

# LARGE LANGUAGE MODEL-DRIVEN LARGE NEIGHBORHOOD SEARCH FOR LARGE-SCALE MILP PROBLEMS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Large Neighborhood Search (LNS) is a widely used method for solving large-scale Mixed Integer Linear Programming (MILP) problems. The effectiveness of LNS crucially depends on the choice of the search neighborhood. However, existing strategies either rely on expert knowledge or computationally expensive Machine Learning (ML) approaches, both of which struggle to scale effectively for large problems. To address this, we propose LLM-LNS, a novel Large Language Model (LLM)-driven LNS framework for large-scale MILP problems. Our approach introduces a dual-layer self-evolutionary LLM agent to automate neighborhood selection, discovering effective strategies with scant small-scale training data that generalize well to large-scale MILPs. The inner layer evolves heuristic strategies to ensure convergence, while the outer layer evolves evolutionary prompt strategies to maintain diversity. Experimental results demonstrate that the proposed dual-layer agent outperforms state-of-the-art agents such as FunSearch and EOH. Furthermore, the full LLM-LNS framework surpasses manually designed LNS algorithms like ACP, ML-based LNS methods like CL-LNS, and large-scale solvers such as Gurobi and SCIP. It also achieves superior performance compared to advanced ML-based MILP optimization frameworks like GNN&GBDT and Light-MILPopt, further validating the effectiveness of our approach.

## 1 INTRODUCTION

Mixed Integer Linear Programming (MILP) is a versatile and widely used mathematical framework for solving complex optimization problems across various domains, including transportation management (Klanšek, 2015), bin packing (Fleszar, 2022), and production planning (Adrio et al., 2023). MILPs are challenging to solve efficiently due to their NP-hard nature (Kim et al., 2021) and the exponential growth of the search space as problem size increases (Vázquez et al., 2018). To address these challenges, researchers have developed two primary approaches (Zhang et al., 2023): exact algorithms, such as branch-and-bound, and heuristic-based approximation methods.

While exact algorithms like branch-and-bound (Boyd & Mattingley, 2007; Morrison et al., 2016) are effective for small to medium-sized problems, they struggle with the computational demands of larger instances. This has led to the rise of heuristic methods, particularly Large Neighborhood Search (LNS) (Ahuja et al., 2002; Mara et al., 2022), which iteratively improves solutions by destroying and repairing parts of the current solution, allowing for exploration of large neighborhoods without full re-optimization (Song et al., 2020; Ye et al., 2023a). However, LNS performance depends heavily on neighborhood selection, which is often hand-crafted and requires significant domain expertise. Designing these operators can be labor-intensive and prone to *cold-start issues*, where limited prior knowledge is available to guide the search (Zhang et al., 2023).

In recent years, machine learning (ML) techniques, including reinforcement learning (Wu et al., 2021; Song et al., 2020) and imitation learning (Sonnerat et al., 2021; Nair et al., 2020), have been applied to automate the design of neighborhood selection strategies. These methods aim to learn heuristic strategies from training datasets, reducing reliance on expert knowledge and allowing the algorithms to adapt to new, homogeneous instances. However, ML-based LNS approaches come with their own challenges. For reinforcement learning, *slow convergence* is a common issue (Beggs,

2005), particularly in large-scale MILP problems, due to the vast search space and the need for extensive exploration before identifying effective strategies. On the other hand, imitation learning requires large amounts of high-quality, labeled data, which can be *computationally expensive* to generate using expert algorithms (Huang et al., 2023b). As a result, both hand-crafted and ML-based methods struggle to efficiently solve large-scale MILP problems.

The rise of Large Language Models (LLMs) offers a promising solution to these challenges. Unlike traditional hand-crafted methods, LLMs come pretrained with vast general knowledge, allowing them to reason about complex tasks and learn problem structures with minimal training data, thus avoiding *cold-start issues*. Additionally, LLMs can adapt to new problems through interactive reasoning, reducing the need for extensive exploration and addressing the *slow convergence* of reinforcement learning. Furthermore, LLMs can dynamically generate heuristic strategies without relying on large labeled datasets, which significantly reduces the *computational overhead* typically associated with imitation learning (Yang et al., 2024; Lange et al., 2024). While LLMs have shown potential in generating strategies for combinatorial optimization problems (Ye et al., 2024; Elhenawy et al., 2024), they often lack the problem-specific refinement needed to produce efficient heuristics without additional guidance (Plaat et al., 2024). Approaches like FunSearch (Romera-Paredes et al., 2024) and Evolution of Heuristic (EOH) (Liu et al., 2024) combine LLMs with evolutionary algorithms (Simon, 2013), but rely on fixed strategies, limiting solution diversity and leading to poor convergence due to insufficient directionality. This underscores the need for a more adaptive framework to fully harness LLMs for large-scale MILP problems.

In this paper, we propose LLM-LNS, a novel Large Language Model-driven Large Neighborhood Search framework designed specifically for solving large-scale MILP problems, which can discover effective neighborhood selection strategies for LNS with scant small-scale training data that generalize well to large-scale MILPs. Our key innovations are as follows:

- **Dual-layer Self-evolutionary LLM Agent:** We propose a novel LLM agent with a dual-layer self-evolutionary mechanism for automatically generating heuristic strategies. The inner layer evolves both thoughts and code representations of heuristic strategies, ensuring convergence, while the outer layer evolves evolutionary prompt strategies to maintain diversity, preventing the search process from getting trapped in local optima.
- **Differential Memory for Directional Evolution:** We introduce differential evolution in the agent to guide both crossover and variation. By feeding the fitness values of parent strategies back into the LLM, we leverage its memory to learn how to evolve from less effective to more effective strategies. This feedback mechanism enables the LLM to act as an optimizer, identifying promising directions and leading to more efficient improvements.
- **Application to Neighborhood Selection in LNS:** We apply the proposed dual-layer LLM agent to the neighborhood selection strategy generation in LNS. By utilizing only a small amount of training data from small-scale problems, the LLM agent can discover new neighborhood selection strategies that generalize well to large-scale MILP problems.
- **Comprehensive Experimental Validation:** We validate the effectiveness of our proposed LLM-LNS at two levels. First, we test its agent’s performance on heuristic generation tasks of combinatorial optimization problems, demonstrating its superiority over state-of-the-art methods such as FunSearch (Romera-Paredes et al., 2024) and EOH (Liu et al., 2024). Second, we evaluate its performance on large-scale MILP problems, where it outperforms traditional LNS methods (e.g., ACP (Ye et al., 2023a)), ML-based LNS methods (e.g., CL-LNS (Huang et al., 2023b)), and leading solvers like Gurobi (Gurobi Optimization, LLC, 2023) and SCIP (Maher et al., 2016). Furthermore, our proposed LLM-LNS surpasses modern ML-based optimization frameworks for large-scale MILP, such as GNN&GBDT (Ye et al., 2023c) and Light-MILPopt (Ye et al., 2023b). These results confirm the effectiveness of our proposed LLM-LNS in solving large-scale optimization problems.

## 2 RELATED WORK

### 2.1 MIXED INTEGER LINEAR PROGRAMMING

Mixed Integer Linear Programming (MILP) problems represent a class of combinatorial optimization problems characterized by a linear objective function subject to a set of linear constraints, where

some or all decision variables are restricted to integer values. An MILP can be defined as follows:

$$\min_x c^T x, \quad \text{subject to} \quad Ax \leq b, l \leq x \leq u, x_i \in \mathbb{Z}, i \in \mathcal{I}, \quad (1)$$

where  $x$  represents the decision variables, with  $n \in \mathbb{Z}$  denoting the dimensionality of the integer variables and  $l, u, c \in \mathbb{R}^n$  corresponding to the lower bounds, upper bounds, and coefficients of the decision variables, respectively. The matrix  $A \in \mathbb{R}^{m \times n}$  and the vector  $b \in \mathbb{R}^m$  define the linear constraints of the problem. The set  $\mathcal{I} \subseteq \{1, 2, \dots, n\}$  denotes the indices of variables that are constrained to integer values. A feasible solution to the MILP problem satisfies all constraints, and the optimal solution minimizes the objective function value. (Artigues et al., 2015; Pizaruk, 2019)

## 2.2 LARGE NEIGHBORHOOD SEARCH

Large Neighborhood Search (LNS) is a widely used heuristic for solving MILP problems. It iteratively improves solutions by exploring predefined neighborhoods around a current solution. However, the effectiveness of LNS heavily relies on the neighborhood selection strategy, as poor choices can lead to stagnation in local optima.

Several approaches have been proposed to address this challenge. One common method is random-LNS (Song et al., 2020), which randomly partitions integer variables into disjoint subsets and optimizes one subset in each iteration while fixing the others. However, random-LNS uses a fixed neighborhood size and overlooks correlations between decision variables, limiting its performance. To overcome these drawbacks, the Adaptive Constraint Partitioning (ACP) framework (Ye et al., 2023a) introduces a dynamic strategy that adjusts the neighborhood size, optimizing all decision variables associated with randomly selected constraints in each iteration. This ensures that highly correlated variables are optimized together, improving performance. Similar strategies have been explored in other works (Huang et al., 2023a; Han et al., 2023), but they still rely on manually designed heuristics, requiring expert knowledge and lacking adaptability to new problem instances.

To address this limitation, machine learning methods have been applied to automate neighborhood selection. Reinforcement learning (RL) approaches define reward functions based on solution improvements, allowing models to learn promising neighborhoods through interaction with the problem (Wu et al., 2021; Song et al., 2020; Nair et al., 2020). Imitation learning, on the other hand, uses large amount of large-scale sampling (Huang et al., 2023b; Zhou et al., 2023) or expert algorithms (Sonnerat et al., 2021) to guide the selection process. While these techniques reduce reliance on handcrafted strategies, RL struggles with convergence in large-scale MILP problems, and imitation learning requires extensive sampling, making it computationally expensive. This highlights the need for more efficient, automatically designed neighborhood selection strategies.

## 2.3 LARGE LANGUAGE MODEL FOR HEURISTIC STRATEGY DESIGN

The rise of Large Language Models (LLMs) has opened new possibilities for generating heuristic strategies to solve combinatorial optimization problems (Yang et al., 2024; Lange et al., 2024). LLMs excel at generating high-level ideas and reasoning over complex tasks, but they often lack problem-specific knowledge, limiting their ability to create effective heuristics without additional guidance (Plaat et al., 2024). To overcome these limitations, recent works have integrated LLMs with evolutionary algorithms (EA) to iteratively refine heuristics.

FunSearch (Romera-Paredes et al., 2024) is a notable attempt that combines LLMs with evolutionary frameworks. FunSearch uses LLMs to generate functions, which are then evolved through an evolutionary search process. This approach has demonstrated success in outperforming hand-crafted algorithms on specific optimization problems. However, FunSearch is computationally expensive, often requiring millions of LLM queries to identify effective heuristic functions, which limits its practicality in many real-world applications. A more recent approach, Evolution of Heuristic (EOH) (Liu et al., 2024), builds on the strengths of LLMs and evolutionary computation while addressing some of FunSearch’s limitations. EOH introduces a novel evolutionary paradigm where heuristics,

Table 1: Comparison of Features Between FunSearch, EOH, and LLM-LNS.

