem setups.

## 5 Reproducibility and Transparency

BLADE is available open-source and fully documented on Github<sup>2</sup>. All code and results of the experiments in this paper, including full prompts, generated algorithms and seeds are available in our Zenodo repository [26].

## 6 Conclusions and Outlook

This paper introduces **BLADE**, a novel benchmarking suite designed to evaluate and compare LLM-driven Automated Algorithm Discovery (AAD) methods for continuous black-box optimisation. BLADE provides a structured framework encompassing various black-box benchmark problem sets, including real-world applications and synthetic suites tailored for capability-focused benchmarking. The *modular design* of BLADE facilitates the integration of different AAD methods, LLMs and evaluation metrics, enabling comprehensive and reproducible experiments.

The use-cases presented in this paper demonstrate the utility of BLADE in analysing the impact of different prompt strategies on LLM-driven algorithm generation and in comparing the performance of various LLMs within the AAD context. The results

<sup>2</sup><https://github.com/XAI-liacs/BLADE>

highlight the potential of LLM-driven AAD to produce competitive optimizers while also revealing the influence of prompt engineering and LLM selection on the effectiveness of the generated algorithms.

Future work will focus on expanding the benchmark suite with additional real-world problems to assess other relevant capabilities of LLM-driven AAD methods. Furthermore, we plan to incorporate a wider range of baseline algorithms and explore automated hyperparameter optimisation within the benchmarking framework.