	FunSearch	EOH	LLM-LNS
Heuristic Evolution	✓	✓	✓
Thought Evolution	×	✓	✓
Evolutionary Prompt	×	×	✓
Strategy Evolution	×	×	✓
Directional Evolution	×	×	✓

162  
163  
164  
165  
166  
167  
168  
169  
170  
171  
172  
173  
174  
175  
176  
177  
178  
179  
180  
181  
182  
183  
184  
185  
186  
187  
188  
189  
190  
191  
192  
193  
194  
195  
196  
197  
198  
199  
200  
201  
202  
203  
204  
205  
206  
207  
208  
209  
210  
211  
212  
213  
214  
215

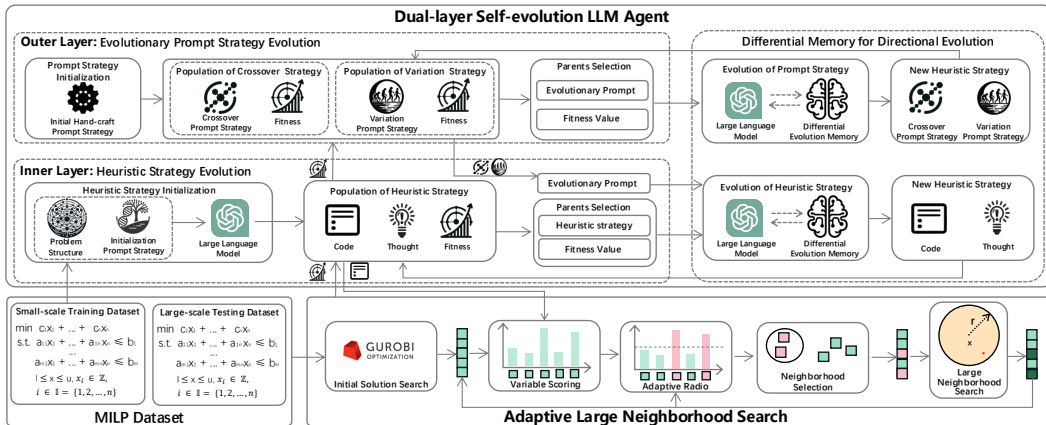


Figure 1: An overview of the proposed LLM-LNS framework. The framework consists of a dual-layer self-evolutionary LLM agent for solving large-scale MILP problems. In the outer layer, evolutionary prompt strategies are generated and passed to the inner layer, where heuristic strategies are evolved. A differential memory mechanism uses fitness feedback to refine these strategies across iterations. The refined strategies are fed into the Adaptive Large Neighborhood Search process, which iteratively improves solutions with the support of solvers like Gurobi.

represented as natural language "thoughts," are translated into executable code by LLMs. These thoughts and their corresponding code are evolved within an EA framework, enabling the efficient generation of high-performance heuristics.

As shown in Table 1, while FunSearch and EOH have advanced the integration of LLMs with evolutionary algorithms, they still have limitations. Both use fixed, manually designed evolutionary strategies that limit solution diversity and often lead to premature convergence. Additionally, they lack directional mechanisms to guide the search, reducing adaptability and improvement over iterations. These challenges highlight the need for more adaptive frameworks to fully harness LLMs in large-scale optimization tasks.

### 3 METHOD

In this section, we introduce LLM-LNS, a Large Language Model-driven Large Neighborhood Search framework designed to solve large-scale MILP problems. As shown in Figure 1, the framework is composed of two main components: a **Dual-layer Self-evolutionary LLM Agent** and a **Adaptive Large Neighborhood Search** process.

#### 3.1 DUAL-LAYER SELF-EVOLUTIONARY LLM AGENT

The Dual-layer Self-evolutionary LLM Agent is the core component of our framework, responsible for generating and evolving heuristic and prompt strategies. The **Dual-layer Self-evolutionary Structure** consists of an Inner Layer that evolves heuristic strategies to accelerate convergence, and an Outer Layer that evolves evolutionary prompt strategies to enhance diversity in heuristic generation. Another key innovation is the incorporation of **Differential Memory for Directional Evolution**, which accelerates convergence by learning the direction of improvement from less effective strategy to better ones. Together, these innovations ensure a balance between exploration and exploitation, significantly improving the efficiency and preventing stagnation in local optima.

##### 3.1.1 DUAL-LAYER SELF-EVOLUTIONARY STRUCTURE

The Dual-layer Self-evolutionary Structure is the core component of the LLM-LNS framework. It is designed to evolve both evolutionary prompt strategies and heuristic strategies in a synergistic manner, leveraging LLMs for automated heuristic design and refinement. This dual-layered structure

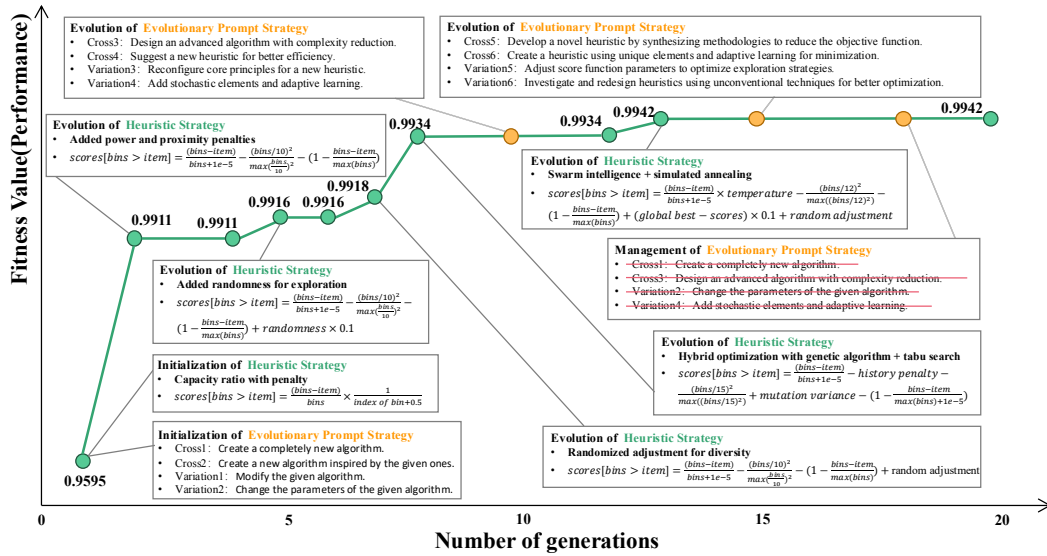


Figure 2: Evolution of Dual-layer Self-evolutionary LLM Agent for online bin packing. We outline the key thoughts and the corresponding code snippets of the best heuristics produced in some generations during the evolution of **heuristic strategies**. Additionally, we highlight the evolution of **evolutionary prompt strategies**, which dynamically adapt the prompt strategies to guide the LLM in generating more effective and diverse heuristics. Both strategies contribute to the overall improvement in performance and convergence throughout the evolutionary process.

mimics the heuristic development process of human experts, ensuring a balance between exploration and exploitation throughout the search process.

**Inner Layer: Heuristic Strategy Evolution.** The Inner Layer focuses on evolving heuristic strategies, which consist of both natural thought and corresponding code implementations, with an emphasis on *convergence*. Key aspects of Inner Layer, as illustrated in Figure 1 and Figure 2, include:

- **Initialization of Heuristic Strategies:** The initial set of heuristics is generated by feeding the structural information from small-scale training problems, along with an initialization prompt strategy, into the LLM. This produces the first generation of heuristic strategies. For example, at generation 1, a basic heuristic is initialized with a fitness value of 0.9595, based on a capacity ratio with penalty calculation, and is expressed both in natural language thought and executable code.
- **Evolution of Heuristic Strategies:** In each generation, new heuristic strategies are evolved by selecting parent strategies from the current heuristic population. Strategies with higher fitness values are more likely to be selected as parents. These parents are then combined with evolutionary strategies, selected from the Outer Layer’s population of prompt strategies (e.g., crossover or variation prompts), to guide the LLM in generating new offspring strategies. For instance, at generation 5, randomness is introduced for exploration, achieving a fitness value of 0.9916. By generation 8, the evolution process incorporates hybrid optimization techniques, such as genetic algorithms combined with tabu search, resulting in a fitness value of 0.9934. This iterative process enables the LLM to continually refine strategies and explore new solution spaces.
- **Evaluation and Final Selection:** After new heuristic strategies are generated, they are evaluated by integrating them into the Adaptive Large Neighborhood Search process, where each heuristic is applied to solve small-scale instances from the training dataset. The performance of each strategy is measured by its objective function value, which serves as its fitness score. After multiple iterations of evolution and evaluation, the best-performing heuristic strategies are identified based on their fitness. By generation 20, advanced techniques like swarm intelligence and simulated annealing are incorporated, and the final best

strategy—achieving a fitness value of 0.9942—is selected for output. This iterative evaluation ensures that the framework converges to the most effective heuristic.

**Outer Layer: Evolutionary Prompt Strategy Evolution.** The Outer Layer focuses on evolving evolutionary prompt strategies, which guide the LLM in generating new heuristic strategies. The emphasis in this layer is on *exploration* to maintain diversity and prevent premature convergence in the heuristic strategy population. The key stages of Outer Layer, as illustrated in Figure 1 and Figure 2, include:

- *Initialization of Prompt Strategies:* The initial set of evolutionary prompt strategies is hand-crafted and designed to perform basic crossover and variation operations, instructing the LLM on how to combine or modify existing heuristic strategies in the inner layer. For example, at generation 1, basic prompt strategies like Cross1 and Cross2 are set, which help the LLM generate new heuristic strategies by recombining or tweaking existing ones.
- *Evolution of Prompt Strategies:* As the evolution progresses, more complex prompt strategies are introduced to address stagnation in the heuristic population. Specifically, if the top- $l$  individuals in the heuristic population remain unchanged for  $t$  consecutive generations, we infer that the evolution may have converged to a local optimum. This triggers the evolution of new prompt strategies. As shown in Figure 2, signs of stagnation were observed in both the 10th and 15th generations. In response, new prompt strategies were generated to overcome the local optimality issue. At generation 10, prompts such as Cross3 and Cross4 were designed to enhance efficiency and reduce algorithmic complexity. By generation 15, even more advanced strategies like Variation5 and Variation6 were introduced, incorporating stochastic elements and adaptive learning to increase diversity and explore new heuristic possibilities. This systematic evolution of prompt strategies helps ensure that the heuristic population continues to evolve and does not get trapped in local optima.
- *Evaluation and Management of Prompt Strategies:* To ensure the efficiency and effectiveness of the prompt strategy population, each prompt strategy is evaluated based on the performance of the heuristic strategies it generates. Specifically, for each prompt strategy, the top- $k$  performing heuristic strategies it produces are tracked, and the average fitness score of these heuristics is used as the fitness score for the prompt strategy itself. This fitness-based evaluation allows us to manage the prompt population and control its size. As the number of prompt strategies increases over generations, underperforming strategies are pruned to prevent excessive growth and focus on the most effective strategies. For example, by generation 18, several underperforming prompt strategies (e.g., the four worst-performing strategies) are removed, as shown in Figure 2. This pruning process ensures that only the most effective prompt strategies continue to evolve, maintaining both diversity and efficiency in the evolutionary process. For parameter details, see Appendix A.

The synergy between the **Inner Layer** and **Outer Layer** drives rapid evolution of effective heuristics and novel evolutionary prompt strategies, as shown in Figure 2. Early generations focus on basic principles, but with the introduction of advanced prompt strategies, such as complexity reduction and adaptive learning, the system quickly adapts to overcome local optima. Notably, the sharp performance improvements between generations 5 to 15 demonstrate the framework’s ability to autonomously discover and refine creative strategies, leading to continuous enhancements in heuristic performance. This dual-layered approach ensures efficient exploration and exploitation, enabling the LLM-LNS framework to tackle large-scale problems with minimal human intervention.

### 3.1.2 DIFFERENTIAL MEMORY FOR DIRECTIONAL EVOLUTION

In our Dual-layer Self-evolutionary LLM Agent, both heuristic strategies and evolutionary prompt strategies evolve through a process that incorporates differential memory. Traditional evolutionary methods often treat the generation of new strategies as independent of the performance history of previous strategies, but our approach introduces a more sophisticated mechanism—Differential Memory for Directional Evolution—which leverages the LLM not only as a generator but also as an optimizer that learns from the evolutionary history of the strategies.

In this process, the LLM is provided with the fitness values of parent strategies and the relationship between higher- and lower-performing strategies. These fitness scores are fed back into the LLM,

**Algorithm 1** Adaptive Large Neighborhood Search (ALNS)

---

**Require:** Initial solution  $\mathbf{x}_0$ , initial neighborhood size  $k$ , time limit  $T$ , threshold  $\epsilon$ , iteration limit  $p$ , minimum and maximum neighborhood sizes  $k_{\min}, k_{\max}$ , decision variable count  $n$ , adjustment rate  $u\%$  (percentage)

```

1: Initialize solution  $\mathbf{x} \leftarrow \mathbf{x}_0$ , set time  $t \leftarrow 0$ 
2: while  $t < T$  do
3:   Compute variable scores using LLM agent
4:   Select top- $k$  variables to form neighborhood
5:   Solve subproblem using solver within neighborhood
6:   Update solution  $\mathbf{x}$  if improved
7:   if time spent in neighborhood exceeds limit then
8:      $k \leftarrow \max(k_{\min}, k - \lceil u\% \cdot n \rceil)$  ▷ Reduce search radius by  $u\%$  of  $n$ 
9:   else if improvement in objective  $< \epsilon$  for  $p$  consecutive iterations then
10:     $k \leftarrow \min(k_{\max}, k + \lceil u\% \cdot n \rceil)$  ▷ Expand search radius by  $u\%$  of  $n$ 
11:   end if
12:   Update time  $t$ 
13: end while
14: return  $\mathbf{x}$ 

```

---

allowing it to learn from past generations and evolve strategies that are more likely to improve over time. By understanding the directional improvement from less effective to more effective strategies, the LLM can guide both crossover and variation operations toward generating more promising solutions. This feedback loop ensures that the LLM continuously adapts as it learns from the evolving population, enabling it to act as an active optimizer.

Rather than simply combining or mutating existing strategies, the LLM uses the differential memory to identify trends and patterns in the evolution of strategies. Given a meta-prompt that describes the optimization task, along with fitness values from previously evaluated strategies, the LLM generates new candidate strategies that are more likely to follow the observed direction of improvement. This directional evolution mechanism allows the LLM to explore the search space more effectively, focusing on areas that have shown promise in earlier generations. Once new candidate strategies are generated, they are evaluated and assigned fitness scores, which are then fed back into the LLM. This iterative process ensures that the LLM becomes increasingly proficient at generating and refining strategies over successive generations. By utilizing differential memory, the LLM can efficiently explore and exploit the solution space, accelerating the convergence towards optimal solutions while avoiding stagnation. The LLM thus serves as both a generator and an optimizer, continuously learning how to evolve strategies that lead to more efficient and targeted improvements.

### 3.2 ADAPTIVE LARGE NEIGHBORHOOD SEARCH

As discussed in Sec.2.2, LNS is a popular heuristic for solving large-scale MILP problems, iteratively improving solutions by exploring neighborhoods around a current solution. However, its effectiveness depends on selecting appropriate neighborhoods; poor choices can lead to stagnation in local optima. To address this, we propose an Adaptive Large Neighborhood Search (ALNS) that leverages the Dual-layer Self-evolutionary LLM Agent for variable scoring and adaptive neighborhood selection, dynamically adjusting the search radius.

As shown in Figure 1, the LLM agent optimizes neighborhood selection by ranking decision variables based on their impact on the objective function. The agent is provided with comprehensive problem information, including the problem formulation, initial solution, and current state. Using this information, the LLM scores each decision variable based on its potential to improve the objective value. In each iteration, a subset of high-scoring variables is selected to define the neighborhood for further exploration. The neighborhood size is controlled by an adaptive radius, which adjusts dynamically based on search progress. If consecutive iterations show little improvement (below a given threshold  $\epsilon$ ), the radius is increased to explore a broader search space. Conversely, if the search within a neighborhood takes excessive time, the radius is reduced to focus on a smaller, more manageable subset of variables. For parameter details, see Appendix A.

A key innovation of our approach is the LLM agent’s ability to generalize neighborhood selection strategies. Initially trained on small-scale MILP problems, the agent learns effective strategies for selecting promising decision variables. Through this training, the LLM refines its understanding of

Table 2: **Online Bin Packing Heuristic Comparison.** This table compares the performance of various bin packing heuristics based on the fraction of excess bins (lower values indicate better performance) across different Weibull distribution instances.

	1k_C100	5k_C100	10k_C100	1k_C500	5k_C500	10k_C500	Avg
First Fit	5.32%	4.40%	4.44%	4.97%	4.27%	4.28%	4.61%
Best Fit	4.87%	4.08%	4.09%	4.50%	3.91%	3.95%	4.23%
FunSearch	3.78%	<b>0.80%</b>	<b>0.33%</b>	6.75%	1.47%	0.74%	2.31%
EOH	4.48%	0.88%	0.83%	4.32%	1.06%	0.97%	2.09%
Ours	<b>3.58%</b>	0.85%	0.41%	<b>3.67%</b>	<b>0.82%</b>	<b>0.42%</b>	<b>1.63%</b>

how different variable selections influence solution quality and runtime. Once trained, the LLM can generalize these strategies to larger, more complex problems. This transfer of knowledge allows the LLM to navigate the vast search space of large-scale MILPs, making adaptive decisions that balance exploration and exploitation. By leveraging the LLM’s ability to generalize from small-scale problems, the method ensures that neighborhood selection is both adaptive and intelligent, focusing computational resources on the most promising regions of the solution space.

This adaptive mechanism significantly improves the efficiency of LNS, enabling faster convergence to high-quality solutions for large-scale MILP problems. The pseudocode in Algorithm 1 outlines the ALNS process, where the LLM agent continuously uses its understanding of the search space, balancing exploration and exploitation through adaptive control of neighborhood size.

## 4 EXPERIMENT

To validate the effectiveness of the proposed LLM-LNS framework, we conduct two sets of experiments. First, we evaluate our proposed Dual-layer Self-evolutionary LLM Agent on heuristic generation tasks for combinatorial optimization problems, comparing it against methods like FunSearch (Romera-Paredes et al., 2024) and EOH (Liu et al., 2024). Second, we assess the full LLM-LNS framework on large-scale MILP problems, where it is compared against traditional LNS methods (e.g., ACP (Ye et al., 2023a)), ML-based LNS approaches (e.g., CL-LNS (Huang et al., 2023b)), the SOTA solvers like Gurobi (Gurobi Optimization, LLC, 2023) and SCIP (Maher et al., 2016), and modern ML optimization frameworks such as GNN&GBDT (Ye et al., 2023c) and Light-MILPopt (Ye et al., 2023b). More experimental results and details are provided in the Appendix A B C D.

### 4.1 HEURISTIC GENERATION FOR COMBINATORIAL OPTIMIZATION PROBLEMS

In this section, we evaluate the performance of the Dual-layer Self-evolutionary LLM Agent in generating heuristic strategies for well-known combinatorial optimization problems. We focus on two widely studied problems: Online Bin Packing (Seiden, 2002) and the Traveling Salesman Problem (TSP) (Hoffman et al., 2013). Our method is compared against several hand-crafted heuristics, state-of-the-art machine learning methods, and other automatically designed heuristics.

#### 4.1.1 ONLINE BIN PACKING

The objective of the Online Bin Packing problem is to allocate a collection of items into the fewest possible bins of fixed capacity. We follow the experimental setup from Romera-Paredes et al. (2024), using Weibull distribution instances with varying numbers of items (1k to 10k) and bin capacities (100 and 500). The performance of each method is measured by the fraction of excess bins used, where lower values indicate better performance. We compare our method against several baselines, including hand-crafted heuristics First Fit (Tang et al., 2016) and Best Fit (Shor, 1991), which are widely used in practice, as well as automatically generated heuristics FunSearch (Romera-Paredes et al., 2024) and EOH (Liu et al., 2024), which represent state-of-the-art approaches.

As shown in Table 2, our method consistently achieves the best performance across different problem sizes and capacities, with an average excess bin fraction of 1.63%, outperforming both hand-crafted heuristics and automatically generated methods. In particular, our approach excels on the 10k items, capacity 500 instance, achieving a fraction of excess bins of 0.42%, outperforming FunSearch (0.74%) and EOH (0.97%). This result highlights the strong scalability and generalization ability of our method, making it particularly effective in handling large-scale, high-capacity scenarios.



Table 3: **Traveling Salesman Problems Heuristic Performance Evaluation.** This table provides a comparison of the relative distance to the best-known solutions for different routing heuristics (lower values indicate better performance) on a subset of TSPLib benchmark instances.

	rd100	pr124	bier127	kroA150	u159	kroB200	Avg
NI	19.91%	15.50%	23.21%	18.17%	23.59%	24.10%	20.75%
FI	9.38%	4.43%	8.04%	8.54%	11.15%	7.54%	8.18%
Or-Tools	<b>0.01%</b>	0.55%	0.66%	0.02%	1.75%	2.57%	0.93%
AM	3.41%	3.68%	5.91%	3.78%	7.55%	7.11%	5.24%
POMO	<b>0.01%</b>	0.60%	13.72%	0.70%	0.95%	1.58%	2.93%
LEHD	<b>0.01%</b>	1.11%	4.76%	1.40%	1.13%	0.64%	1.51%
EOH	<b>0.01%</b>	<b>0.00%</b>	0.42%	0.29%	<b>-0.01%</b>	<b>0.26%</b>	0.16%
Ours	<b>0.01%</b>	<b>0.00%</b>	<b>0.01%</b>	<b>0.00%</b>	<b>-0.01%</b>	0.44%	<b>0.08%</b>

#### 4.1.2 TRAVELING SALESMAN PROBLEM

The Traveling Salesman Problem (TSP) is a classic combinatorial optimization problem where the goal is to find the shortest route that visits all given locations exactly once. We evaluate our method on a subset of TSPLib benchmark instances (Reinelt, 1991), with performance measured by the relative distance to the best-known solutions (lower values indicate better performance). We compare our method against two types of baselines: hand-crafted heuristics and AI-generated heuristics. The hand-crafted heuristics include Nearest Insertion (NI) and Farthest Insertion (FI) (Rosenkrantz et al., 1977), two widely used constructive heuristics. We also include Google OR-Tools (Perron & Furnon), a popular solver, using its default settings and the recommended local search option. Beyond EOH (Liu et al., 2024), we compare against the Attention Model (AM) (Kool et al., 2018), POMO (Kwon et al., 2020), and LEHD (Luo et al., 2023), all of which are ML-based methods.

As shown in Table 3, our method achieves the best average performance with a 0.08% gap to the best-known solutions, outperforming both hand-crafted heuristics and neural network-based methods. Notably, on the bier127 instance, our method achieves a relative distance of just 0.01% to the best-known solution, significantly outperforming EOH (0.42%) and other baselines, including LEHD (4.76%) and AM (5.91%). This substantial improvement highlights the effectiveness of our approach in solving challenging instances of the TSP.

It is important to note that both the Online Bin Packing and TSP problems use the same GPT-4o-mini LLM, with identical settings: 20 iterations and a population size of 20 for Online Bin Packing, and 10 for the TSP problem. Despite these identical settings, our method consistently outperforms EOH in both problems, showcasing the superior efficiency of the dual-layer self-evolutionary mechanism in exploring the solution space. This mechanism allows our method to dynamically adapt and refine solutions, resulting in better overall performance with the same computational resources. These results underscore the robustness and scalability of our approach, offering a promising direction for solving large-scale combinatorial optimization problems using LLMs.

#### 4.2 PERFORMANCE OF LLM-LNS ON LARGE-SCALE MILP PROBLEMS

To validate the effectiveness of the proposed LLM-LNS framework for large-scale MILP problems, we evaluate its performance on four widely-used benchmark datasets: Set Covering (SC) (Caprara et al., 2000), Minimum Vertex Cover (MVC) (Dinur & Safra, 2005), Maximum Independent Set (MIS) (Tarjan & Trojanowski, 1977), and Mixed Integer Knapsack Set (MIKS) (Atamtürk, 2003). Initially, LLM-LNS is trained on small-scale problems with tens of thousands of variables and constraints and then tested on large-scale instances (Table 4) to assess its scalability and generalization.