## References

- [1] Paul CH Albers and Han de Vries. 2001. Elo-rating as a tool in the sequential estimation of dominance strengths. 489–495 pages.
- [2] Thomas Bartz-Beielstein, Carola Doerr, Daan van den Berg, Jakob Bossek, Sowmya Chandrasekaran, Tome Eftimov, Andreas Fischbach, Pascal Kerschke, William La Cava, Manuel Lopez-Ibanez, Katherine M. Malan, Jason H. Moore, Boris Naujoks, Patryk Orzechowski, Vanessa Volz, Markus Wagner, and Thomas Weise. 2020. Benchmarking in Optimization: Best Practice and Open Issues. arXiv:2007.03488 [cs.NE] <https://arxiv.org/abs/2007.03488>
- [3] Pauline Bennet, Denis Langevin, Chayma Essoual, Abdourahman Khairah-Walieh, Olivier Teytaud, Peter Wiecha, and Antoine Moreau. 2024. Illustrated tutorial on global optimization in nanophotonics. *JOSA B* 41, 2 (2024), A126–A145.
- [4] Jacob de Nobel, Diederick Vermetten, Hao Wang, Carola Doerr, and Thomas Bäck. 2021. Tuning as a means of assessing the benefits of new ideas in interplay with existing algorithmic modules. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO'21, Companion material)*. ACM, 1375–1384. doi:10.1145/3449726.3463167
- [5] Konstantin Dietrich, Diederick Vermetten, Carola Doerr, and Pascal Kerschke. 2024. Impact of Training Instance Selection on Automated Algorithm Selection Models for Numerical Black-box Optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference (Melbourne, VIC, Australia) (GECCO '24)*. Association for Computing Machinery, New York, NY, USA, 1007–1016. doi:10.1145/3638529.3654100
- [6] Vitaliy Feoktistov. 2006. *Differential evolution*. Springer.
- [7] Nikolaus Hansen, Anne Auger, Raymond Ros, Olaf Mersmann, Tea Tušar, and Dimo Brockhoff. 2021. COCO: a platform for comparing continuous optimizers in a black-box setting. *Optimization Methods and Software* 36, 1 (2021), 114–144. doi:10.1080/10556788.2020.1808977
- [8] Nikolaus Hansen, Steffen Finck, Raymond Ros, and Anne Auger. 2009. *Real-Parameter Black-Box Optimization Benchmarking 2009: Noiseless Functions Definitions*. Research Report RR-6829. INRIA. <https://hal.inria.fr/inria-00362633>
- [9] Nikolaus Hansen, Sibylle D Müller, and Petros Koumoutsakos. 2003. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolutionary computation* 11, 1 (2003), 1–18.
- [10] J. Liang, B. Qu, and P. Suganthan. 2013. *Problem Definitions and Evaluation Criteria for the CEC 2014 Special Session and Competition on Single Objective Real-Parameter Numerical Optimization*. Technical Report. Computational Intelligence Laboratory, Zhengzhou University, Zhengzhou, China and Technical Report, Nanyang Technological University, Singapore.
- [11] J. Liang, B. Qu, P. Suganthan, and Q. Chen. 2014. *Problem Definitions and Evaluation Criteria for the CEC 2015 Competition on Learning-Based Real-Parameter Single Objective Optimization*. Technical Report. Computational Intelligence Laboratory, Zhengzhou University, Zhengzhou, China and Technical Report, Nanyang Technological University, Singapore.
- [12] J. Liang, B. Qu, P. Suganthan, and A. Hernández-Díaz. 2013. *Problem Definitions and Evaluation Criteria for the CEC 2013 Special Session on Real-Parameter Optimization*. Technical Report. Computational Intelligence Laboratory, Zhengzhou University, Zhengzhou, China and Technical Report, Nanyang Technological University, Singapore.
- [13] Marius Lindauer, Katharina Eggensperger, Matthias Feurer, André Biedenkapp, Difan Deng, Carolin Benjamins, Tim Rühkopf, René Sass, and Frank Hutter. 2022. SMAC3: A Versatile Bayesian Optimization Package for Hyperparameter Optimization. *Journal of Machine Learning Research* 23, 54 (2022), 1–9. <http://jmlr.org/papers/v23/21-0888.html>
- [14] Fei Liu, Tong Xialiang, Mingxuan Yuan, Xi Lin, Fu Luo, Zhenkun Wang, Zhichao Lu, and Qingfu Zhang. 2024. Evolution of Heuristics: Towards Efficient Automatic Algorithm Design Using Large Language Model. In *Forty-first International Conference on Machine Learning*.
- [15] Fei Liu, Yiming Yao, Ping Guo, Zhiyuan Yang, Zhe Zhao, Xi Lin, Xialiang Tong, Mingxuan Yuan, Zhichao Lu, Zhenkun Wang, and Qingfu Zhang. 2024. A Systematic Survey on Large Language Models for Algorithm Design. arXiv:2410.14716 [cs.LG] <https://arxiv.org/abs/2410.14716>
- [16] Fei Liu, Rui Zhang, Zhuoliang Xie, Rui Sun, Kai Li, Xi Lin, Zhenkun Wang, Zhichao Lu, and Qingfu Zhang. 2024. LLM4AD: A Platform for Algorithm Design with Large Language Model. arXiv:2412.17287 [cs.AI] <https://arxiv.org/abs/2412.17287>
- [17] Fu Xing Long, Diederick Vermetten, Bas van Stein, and Anna V. Kononova. 2023. BBOB Instance Analysis: Landscape Properties and Algorithm Performance Across Problem Instances. In *Applications of Evolutionary Computation*, João Correia, Stephen Smith, and Raneem Qaddoura (Eds.). Springer Nature Switzerland, Cham, 380–395.
- [18] Manuel López-Ibáñez, Juergen Branke, and Luis Paquete. 2021. Reproducibility in Evolutionary Computation. *ACM Trans. Evol. Learn. Optim.* 1, 4, Article 14 (Oct. 2021), 21 pages. doi:10.1145/3466624