We compare LLM-LNS with several state-of-the-art baselines, including heuristic LNS methods like Random-LNS (Song et al., 2020), Adaptive Constraint Propagation (ACP) (Ye et al., 2023a), and the learning-based CL-LNS framework

Table 4: The size of one real-world case study in the internet domain and four widely used NP-hard benchmark MILPs.

Problem	Scale	Number of Variables	Number of Constraints
SC (Minimize)	SC <sub>1</sub>	200000	200000
	SC <sub>2</sub>	2000000	2000000
MVC (Minimize)	MVC <sub>1</sub>	100000	300000
	MVC <sub>2</sub>	1000000	3000000
MIS (Maximize)	MIS <sub>1</sub>	100000	300000
	MIS <sub>2</sub>	1000000	3000000
MIKS (Maximize)	MIKS <sub>1</sub>	200000	200000
	MIKS <sub>2</sub>	2000000	2000000

Table 5: **Comparison of objective values on large-scale MILP instances across different methods.** For each instance, the best-performing objective value is highlighted in bold. The - symbol indicates that the method was unable to generate samples for any instance within 30,000 seconds, while \* indicates that the GNN&GBDT framework could not solve the MILP problem.

	SC <sub>1</sub>	SC <sub>2</sub>	MVC <sub>1</sub>	MVC <sub>2</sub>	MIS <sub>1</sub>	MIS <sub>2</sub>	MIKS <sub>1</sub>	MIKS <sub>2</sub>
Random-LNS	16140.6	169417.5	27031.4	276467.5	22892.9	223748.6	36011.0	351964.2
ACP	17672.1	182359.4	26877.2	274013.3	23058.0	226498.2	34190.8	332235.6
CL-LNS	-	-	31285.0	-	15000.0	-	-	-
Gurobi	17934.5	320240.4	28151.3	283555.8	21789.0	216591.3	32960.0	329642.4
SCIP	25191.2	385708.4	31275.4	491042.9	18649.9	9104.3	29974.7	168289.9
GNN&GBDT	16728.8	252797.2	27107.9	271777.2	22795.7	227006.4	*	*
Light-MILPOPT	16108.1	160015.5	26950.7	269571.5	22966.5	230432.9	36125.5	362265.1
LLM-LNS(Ours)	<b>15802.7</b>	<b>158878.9</b>	<b>26725.3</b>	<b>268033.7</b>	<b>23169.3</b>	<b>231636.9</b>	<b>36479.8</b>	<b>363749.5</b>

(Huang et al., 2023b). Additionally, we include traditional solvers like Gurobi (Gurobi Optimization, LLC, 2023) and SCIP (Maher et al., 2016), as well as modern ML-based frameworks such as GNN&GBDT (Ye et al., 2023c) and Light-MILPOpt (Ye et al., 2023b). To ensure a fair comparison, Gurobi is used as the sub-solver in the neighborhood search step across all methods. For LLM-LNS, the neighborhood selection strategy is trained over 20 iterations on smaller problems before being applied to larger instances. Detailed results and discussions are provided in the Appendix D.

The experimental results, summarized in Table 5, show that LLM-LNS consistently outperforms traditional LNS-based heuristics and learning-based methods. Unlike hand-crafted LNS strategies, which are typically static and less effective as problem complexity increases, LLM-LNS dynamically adapts through its dual-layer self-evolutionary mechanism, enabling more efficient exploration of the solution space. Even compared to state-of-the-art learning-based LNS methods like CL-LNS, LLM-LNS demonstrates superior performance. Although CL-LNS represents one of the most advanced learning-based approaches, it often fails to complete sampling within an acceptable time for large-scale instances, and even when results are obtained, the solution quality is significantly lower. This highlights the challenges faced by existing LNS-based methods when dealing with large and complex MILP problems, while underscoring the robustness and adaptability of LLM-LNS.

In addition, LLM-LNS shows a clear advantage over traditional solvers like Gurobi and SCIP, as well as learning-based methods such as GNN&GBDT and Light-MILPOpt. While traditional solvers perform competitively on smaller instances, their performance degrades significantly as the problem size increases. Similarly, learning-based methods struggle with large-scale MILPs, finding it difficult to efficiently explore the exponentially growing solution space. In contrast, LLM-LNS consistently delivers superior results across both small and large-scale problems, offering a scalable and efficient solution. These findings suggest that LLM-LNS not only bridges the gap between traditional and learning-based methods, but also opens new avenues for scalable optimization in large-scale MILPs.

Overall, the experimental results demonstrate the effectiveness of our proposed innovations. In the first set of experiments, we validate the capability of the **Dual-layer Self-evolutionary LLM Agent** to autonomously generate competitive heuristic strategies for combinatorial optimization problems, consistently outperforming state-of-the-art methods such as FunSearch and EOH. This confirms the agent’s ability to balance exploration and exploitation, as guided by the **Differential Memory for Directional Evolution**. In the second set, we apply the **LLM-LNS framework** to large-scale MILP problems, where it not only surpasses traditional LNS methods and advanced solvers like Gurobi and SCIP, but also demonstrates superior scalability compared to modern ML-based frameworks. These results highlight the success of applying our LLM agent to **neighborhood selection in LNS**, showcasing its generalization to complex, large-scale problems with minimal training data.

## 5 CONCLUSION

In this paper, we propose **LLM-LNS**, a Large Language Model-driven LNS framework for solving large-scale MILP problems, utilizing a dual-layer self-evolutionary LLM agent to automate heuristic strategy generation. Experiments show that LLM-LNS consistently outperforms traditional solvers, learning-based methods, and state-of-the-art LNS frameworks. Future work will explore new agent architectures and broader optimization problems, aiming to further enhance the integration of LLMs with optimization techniques. The code of LLM-LNS will be open-sourced after the paper review.

## REFERENCES

- 540  
541  
542 Gerardo Adrio, Alberto García-Villoria, Marc Juanpera, and Rafael Pastor. Milp model for the  
543 mid-term production planning in a chemical company with non-constant consumption of raw  
544 materials. an industrial application. *Computers & Chemical Engineering*, 177:108361, 2023.
- 545  
546 Ravindra K Ahuja, Özlem Ergun, James B Orlin, and Abraham P Punnen. A survey of very large-  
547 scale neighborhood search techniques. *Discrete Applied Mathematics*, 123(1-3):75–102, 2002.
- 548  
549 Christian Artigues, Oumar Koné, Pierre Lopez, and Marcel Mongeau. Mixed-integer linear pro-  
550 gramming formulations. *Handbook on project management and scheduling vol. 1*, pp. 17–41,  
2015.
- 551  
552 Alper Atamtürk. On the facets of the mixed-integer knapsack polyhedron. *Mathematical Program-*  
553 *ming*, 98(1-3):145–175, 2003.
- 554  
555 Alan W Beggs. On the convergence of reinforcement learning. *Journal of economic theory*, 122(1):  
1–36, 2005.
- 556  
557 Stephen Boyd and Jacob Mattingley. Branch and bound methods. *Notes for EE364b, Stanford*  
558 *University*, 2006:07, 2007.
- 559  
560 Alberto Caprara, Paolo Toth, and Matteo Fischetti. Algorithms for the set covering problem. *Annals*  
561 *of Operations Research*, 98:353–371, 2000.
- 562  
563 Irit Dinur and Samuel Safra. On the hardness of approximating minimum vertex cover. *Annals of*  
564 *mathematics*, pp. 439–485, 2005.
- 565  
566 Mohammed Elhenawy, Ahmad Abutahoun, Taqwa I Alhadidi, Ahmed Jaber, Huthaifa I Ashqar,  
567 Shadi Jaradat, Ahmed Abdelhay, Sebastien Glaser, and Andry Rakotonirainy. Visual reasoning  
568 and multi-agent approach in multimodal large language models (mllms): Solving tsp and mtsp  
569 combinatorial challenges. *arXiv preprint arXiv:2407.00092*, 2024.
- 570  
571 Krzysztof Fleszar. A milp model and two heuristics for the bin packing problem with conflicts and  
572 item fragmentation. *European Journal of Operational Research*, 303(1):37–53, 2022.
- 573  
574 Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2023. URL <https://www.gurobi.com>.
- 575  
576 Qingyu Han, Linxin Yang, Qian Chen, Xiang Zhou, Dong Zhang, Akang Wang, Ruoyu Sun, and Xi-  
577 aodong Luo. A gnn-guided predict-and-search framework for mixed-integer linear programming.  
578 *arXiv preprint arXiv:2302.05636*, 2023.
- 579  
580 Karla L Hoffman, Manfred Padberg, Giovanni Rinaldi, et al. Traveling salesman problem. *Encyclo-*  
581 *pedia of operations research and management science*, 1:1573–1578, 2013.
- 582  
583 Taoan Huang, Aaron Ferber, Yuandong Tian, Bistra Dilikina, and Benoit Steiner. Local branching  
584 relaxation heuristics for integer linear programs. In *International Conference on Integration of*  
585 *Constraint Programming, Artificial Intelligence, and Operations Research*, pp. 96–113. Springer,  
586 2023a.
- 587  
588 Taoan Huang, Aaron M Ferber, Yuandong Tian, Bistra Dilikina, and Benoit Steiner. Searching large  
589 neighborhoods for integer linear programs with contrastive learning. In *International Conference*  
590 *on Machine Learning*, pp. 13869–13890. PMLR, 2023b.
- 591  
592 Minsu Kim, Jinkyoo Park, et al. Learning collaborative policies to solve np-hard routing problems.  
593 *Advances in Neural Information Processing Systems*, 34:10418–10430, 2021.
- 594  
595 Uroš Klanšek. A comparison between milp and minlp approaches to optimal solution of nonlinear  
596 discrete transportation problem. *Transport*, 30(2):135–144, 2015.
- 597  
598 Wouter Kool, Herke Van Hoof, and Max Welling. Attention, learn to solve routing problems! *arXiv*  
599 *preprint arXiv:1803.08475*, 2018.

- 594 Yeong-Dae Kwon, Jinho Choo, Byoungjip Kim, Iljoo Yoon, Youngjune Gwon, and Seungjai Min.  
595 Pomo: Policy optimization with multiple optima for reinforcement learning. *Advances in Neural*  
596 *Information Processing Systems*, 33:21188–21198, 2020.
- 597 Robert Lange, Yingtao Tian, and Yujin Tang. Large language models as evolution strategies. In  
598 *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp. 579–582,  
599 2024.
- 600  
601 Fei Liu, Tong Xialiang, Mingxuan Yuan, Xi Lin, Fu Luo, Zhenkun Wang, Zhichao Lu, and Qingfu  
602 Zhang. Evolution of heuristics: Towards efficient automatic algorithm design using large language  
603 model. In *Forty-first International Conference on Machine Learning*, 2024.
- 604  
605 Fu Luo, Xi Lin, Fei Liu, Qingfu Zhang, and Zhenkun Wang. Neural combinatorial optimization with  
606 heavy decoder: Toward large scale generalization. *Advances in Neural Information Processing*  
607 *Systems*, 36:8845–8864, 2023.
- 608 Stephen Maher, Matthias Miltenberger, João Pedro Pedroso, Daniel Rehfeldt, Robert Schwarz, and  
609 Felipe Serrano. PySCIPOpt: Mathematical programming in python with the SCIP optimization  
610 suite. In *Mathematical Software – ICMS 2016*, pp. 301–307. Springer International Publishing,  
611 2016. doi: 10.1007/978-3-319-42432-3\_37.
- 612  
613 Setyo Tri Windras Mara, Rachmadi Norcahyo, Panca Jodiawan, Luluk Lusiantoro, and  
614 Achmad Pratama Rifai. A survey of adaptive large neighborhood search algorithms and applica-  
615 tions. *Computers & Operations Research*, 146:105903, 2022.
- 616 David R Morrison, Sheldon H Jacobson, Jason J Sauppe, and Edward C Sewell. Branch-and-bound  
617 algorithms: A survey of recent advances in searching, branching, and pruning. *Discrete Opti-*  
618 *mization*, 19:79–102, 2016.
- 619  
620 Vinod Nair, Mohammad Alizadeh, et al. Neural large neighborhood search. In *Learning Meets*  
621 *Combinatorial Algorithms at NeurIPS2020*, 2020.
- 622 L. Perron and V. Furnon. Or-tools. URL [https://developers.google.com/](https://developers.google.com/optimization/)  
623 [optimization/](https://developers.google.com/optimization/).
- 624  
625 Nikolai Pizaruk. Mixed integer programming: Models and methods. *Minsk: BSU*, 2019.
- 626  
627 Aske Plaat, Annie Wong, Suzan Verberne, Joost Broekens, Niki van Stein, and Thomas Back. Rea-  
628 soning with large language models, a survey. *arXiv preprint arXiv:2407.11511*, 2024.
- 629  
630 Gerhard Reinelt. Tsplib—a traveling salesman problem library. *ORSA journal on computing*, 3(4):  
376–384, 1991.
- 631  
632 Bernardino Romera-Paredes, Mohammadamin Barekatin, Alexander Novikov, Matej Balog,  
633 M Pawan Kumar, Emilien Dupont, Francisco JR Ruiz, Jordan S Ellenberg, Pengming Wang,  
634 Omar Fawzi, et al. Mathematical discoveries from program search with large language models.  
635 *Nature*, 625(7995):468–475, 2024.
- 636  
637 Daniel J Rosenkrantz, Richard E Stearns, and Philip M Lewis, II. An analysis of several heuristics  
638 for the traveling salesman problem. *SIAM journal on computing*, 6(3):563–581, 1977.
- 639  
640 Steven S Seiden. On the online bin packing problem. *Journal of the ACM (JACM)*, 49(5):640–671,  
641 2002.
- 642  
643 Peter W Shor. How to pack better than best fit: tight bounds for average-case online bin packing. In  
644 *[1991] Proceedings 32nd Annual Symposium of Foundations of Computer Science*, pp. 752–759.  
645 IEEE Computer Society, 1991.
- 646  
647 Dan Simon. *Evolutionary optimization algorithms*. John Wiley & Sons, 2013.
- 648  
649 Jialin Song, Yisong Yue, Bistra Dilkina, et al. A general large neighborhood search framework  
650 for solving integer linear programs. *Advances in Neural Information Processing Systems*, 33:  
20012–20023, 2020.

- 648 Nicolas Sonnerat, Pengming Wang, Ira Ktena, Sergey Bartunov, and Vinod Nair. Learning a large  
649 neighborhood search algorithm for mixed integer programs. *arXiv preprint arXiv:2107.10201*,  
650 2021.
- 651 Xueyan Tang, Yusen Li, Runtian Ren, and Wentong Cai. On first fit bin packing for online cloud  
652 server allocation. In *2016 IEEE International Parallel and Distributed Processing Symposium*  
653 *(IPDPS)*, pp. 323–332. IEEE, 2016.
- 654 Robert Endre Tarjan and Anthony E Trojanowski. Finding a maximum independent set. *SIAM*  
655 *Journal on Computing*, 6(3):537–546, 1977.
- 656 Daniel Vázquez, María J Fernández-Torres, Rubén Ruiz-Femenia, Laureano Jiménez, and José A  
657 Caballero. Milp method for objective reduction in multi-objective optimization. *Computers &*  
658 *Chemical Engineering*, 108:382–394, 2018.
- 659 Yaoxin Wu, Wen Song, Zhiguang Cao, and Jie Zhang. Learning large neighborhood search policy  
660 for integer programming. *Advances in Neural Information Processing Systems*, 34:30075–30087,  
661 2021.
- 662 Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V Le, Denny Zhou, and Xinyun  
663 Chen. Large language models as optimizers. In *The Twelfth International Conference on Learning*  
664 *Representations*, 2024. URL <https://openreview.net/forum?id=Bb4VGOWELI>.
- 665 Haoran Ye, Jiarui Wang, Zhiguang Cao, and Guojie Song. Reevo: Large language models as hyper-  
666 heuristics with reflective evolution. *arXiv preprint arXiv:2402.01145*, 2024.
- 667 Huigen Ye, Hongyan Wang, Hua Xu, Chengming Wang, and Yu Jiang. Adaptive constraint partition  
668 based optimization framework for large-scale integer linear programming (student abstract). In  
669 *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 16376–16377,  
670 2023a.
- 671 Huigen Ye, Hua Xu, and Hongyan Wang. Light-milpopt: Solving large-scale mixed integer linear  
672 programs with lightweight optimizer and small-scale training dataset. In *The Twelfth International*  
673 *Conference on Learning Representations*, 2023b.
- 674 Huigen Ye, Hua Xu, Hongyan Wang, Chengming Wang, and Yu Jiang. GNN&GBDT-guided fast  
675 optimizing framework for large-scale integer programming. In *Proceedings of the 40th Interna-*  
676 *tional Conference on Machine Learning*, volume 202, pp. 39864–39878. PMLR, 2023c.
- 677 Jiayi Zhang, Chang Liu, Xijun Li, Hui-Ling Zhen, Mingxuan Yuan, Yawen Li, and Junchi Yan.  
678 A survey for solving mixed integer programming via machine learning. *Neurocomputing*, 519:  
679 205–217, 2023.
- 680 Jianan Zhou, Yaoxin Wu, Zhiguang Cao, Wen Song, Jie Zhang, and Zhenghua Chen. Learning  
681 large neighborhood search for vehicle routing in airport ground handling. *IEEE Transactions on*  
682 *Knowledge and Data Engineering*, 35(9):9769–9782, 2023.
- 683  
684  
685  
686  
687  
688  
689  
690  
691  
692  
693  
694  
695  
696  
697  
698  
699  
700  
701

## APPENDIX

This Appendix contains four sections, each addressing a specific aspect of the experimental setup and results. Below is a brief overview of each section:

- **Parameter Settings** (Appendix A): This section provides detailed descriptions of the key parameters used in our experiments, including the number of top-performing heuristic strategies to evaluate, the threshold for stagnation detection, and the criteria for evolutionary convergence. Specific parameter values for different combinatorial optimization problems, such as Bin Packing (BP), Maximum Vertex Covering (MVC), and Mixed Integer Knapsack Set (MIKS), are also outlined.
- **Evolutionary Process of LLM-LNS** (Appendix B): This section focuses on the evolutionary process of the proposed Dual-layer Self-Evolutionary LLM Agent. It explains how the inner and outer layers of the model co-evolve to generate and refine heuristic strategies across several combinatorial optimization problems. Subsections include detailed comparisons between the Evolution of Heuristic (EoH) method and the proposed dual-layer approach for problems like Bin Packing and Traveling Salesman Problem (TSP).
- **Convergence Analysis of LLM-LNS** (Appendix C): This section presents an analysis of the convergence behavior of the LLM-LNS method compared to the EoH approach. It highlights the superior convergence rates and solution quality achieved by LLM-LNS in both the Online Bin Packing and Traveling Salesman Problem. Graphs and figures demonstrate the evolutionary progress and show how LLM-LNS outperforms EoH in terms of stability and final objective scores.
- **Supplementary Experiments for LLM-LNS on Large-Scale MILP Problems** (Appendix D): This section includes additional experiments where LLM-LNS is applied to large-scale Mixed Integer Linear Programming (MILP) problems. It evaluates the performance of LLM-LNS using different solvers, such as SCIP, and compares the results with various traditional and learning-based methods. The section also includes a detailed analysis of the consistency and stability of each method, as well as error bar comparisons to assess the reliability of the solutions.

These appendices provide a comprehensive overview of the experimental setup, evolutionary process, convergence analysis, and supplementary experiments, offering a deeper understanding of the performance and robustness of the LLM-LNS method in solving complex combinatorial optimization problems.

## A EXPERIMENTAL SETTINGS

In this section, we detail the parameter settings used in our experiments for both the Dual-layer Self-evolutionary LLM Agent and the Adaptive Large Neighborhood Search (ALNS). We also provide an overview of the standard MILP problem instances used in this study.

### A.1 DUAL-LAYER SELF-EVOLUTIONARY LLM AGENT PARAMETERS

The following key parameters were used for the evolutionary process of the LLM agent:

- $k$ : Represents the number of top-performing heuristic strategies used to evaluate each prompt strategy. For each prompt strategy, the top- $k$  heuristics it generates are tracked, and their average fitness score is used as the fitness score for the prompt strategy. In our experiments,  $k$  is set to half of the population size. Specifically:
  - For **Bin Packing (BP)** and **Traveling Salesman Problem (TSP)**, the population sizes are 20 and 10, respectively, so  $k$  is set to 10 and 5.
  - For the four MILP problems—**Maximum Vertex Covering (MVC)**, **Set Covering (SC)**, **Independent Set (IS)**, and **Mixed Integer Knapsack Set (MIKS)**—the population size is 4, so  $k$  is set to 2.
- $l$ : Denotes the number of top individuals in the heuristic population that are monitored for stagnation. If the top- $l$  individuals remain unchanged for  $t$  generations, we infer that the

756 evolution has potentially converged to a local optimum, triggering the introduction of new  
 757 prompt strategies. In all our experiments,  $l$  is set to 4.

- 758 •  $t$ : The number of consecutive generations during which the top- $l$  individuals must remain  
 759 unchanged before stagnation is detected. In all our experiments,  $t$  is set to 3.

## 761 A.2 ADAPTIVE LARGE NEIGHBORHOOD SEARCH (ALNS) PARAMETERS

762 For ALNS, we use the following parameters:

- 763 • **Neighborhood size**  $k$ : Set to half of the decision variable count  $n$ . This represents the  
 764 number of decision variables selected to form the search neighborhood in each iteration.
- 765 • **Time limit**  $T$ : The maximum allowed runtime for solving the problem.
- 766 • **Threshold**  $\epsilon$ : Represents the minimum improvement in the objective function to continue  
 767 exploring the current neighborhood. We set  $\epsilon = 1e-3$ .
- 768 • **Iteration limit**  $p$ : The number of consecutive iterations with improvements below the  
 769 threshold  $\epsilon$  before expanding the neighborhood size. We set  $p = 3$ .
- 770 • **Minimum and maximum neighborhood sizes**  $k_{\min}, k_{\max}$ : These are set to  $k_{\min} = 0$  and  
 771  $k_{\max} = n$  (the total number of decision variables in the problem).
- 772 • **Adjustment rate**  $u\%$ : Specifies the percentage of decision variables  $n$  by which the  
 773 neighborhood size is adjusted during expansion or reduction. In our experiments, we set  
 774  $u\% = 10$ .

## 781 A.3 MILP PROBLEM OVERVIEW

782 We use a set of standard problem instances based on four canonical MILP problems: Maximum  
 783 Independent Set (MIS), Minimum Vertex Covering (MVC), Set Covering (SC), and Mixed Integer  
 784 Knapsack Set (MIKS). Below are the formal definitions of these problems.

785 **Maximum Independent Set problem (MIS)**: The Maximum Independent Set problem has applica-  
 786 tions in network design, where one might need to select the largest subset of mutually non-interacting  
 787 entities, such as devices in a wireless network to avoid interference. Another common application  
 788 is in social network analysis, where independent sets can represent groups of users who do not have  
 789 direct connections, useful for targeting non-overlapping communities.