- [19] Manuel López-Ibáñez, Diederick Vermetten, Johann Dreö, and Carola Doerr. 2024. Using the empirical attainment function for analyzing single-objective black-box optimization algorithms. *IEEE Transactions on Evolutionary Computation* (2024).
- [20] Manuel López-Ibáñez, Jérémie Dubois-Lacoste, Leslie Pérez Cáceres, Mauro Bittaratti, and Thomas Stützle. 2016. The irace package: Iterated racing for automatic algorithm configuration. *Operations Research Perspectives* 3 (2016), 43–58. doi:10.1016/j.orp.2016.09.002
- [21] Ana Nikolić, Ana Kostovska, Gjorgjina Cenikj, Carola Doerr, and Tome Eftimov. 2024. Generalization Ability of Feature-Based Performance Prediction Models: A Statistical Analysis Across Benchmarks. In *IEEE Congress on Evolutionary Computation, CEC 2024, Yokohama, Japan, June 30 - July 5, 2024*. IEEE, 1–8. doi:10.1109/CEC60901.2024.10611952
- [22] Bernardino Romera-Paredes, Mohammadamin Barekatin, Alexander Novikov, Matej Balog, M Pawan Kumar, Emilien Dupont, Francisco JR Ruiz, Jordan S Ellenberg, Pengming Wang, Omar Fawzi, Pushmeet Kohli, and Alhussein Fawzi. 2024. Mathematical discoveries from program search with large language models. *Nature* 625 (01 2024), 468–475. Issue 7995.
- [23] Kevin Sim, Quentin Renau, and Emma Hart. 2025. Beyond the Hype: Benchmarking LLM-Evolved Heuristics for Bin Packing. arXiv:2501.11411 [cs.NE] <https://arxiv.org/abs/2501.11411>
- [24] Yiwen Sun, Furong Ye, Xianyin Zhang, Shiyu Huang, Bingzhen Zhang, Ke Wei, and Shaowei Cai. 2024. AutoSAT: Automatically Optimize SAT Solvers via Large Language Models. arXiv:2402.10705 [cs.AI] <https://arxiv.org/abs/2402.10705>
- [25] Gemini Team. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. arXiv:2403.05530 [cs.CL] <https://arxiv.org/abs/2403.05530>
- [26] Niki van Stein. 2025. *BLADE - Code and Results for the paper*. doi:10.5281/zenodo.15119985
- [27] Niki van Stein and Thomas Bäck. 2024. LLaMEA: A Large Language Model Evolutionary Algorithm for Automatically Generating Metaheuristics. *IEEE Transactions on Evolutionary Computation* (2024), 1–1. doi:10.1109/TEVC.2024.3497793
- [28] Niki van Stein, Anna V. Kononova, Lars Kotthoff, and Thomas Bäck. 2025. Code Evolution Graphs: Understanding Large Language Model Driven Design of Algorithms. arXiv:2503.16668 [cs.NE] <https://arxiv.org/abs/2503.16668>
- [29] Niki van Stein, Diederick Vermetten, and Thomas Bäck. 2024. In-the-loop Hyper-Parameter Optimization for LLM-Based Automated Design of Heuristics. arXiv:2410.16309 [cs.NE] <https://arxiv.org/abs/2410.16309>
- [30] Niki van Stein, Diederick Vermetten, Anna V. Kononova, and Thomas Bäck. 2025. Explainable Benchmarking for Iterative Optimization Heuristics. *ACM Trans. Evol. Learn. Optim.* (Feb. 2025). doi:10.1145/3716638 Just Accepted.
- [31] Diederick Vermetten, Fabio Caraffini, Anna V. Kononova, and Thomas Bäck. 2023. Modular Differential Evolution. In *Proceedings of the Genetic and Evolutionary Computation Conference (Lisbon, Portugal) (GECCO '23)*. Association for Computing Machinery, New York, NY, USA, 864–872. doi:10.1145/3583131.3590417
- [32] Diederick Vermetten, Fabio Caraffini, Bas van Stein, and Anna V. Kononova. 2022. Using structural bias to analyse the behaviour of modular CMA-ES. In *GECCO '22: Genetic and Evolutionary Computation Conference, Companion Volume, Boston, Massachusetts, USA, July 9 - 13, 2022*, Jonathan E. Fieldsend and Markus Wagner (Eds.). ACM, 1674–1682. doi:10.1145/3520304.3534035
- [33] Diederick Vermetten, Manuel López-Ibáñez, Olaf Mersmann, Richard Allmendinger, and Anna V Kononova. 2023. Analysis of modular CMA-ES on strict box-constrained problems in the SBOX-COST benchmarking suite. In *Proceedings of the Companion Conference on Genetic and Evolutionary Computation*. 2346–2353.
- [34] Diederick Vermetten, Furong Ye, Thomas Bäck, and Carola Doerr. 2024. MA-BBOB: A problem generator for black-box optimization using affine combinations and shifts. *ACM Transactions on Evolutionary Learning* (2024).
- [35] Diederick Vermetten, Furong Ye, Thomas Bäck, and Carola Doerr. 2025. MA-BBOB: A Problem Generator for Black-Box Optimization Using Affine Combinations and Shifts. *ACM Trans. Evol. Learn. Optim.* 5, 1, Article 5 (March 2025), 19 pages. doi:10.1145/3673908
- [36] Hao Wang, Diederick Vermetten, Furong Ye, Carola Doerr, and Thomas Bäck. 2022. IOHanalyzer: Detailed Performance Analyses for Iterative Optimization Heuristics. *ACM Transactions on Evolutionary Learning and Optimization* 2, 1 (apr 2022), 29 pages. <https://doi.org/10.1145/3510426>
- [37] David H Wolpert, William G Macready, et al. 1995. *No free lunch theorems for search*. Technical Report. Citeseer.

- [38] Haoran Yin, Anna V Kononova, Thomas Bäck, and Niki van Stein. 2024. Controlling the Mutation in Large Language Models for the Efficient Evolution of Algorithms. *arXiv preprint arXiv:2412.03250* (2024).
- [39] Haoran Yin, Anna V. Kononova, Thomas Bäck, and Niki van Stein. 2025. Optimizing Photonic Structures with Large Language Model Driven Algorithm Discovery. arXiv:2503.19742 [cs.NE] <https://arxiv.org/abs/2503.19742>