790 Consider an undirected graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , where a subset of nodes  $\mathcal{S} \subseteq \mathcal{V}$  is called an independent  
 791 set if no edge  $e \in \mathcal{E}$  exists between any pair of nodes in  $\mathcal{S}$ . The MIS problem seeks to find an  
 792 independent set of maximum cardinality. The binary decision variable  $x_v$  indicates whether node  
 793  $v \in \mathcal{V}$  is part of the independent set ( $x_v = 1$ ) or not ( $x_v = 0$ ). The problem can be formulated as:

$$\begin{aligned}
 & \max \sum_{v \in \mathcal{V}} x_v \\
 & \text{s.t. } x_u + x_v \leq 1, \quad \forall (u, v) \in \mathcal{E}, \\
 & \quad x_v \in \{0, 1\}, \quad \forall v \in \mathcal{V}.
 \end{aligned} \tag{2}$$

802 **Minimum Vertex Covering problem (MVC)**: The Minimum Vertex Covering problem is widely  
 803 used in resource allocation, where one needs to ensure that every interaction (edge) between pairs  
 804 of objects (nodes) is covered by a resource. For example, in network security, this problem can be  
 805 used to efficiently place security agents or sensors such that all communication links are monitored.

806 Given an undirected graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , a subset of nodes  $\mathcal{S} \subseteq \mathcal{V}$  is called a covering set if for any  
 807 edge  $e \in \mathcal{E}$ , at least one of its endpoints is included in  $\mathcal{S}$ . The MVC problem aims to find a covering  
 808 set of minimum cardinality. The binary decision variable  $x_v$  indicates whether node  $v \in \mathcal{V}$  is part  
 809 of the covering set ( $x_v = 1$ ) or not ( $x_v = 0$ ). The problem is formulated as:

$$\begin{aligned}
& \min \sum_{v \in \mathcal{V}} x_v \\
& \text{s.t. } x_u + x_v \geq 1, \quad \forall (u, v) \in \mathcal{E}, \\
& \quad x_v \in \{0, 1\}, \quad \forall v \in \mathcal{V}.
\end{aligned} \tag{3}$$

**Set Covering problem (SC):** The Set Covering problem is fundamental in facility location, where one must select the minimum number of locations (subsets) to serve all customers (elements of the universal set). It is also used in airline crew scheduling, where the goal is to assign the minimum number of crews to cover all flights.

Given a finite universal set  $\mathcal{U} = \{1, 2, \dots, n\}$  and a collection of  $m$  subsets  $S_1, \dots, S_m$  of  $\mathcal{U}$ , each subset  $S_i$  is associated with a cost  $c_i$ . The SC problem involves selecting a combination of these subsets such that every element in  $\mathcal{U}$  is covered by at least one of the selected subsets, while minimizing the total cost. The binary decision variable  $x_i$  indicates whether subset  $S_i$  is selected ( $x_i = 1$ ) or not ( $x_i = 0$ ). The problem is formulated as:

$$\begin{aligned}
& \min \sum_{i=1}^m c_i x_i \\
& \text{s.t. } \sum_{i=1}^m x_i \cdot \mathbf{1}_{\{j \in S_i\}} \geq 1, \quad \forall j \in \mathcal{U}, \\
& \quad x_i \in \{0, 1\}, \quad \forall i \in \{1, \dots, m\}.
\end{aligned} \tag{4}$$

**Mixed Integer Knapsack Set problem (MIKS):** The Mixed Integer Knapsack Set problem is commonly used in logistics, resource allocation, and portfolio selection problems. It models situations where some resources can be allocated fractionally while others must be fully included or excluded. For example, in supply chain management, some goods can be shipped partially, while others must be shipped as a whole.

The MIKS problem is a generalization of the knapsack problem that involves both continuous and binary decision variables. Given  $N$  sets and  $M$  items, each item must be covered by at least one of the sets. The objective is to minimize the total cost of the selected sets, where some sets can be partially selected. Let  $x_i$  represent the decision variable for set  $i$ , where  $x_i = 1$  indicates full selection, and  $0 \leq x_i \leq 1$  allows partial selection. The problem is formulated as:

$$\begin{aligned}
& \min \sum_{i=1}^N c_i x_i \\
& \text{s.t. } \sum_{i: j \in S_i} x_i \geq 1, \quad \forall j \in \{1, 2, \dots, M\}, \\
& \quad 0 \leq x_i \leq 1, \quad \forall i \in \{1, 2, \dots, N\}, \\
& \quad x_i \in \{0, 1\} \text{ or } [0, 1], \quad \forall i \in \{1, 2, \dots, N\}.
\end{aligned} \tag{5}$$

## B EVOLUTIONARY PROCESS OF LLM-LNS

### B.1 EVOLUTIONARY PROCESS OVERVIEW

In this appendix, we provide a detailed breakdown of the experimental results and the evolution of heuristic strategies generated by our proposed **Dual-layer Self-Evolutionary LLM Agent**. The following sections offer a comprehensive analysis of how the inner and outer layers of the LLM agent collaborate to generate and refine heuristic strategies across various combinatorial optimization problems, including **Online Bin Packing (bp\_online)**, the **Traveling Salesman Problem (TSP)**, and large-scale MILP instances such as **Maximum Vertex Covering (MVC)**, **Set Covering (SC)**, **Independent Set (IS)**, and **Mixed Integer Knapsack Set (MIKS)**.



- 864
- 865
- 866
- 867
- 868
- 869
- 870
- 871
- 872
- 873
- 874
- **Inner and Outer Layer Prompt Initialization and Evolution:** As shown in Sec. B.2, our approach leverages a dual-layer architecture, where the **inner layer** evolves heuristic strategies by modifying solution components, while the **outer layer** evolves the prompt structure guiding the inner layer, balancing exploration and exploitation. The inner layer prompts iteratively generate heuristics by scoring decision variables based on their contributions to the objective function and constraints, with randomness included to avoid local optima. This enables the LLM to reason about the problem structure and generate high-quality strategies, even without extensive domain expertise. The **outer layer** maintains diversity by evolving prompt structures to prevent premature convergence on suboptimal solutions. Both layers adapt based on past performance, allowing the LLM to refine its strategy generation over time.
  - **Heuristic Improvement Through Dual-layer Self-evolutionary LLM Agent:** As shown in Sec. B.3, we demonstrates the progression of heuristic strategies, starting from initial random strategies and gradually evolving into more effective ones through the dual-layer self-evolutionary process. The initial strategies are simple and focus on ranking decision variables based on their contributions to the objective function and constraints. Over time, the LLM agent introduces additional complexity, such as incorporating randomness and penalizing larger deviations from the current solution, improving the robustness of the generated heuristics. The progression of the population is guided by the outer layer, which adjusts the structure and focus of prompts to encourage exploration and avoid premature convergence. The inner layer then refines specific solution components in response to the prompts, iteratively improving the performance of the heuristic strategies. As seen from the evolution of objective scores, the dual-layer system enables the generation of increasingly effective heuristics, balancing exploration with exploitation to achieve superior results in various problem instances.
  - **Heuristic Strategies for Bin Packing Online: EoH vs. Dual-Layer Self-Evolution LLM Agent:** As shown in Sec. B.4, both the *Evolution of Heuristic (EoH)* method and our Dual-layer Self-Evolution LLM Agent utilize LLM-based evolutionary processes to generate heuristic strategies for the *Bin Packing Online* problem. The strategy generated by EoH approach, while leveraging LLM to evolve heuristics, focuses primarily on a hybrid scoring system that combines utilization ratios, dynamic adjustments, and an exponentially decaying factor. This method is effective but tends to rely on a more static set of features and parameters, which limits its adaptability across diverse problem instances. In contrast, our Dual-layer Self-Evolution LLM Agent incorporates a more dynamic and adaptive strategy. By combining nonlinear capacity scaling, relative size assessment, and historical penalties for overutilized bins, our approach allows for greater flexibility and adaptability. Specifically, the generated heuristics dynamically adjust based on remaining capacity, item size, and previous bin usage, thereby balancing local search with global optimization. This adaptability enables our agent to discover and refine more efficient strategies that minimize the number of bins used. The results clearly demonstrate that while both methods use LLM-based evolution, our dual-layer approach consistently outperforms the EoH method in terms of solution quality and computational efficiency. The dual-layer system’s ability to evolve both the heuristic strategies and the prompt structures ensures that it can fine-tune solutions more effectively, leading to superior bin utilization and fewer bins required overall. This highlights the strength of our approach in generating more robust and context-aware heuristics.
  - **Heuristic Strategies for Traveling Salesman Problem (TSP): EoH vs. Dual-Layer Self-Evolution LLM Agent:** Similar to the *Bin Packing Online* problem, both the *Evolution of Heuristic (EoH)* method and our Dual-layer Self-Evolution LLM Agent use LLM-based evolutionary processes to generate heuristic strategies for the *Traveling Salesman Problem (TSP)*. As shown in Sec. B.5, the strategy generated by EoH method employs a randomized approach that adjusts the edge distance matrix by increasing the distances of a random proportion of edges, while rewarding a smaller subset of unused edges. This method encourages exploration but tends to apply uniform adjustments without fully accounting for the global structure of the solution. In contrast, strategy generated by our Dual-layer Self-Evolution LLM Agent introduces a more sophisticated edge distance adjustment mechanism. It dynamically explores alternative routes by incorporating an inverse frequency factor, which penalizes frequently used edges and rewards less frequently used ones. This
- 875
- 876
- 877
- 878
- 879
- 880
- 881
- 882
- 883
- 884
- 885
- 886
- 887
- 888
- 889
- 890
- 891
- 892
- 893
- 894
- 895
- 896
- 897
- 898
- 899
- 900
- 901
- 902
- 903
- 904
- 905
- 906
- 907
- 908
- 909
- 910
- 911
- 912
- 913
- 914
- 915
- 916
- 917

918 adaptive mechanism gradually resets excessively amplified distances, promoting diversi-  
919 fication and improving the exploration of the solution space. Furthermore, it balances  
920 exploitation by focusing on refining the most promising routes based on past tours, lead-  
921 ing to faster convergence towards a global optimum. The results clearly demonstrate that  
922 while both methods are effective in exploring new routes, the dual-layer approach consis-  
923 tently outperforms the EoH method in terms of solution quality and convergence speed.  
924 By incorporating a more nuanced edge adjustment process and dynamically adapting to  
925 the problem context, the Dual-layer Self-Evolution LLM Agent achieves superior results  
926 in minimizing the total distance, making it a more robust and efficient solution for the TSP.

- 927 • **Evolutionary Path of the Dual-Layer Self-Evolution LLM Agent:** As illustrated in Sec.  
928 B.6, we trace the evolutionary process of the LLM agent in solving Maximum Vertex Cover  
929 (MVC) problem, detailing how heuristic strategies evolve step by step through the inner  
930 and outer layers, gradually converging to optimized solutions. Initially, the agent gener-  
931 ates simple heuristics that focus on ranking decision variables based on their impact on  
932 the objective function and constraint violation, incorporating randomness to encourage  
933 exploration. These early strategies serve as a foundation for further refinement. As the  
934 process evolves, the outer layer refines the prompt instructions, guiding the inner layer to  
935 develop more sophisticated heuristics. The LLM begins to incorporate additional factors,  
936 such as the absolute difference from the initial solution and a more nuanced treatment of  
937 constraints. This results in improved exploration of the solution space, as well as better  
938 handling of both the objective function and constraints. In the later stages, the agent in-  
939 tegrates more advanced techniques, such as hybrid methods combining genetic algorithms  
940 with local search, to enhance convergence speed and solution quality. The final heuristics  
941 represent a co-evolutionary approach that balances exploration and exploitation, leading  
942 to significantly optimized solutions. The evolution of prompts, from the initial simplistic  
943 forms to highly specialized instructions, demonstrates the power of the dual-layer architec-  
944 ture in improving both the heuristic strategies and the problem-solving process itself.
- 945 • **Evolutionary Result of the Dual-Layer Self-Evolution LLM Agent:** Finally, we present  
946 the results achieved by the LLM agent after the completion of the entire evolutionary pro-  
947 cess across three challenging combinatorial optimization problems: Set Covering (SC),  
948 Maximum Independent Set (MIS), and Mixed Integer Knapsack Set (MIKS). As detailed  
949 in Sec. B.7, the final heuristics generated by the Dual-layer Self-Evolution LLM Agent  
950 are compared with those produced by traditional methods and state-of-the-art approaches,  
951 demonstrating significant improvements in solution quality and computational efficiency.  
952 For the Set Covering problem (SC), the LLM agent’s final heuristic achieves a superior  
953 balance between minimizing the number of selected sets and satisfying the constraints. By  
954 dynamically adjusting penalties and incorporating random exploration, the agent efficiently  
955 navigates the solution space, outperforming traditional methods in both the objective score  
956 and constraint satisfaction. In the Maximum Independent Set (MIS) problem, the LLM  
957 agent leverages simulated annealing principles combined with adaptive scoring of decision  
958 variables. This approach not only ensures thorough exploration but also accelerates conver-  
959 gence towards high-quality solutions. The agent’s ability to balance objective contributions  
960 with constraint violations leads to a considerable reduction in the total error, as reflected in  
961 the final objective score. Lastly, for the Mixed Integer Knapsack Set (MIKS) problem, the  
962 LLM agent adopts a hybrid strategy that integrates genetic algorithms and simulated an-  
963 nealing. This allows for a more diversified search process, strategically selecting decision  
964 variables based on their contributions to the objective function and constraint interactions.  
965 The agent’s solution demonstrates a significant improvement over existing methods, par-  
966 ticularly in how it dynamically adapts to varying problem constraints while maintaining  
967 computational efficiency.

966 In summary, the proposed Dual-layer Self-Evolutionary LLM Agent exhibits significant advantages  
967 in generating and refining heuristic strategies across a wide range of combinatorial optimization  
968 problems. By leveraging the complementary roles of the inner and outer layers, the system is able to  
969 balance exploration and exploitation effectively, allowing for the discovery of high-quality, context-  
970 aware heuristic strategies. The adaptability of the dual-layer approach, particularly its ability to  
971 evolve both the problem-solving heuristics and the prompt structures that guide them, leads to su-  
perior performance in terms of both solution quality and computational efficiency. Across different

problem domains, from online bin packing to large-scale MILP problems, the LLM agent consistently outperforms traditional methods and state-of-the-art approaches, demonstrating its robustness, scalability, and evolutionary refinement capabilities.

## B.2 INNER AND OUTER LAYER PROMPT INITIALIZATION AND EVOLUTION

### Prompt for Generating Initial Heuristic Strategies

Given an initial feasible solution and a current solution to a Mixed-Integer Linear Programming (MILP) problem, with variables' lower\_bound, upper\_bound and coefficient in objective function. We want to improve the current solution using Large Neighborhood Search (LNS).

The task can be solved step-by-step by starting from the current solution and iteratively selecting a subset of decision variables to relax and re-optimize. In each step, most decision variables are fixed to their values in the current solution, and only a small subset is allowed to change. You need to score all the decision variables based on the information I give you, and I will choose the decision variables with high scores as neighborhood selection. To avoid getting stuck in local optima, the choice of the subset can incorporate a degree of randomness.

First, describe your new algorithm and main steps in one sentence. The description must be inside a brace. Next, implement it in Python as a function named `select_neighborhood`. This function should accept 5 input(s): `'initial_solution'`, `'current_solution'`, `'lower_bound'`, `'upper_bound'`, `'objective_coefficient'`. The function should return 1 output(s): `'neighbor_score'`. `'initial_solution'`, `'current_solution'`, `'lower_bound'`, `'upper_bound'` and `'objective_coefficient'` are numpy arrays. `'neighbor_score'` is also a numpy array that you need to create manually. The *i*-th element of the arrays corresponds to the *i*-th decision variable. All are Numpy arrays. I don't give you `'neighbor_score'` so that you need to create it manually. The length of the `'neighbor_score'` array is the same as the length of the other arrays.

Do not give additional explanations.

### (Cross) Initial Prompt for Heuristic Strategies Evolution

Given an initial feasible solution and a current solution to a Mixed-Integer Linear Programming (MILP) problem, with variables' lower\_bound, upper\_bound and coefficient in objective function. We want to improve the current solution using Large Neighborhood Search (LNS).

The task can be solved step-by-step by starting from the current solution and iteratively selecting a subset of decision variables to relax and re-optimize. In each step, most decision variables are fixed to their values in the current solution, and only a small subset is allowed to change. You need to score all the decision variables based on the information I give you, and I will choose the decision variables with high scores as neighborhood selection. To avoid getting stuck in local optima, the choice of the subset can incorporate a degree of randomness.

I have 5 existing algorithm's thought, objective function value with their codes as follows: No.1 algorithm's thought, objective function value, and the corresponding code are: ...

No.2 algorithm's thought, objective function value, and the corresponding code are: ...

...

No.5 algorithm's thought, objective function value, and the corresponding code are: ...

Please help me create a new algorithm that has a totally different form from the given ones.

First, describe your new algorithm and main steps in one sentence. The description must be inside a brace. Next, implement it in Python as a function named `select_neighborhood`. This function should accept 5 input(s): `'initial_solution'`, `'current_solution'`, `'lower_bound'`, `'upper_bound'`, `'objective_coefficient'`. The function should return 1 output(s): `'neighbor_score'`. `'initial_solution'`, `'current_solution'`, `'lower_bound'`, `'upper_bound'` and `'objective_coefficient'` are numpy arrays. `'neighbor_score'` is also a numpy array that you need to create manually. The *i*-th element of the arrays corresponds to the *i*-th decision variable. All are Numpy arrays. I don't give you `'neighbor_score'` so that you need to create it manually. The length of the `'neighbor_score'` array is the same as the length of the other arrays.

Do not give additional explanations.

### (Cross) Initial Prompt Strategies

1. Please help me create a new algorithm that has a totally different form from the given ones.
2. Please help me create a new algorithm that has a totally different form from the given ones but can be motivated from them.

### (Cross) Prompt for Prompt Strategies Evolution

We are working on solving a minimization problem. Our objective is to leverage the capabilities of the Language Model (LLM) to generate heuristic algorithms that can efficiently tackle this problem. We have already developed a set of initial prompts and observed the corresponding outputs. However, to improve the effectiveness of these algorithms, we need your assistance in carefully analyzing the existing prompts and their results. Based on this analysis, we ask you to generate new prompts that will help us achieve better outcomes in solving the minimization problem.

I have 5 existing prompts with objective function value as follows:

No.1 prompt's tasks assigned to LLM, and objective function value are: ...

No.2 prompt's tasks assigned to LLM, and objective function value are: ...

...

No.5 prompt's tasks assigned to LLM, and objective function value are: ...

Please help me create a new prompt that has a totally different form from the given ones but can be motivated from them.

Please describe your new prompt and main steps in one sentence. Do not give additional explanations.

### B.3 HEURISTIC IMPROVEMENT THROUGH DUAL-LAYER SELF-EVOLUTIONARY LLM AGENT

#### Heuristic 1 (Obj Score: 5375.52145)

Rank decision variables based on their penalty contribution and the difference from current solution, incorporating randomness in scoring.

```
import numpy as np
def select_neighborhood(n, m, k, site, value,
                       constraint, initial_solution,
                       current_solution, objective_coefficient):
    neighbor_score = np.zeros(n)
    variable_difference = np.zeros(n)
    for i in range(m):
        lhs = sum(value[i][j] * current_solution[
            site[i][j]] for j in range(k[i]))
        penalty = max(0, lhs - constraint[i])
        for j in range(k[i]):
            var_index = site[i][j]
            difference = current_solution[
                var_index] - initial_solution[
                var_index]
            neighbor_score[var_index] += penalty *
                difference
    neighbor_score += objective_coefficient * np.
        random.rand(n)
    return neighbor_score
```

#### Heuristic 2 (Obj Score: 5383.05876)

Rank decision variables based on their objective contribution and impact on current solution deviation, with randomness included in the scoring process.

```
import numpy as np
def select_neighborhood(n, m, k, site, value,
                       constraint, initial_solution,
                       current_solution, objective_coefficient):
    neighbor_score = np.zeros(n)
    variable_contribution = np.zeros(n)
    for i in range(m):
        lhs = sum(value[i][j] * current_solution[
            site[i][j]] for j in range(k[i]))
        deviation = lhs - constraint[i]
        for j in range(k[i]):
            var_index = site[i][j]
            contribution = value[i][j] * (
                initial_solution[var_index] -
                current_solution[var_index])
            neighbor_score[var_index] +=
                contribution
    neighbor_score += objective_coefficient * np.
        random.rand(n)
    return neighbor_score
```

#### Heuristic 3 (Obj Score: 5384.8486)

This modified algorithm ranks decision variables based on their contribution to the total current solution's objective function value and their degree of constraint satisfaction.

```
import numpy as np
def select_neighborhood(n, m, k, site, value,
                       constraint, initial_solution,
                       current_solution, objective_coefficient):
    neighbor_score = np.zeros(n)
    for i in range(m):
        lhs = sum(value[i][j] * current_solution[
            site[i][j]] for j in range(k[i]))
        for j in range(k[i]):
            if lhs > constraint[i]:
                neighbor_score[site[i][j]] +=
                    objective_coefficient[site[i]
                        ][j]] * (lhs - constraint[i])
            else:
                neighbor_score[site[i][j]] +=
                    objective_coefficient[site[i]
                        ][j]] * (constraint[i] - lhs)
    neighbor_score += np.random.rand(n) * 0.1
    return neighbor_score
```

#### Heuristic 4 (Obj Score: 5384.95417)

Rank decision variables by their contribution to the objective function and difference from initial values, while also weighing their frequency of use in the constraints.

```
import numpy as np
def select_neighborhood(n, m, k, site, value,
                       constraint, initial_solution,
                       current_solution, objective_coefficient):
    score = np.zeros(n)
    frequency = np.zeros(n)
    for i in range(m):
        lhs = sum(value[i][j] * current_solution[
            site[i][j]] for j in range(k[i]))
        deviation = lhs - constraint[i]
        for j in range(k[i]):
            var_index = site[i][j]
            contribution = value[i][j] * np.abs(
                initial_solution[var_index] -
                current_solution[var_index])
            score[var_index] += contribution
            frequency[var_index] += 1
    neighbor_score = score / (frequency + 1e-5) +
        objective_coefficient * np.random.rand(n)
    return neighbor_score
```

#### Prompt Designed by LLM

Develop an algorithm that combines the strengths of existing heuristics while introducing random perturbations to enhance exploration and minimize the objective function more effectively.

#### Heuristic (Obj Score: 5374.19865)

Rank decision variables based on their contribution to the objective function and incorporate the absolute difference from the initial solution while adding a degree of randomness to the scores.

```
import numpy as np
def select_neighborhood(n, m, k, site, value, constraint, initial_solution, current_solution,
                       objective_coefficient):
    neighbor_score = np.zeros(n)
    for i in range(m):
        lhs = sum(value[i][j] * current_solution[site[i][j]] for j in range(k[i]))
        for j in range(k[i]):
            var_index = site[i][j]
            difference = np.abs(current_solution[var_index] - initial_solution[var_index])
            neighbor_score[var_index] += (constraint[i] - lhs) * difference
    neighbor_score += objective_coefficient * np.random.rand(n)
    return neighbor_score
```

#### B.4 HEURISTIC STRATEGIES FOR BIN PACKING ONLINE: EOH VS. DUAL-LAYER SELF-EVOLUTION LLM AGENT

##### Heuristic Designed by EoH

###### Description

The heuristic incorporates a weighted average of the utilization ratio, dynamic adjustment, and an exponentially decaying factor, with different parameter settings to minimize the number of used bins.

###### Code

```
import numpy as np
def heuristic(item, bins):
    diff = bins-item # remaining capacity
    exp = np.exp(diff) # exponent term
    sqrt = np.sqrt(diff) # square root term
    ulti = 1-diff/bins # utilization term
    comb = ulti * sqrt # combination of utilization and square root
    adjust = np.where(diff > (item * 3), comb + 0.8, comb + 0.3)
    # hybrid adjustment term to penalize large bins
    hybrid_exp = bins / ((exp + 0.7) *exp)
    # hybrid score based on exponent term
    scores = hybrid_exp + adjust
    # sum of hybrid score and adjustment
    return scores
```

##### Heuristic Designed by Dual-layer Self-evolution LLM Agent

###### Description

The new algorithm employs a hybrid optimization strategy that combines nonlinear penalties for historical usage, adaptive capacity scaling, and a relative size assessment, facilitating a balance between local and global search for optimal bin assignment.

###### Code

```
import numpy as np
def score(item, bins):
    feasible_bins = bins[bins > item]
    scores = np.zeros_like(bins)
    if len(feasible_bins) == 0:
        return scores
    # Nonlinear capacity scaling that enhances the desire for larger
    # spaces
    remaining_capacity = feasible_bins - item
    capacity_scaling = np.log1p(remaining_capacity) * (
        remaining_capacity / np.max(remaining_capacity))
    # Relative size assessment: quadratic term comparing item size with
    # bin capacities
    relative_size_effect = (item ** 2 / feasible_bins) * 50 # Scale to
    # moderate impact
    # Nonlinear penalty based on historical usage counts to deter
    # overutilization
    historical_count = np.arange(len(feasible_bins)) + 1 # Simulating
    # historical usage
    penalty_factor = np.power(1.5, historical_count) # Exponential
    # penalty for higher usage
    # Combining scores: enhanced capacity scaling, moderated size
    # assessment, and historical penalties
    scores[bins > item] = capacity_scaling - relative_size_effect -
    penalty_factor
    return scores
```

## B.5 HEURISTIC STRATEGIES FOR TRAVELING SALESMAN PROBLEM: EOH VS. DUAL-LAYER SELF-EVOLUTION LLM AGENT

### Heuristic Designed by EoH

#### Description

This algorithm uses a randomized approach to update the edge distance matrix by randomly selecting a proportion of edges to increase their distances while uniformly rewarding a smaller proportion of unused edges to encourage exploration.

#### Code

```

import numpy as np
def update_edge_distance(edge_distance, local_opt_tour, edge_n_used):
    N = edge_distance.shape[0]
    updated_edge_distance = edge_distance.copy()
    # Parameters for randomization
    increase_factor = 2.0
    decrease_factor = 0.9
    random_selection_ratio = 0.3 # percentage of edges to randomly adjust
    # Identify all edges used in the local optimal tour
    used_edges = set()
    for i in range(len(local_opt_tour)):
        start = local_opt_tour[i]
        end = local_opt_tour[(i + 1) % len(local_opt_tour)]
        used_edges.add((min(start, end), max(start, end)))
    # Randomly select a proportion of edges to increase distance
    all_edges = [(i, j) for i in range(N) for j in range(N) if i != j]
    np.random.shuffle(all_edges)
    num_edges_to_increase = int(len(all_edges) * random_selection_ratio)
    for edge in all_edges[:num_edges_to_increase]:
        start, end = edge
        # If the edge is used in the local optimal tour, apply a higher increase
        if (min(start, end), max(start, end)) in used_edges:
            updated_edge_distance[start, end] *= increase_factor
            updated_edge_distance[end, start] *= increase_factor
        else:
            updated_edge_distance[start, end] *= decrease_factor
            updated_edge_distance[end, start] *= decrease_factor
    return updated_edge_distance

```

### Heuristic Designed by Dual-layer Self-evolution LLM Agent

#### Description

The new algorithm refines the edge distance adjustment mechanism by incorporating an acceptance heuristic that dynamically explores alternative routes while gradually resetting excessively amplified distances, thus promoting diversification and improved convergence towards a global optimum.

#### Code

```

import numpy as np
def update_edge_distance(edge_distance, local_opt_tour, edge_n_used):
    # Create a copy of the edge distance matrix for updates
    updated_edge_distance = np.copy(edge_distance)
    # Extract the number of nodes
    num_nodes = edge_distance.shape[0]
    # Calculate the inverse frequency factor for each edge
    inverse_frequency_factor = np.max(edge_n_used) - edge_n_used + 1
    # Update the edge distance based on the local optimal tour
    for i in range(len(local_opt_tour)):
        # Get the current and next node in the local optimal tour
        current_node = local_opt_tour[i]
        next_node = local_opt_tour[(i + 1) % len(local_opt_tour)]
        # Apply the inverse frequency factor to decrease the edge weight
        updated_edge_distance[current_node, next_node] *= inverse_frequency_factor[
            current_node, next_node]
        updated_edge_distance[next_node, current_node] *= inverse_frequency_factor[
            next_node, current_node]
    return updated_edge_distance

```

## B.6 EVOLUTIONARY PATH OF THE DUAL-LAYER SELF-EVOLUTION LLM AGENT

1188  
1189  
1190  
1191  
1192  
1193  
1194  
1195  
1196  
1197  
1198  
1199  
1200  
1201  
1202  
1203  
1204  
1205  
1206  
1207  
1208  
1209  
1210  
1211  
1212  
1213  
1214  
1215  
1216  
1217  
1218  
1219  
1220  
1221  
1222  
1223  
1224  
1225  
1226  
1227  
1228  
1229  
1230  
1231  
1232  
1233  
1234  
1235  
1236  
1237  
1238  
1239  
1240  
1241

**Heuristic (Obj Score: 5400.48176)**

The algorithm ranks decision variables based on their impact on the objective function and how they relate to the violated constraints, incorporating a degree of randomness.

**Code**

```
import numpy as np
def select_neighborhood(n, m, k, site, value, constraint,
    initial_solution, current_solution, objective_coefficient):
    neighbor_score = np.zeros(n)
    violated_constraints = 0
    for i in range(m):
        lhs = sum(value[i][j] * current_solution[site[i][j]] for
            j in range(k[i]))
        if lhs > constraint[i]:
            violated_constraints += 1
            for j in range(k[i]):
                neighbor_score[site[i][j]] +=
                    objective_coefficient[site[i][j]]
    if violated_constraints > 0:
        neighbor_score /= violated_constraints
        randomness = np.random.rand(n) * 0.1
        neighbor_score += randomness
    return neighbor_score
```

**Heuristic (Obj Score: 5374.19865)**

Rank decision variables based on their contribution to the objective function and incorporate the absolute difference from the initial solution while adding a degree of randomness to the scores.

**Code**

```
import numpy as np
def select_neighborhood(n, m, k, site, value, constraint,
    initial_solution, current_solution, objective_coefficient):
    neighbor_score = np.zeros(n)
    for i in range(m):
        lhs = sum(value[i][j] * current_solution[site[i][j]] for
            j in range(k[i]))
        for j in range(k[i]):
            var_index = site[i][j]
            difference = np.abs(current_solution[var_index] -
                initial_solution[var_index])
            neighbor_score[var_index] += (constraint[i] - lhs) *
                difference
    neighbor_score += objective_coefficient * np.random.rand(n)
    return neighbor_score
```

**Heuristic (Obj Score: 5373.34904)**

Develop a co-evolutionary heuristic approach that integrates genetic algorithms with local search techniques to enhance convergence speed and minimize the objective function for the specified problem.

**Code**

```
import numpy as np
def select_neighborhood(n, m, k, site, value, constraint,
    initial_solution, current_solution, objective_coefficient):
    neighbor_score = np.zeros(n)
    for i in range(m):
        lhs = sum(value[i][j] * current_solution[site[i][j]] for
            j in range(k[i]))
        for j in range(k[i]):
            var_index = site[i][j]
            difference = np.abs(current_solution[var_index] -
                initial_solution[var_index])
            neighbor_score[var_index] += (constraint[i] - lhs) *
                difference
    random_adjustment = np.random.rand(n)
    adaptive_mutation_rate = np.clip(np.abs(objective_coefficient
        ), 0.1, 1.0)
    neighbor_score += adaptive_mutation_rate * random_adjustment
    return neighbor_score
```

**Initial Prompts**

- (Cross) Please help me create a new algorithm that has a totally different form from the given ones.
- (Cross) Please help me create a new algorithm that has a totally different form from the given ones but can be motivated from them.
- (Variation) Please assist me in creating a new algorithm that has a different form but can be a modified version of the algorithm provided.
- (Variation) Please identify the main algorithm parameters and assist me in creating a new algorithm that has a different parameter settings of the score function provided.

**Current Prompts**

- (Cross) Develop a modified heuristic algorithm that utilizes a hybrid approach, combining elements of simulated annealing and genetic algorithms, to optimize the given minimization problem.
- (Cross) Design a modified heuristic algorithm for the minimization problem by incorporating elements of simulated annealing with a unique cooling schedule.
- (Variation) Please identify the main algorithm parameters and assist me in creating a new algorithm that has a different parameter settings of the score function provided.
- (Variation) Develop an algorithm that combines the strengths of existing heuristics while introducing random perturbations to enhance exploration and minimize the objective function more effectively.

**Final Prompts**

- (Cross) Develop a hybrid heuristic algorithm for the minimization problem that combines genetic algorithms with tabu search to enhance local search capabilities while maintaining diversity in the solution population.
- (Cross) Develop a co-evolutionary heuristic approach that integrates genetic algorithms with local search techniques to enhance convergence speed and minimize the objective function for the specified problem.
- (Variation) Design a novel optimization strategy that integrates genetic algorithms with dynamic programming principles to enhance the search for optimal solutions, focusing on adaptive mutation rates to effectively minimize the objective function value.
- (Variation) Design a novel optimization framework that integrates particle swarm optimization with genetic algorithms, focusing on adaptive mutation strategies to enhance convergence speed and minimize the objective function value.

## B.7 EVOLUTIONARY RESULT OF THE DUAL-LAYER SELF-EVOLUTION LLM AGENT

### B.7.1 EVOLUTIONARY RESULT OF SET COVERING PROBLEM

#### Heuristic (Obj Score: 3339.39339)

This algorithm computes scores based on the penalty incurred by each variable when deviating from the current solution and evaluates the impact on constraint satisfaction.

#### Code

```
import numpy as np
def select_neighborhood(n, m, k, site, value, constraint, initial_solution, current_solution,
    objective_coefficient):
    neighbor_score = np.zeros(n)
    for i in range(m):
        lhs_value = sum(value[i][j] * current_solution[site[i][j]] for j in range(k[i]))
        for j in range(k[i]):
            variable_index = site[i][j]
            if lhs_value >= constraint[i]:
                penalty = lhs_value - constraint[i]
                contribution = penalty * value[i][j]
                neighbor_score[variable_index] += contribution
            else:
                contribution = value[i][j]
                neighbor_score[variable_index] -= contribution
    costs = np.abs(current_solution - initial_solution) * (objective_coefficient + 1e-5)
    with np.errstate(divide='ignore', invalid='ignore'):
        neighbor_score = np.divide(neighbor_score, costs, where=costs != 0)
    neighbor_score -= np.min(neighbor_score)
    neighbor_score /= np.max(neighbor_score) if np.max(neighbor_score) != 0 else 1
    rand_factor = np.random.rand(n) * 0.1
    neighbor_score += rand_factor
    return neighbor_score
```

#### Final Prompts

- (Cross) Please help me create a new algorithm that has a totally different form from the given ones.
- (Cross) Please help me create a new algorithm that has a totally different form from the given ones but can be motivated from them.
- (Variation) Please assist me in creating a new algorithm that has a different form but can be a modified version of the algorithm provided.
- (Variation) Please identify the main algorithm parameters and assist me in creating a new algorithm that has a different parameter settings of the score function provided.

### B.7.2 EVOLUTIONARY RESULT OF MAXIMUM INDEPENDENT SET PROBLEM

#### Heuristic (Obj Score: -4634.0636)

This new heuristic approach combines the principles of simulated annealing with the adaptive scoring of decision variables based on their contributions to violated constraints while incorporating randomness to enhance exploration of the solution space.

#### Code

```
import numpy as np
def select_neighborhood(n, m, k, site, value, constraint, initial_solution, current_solution,
    objective_coefficient):
    neighbor_score = np.zeros(n)
    current_objective_value = np.dot(current_solution, objective_coefficient)
    variable_contributions = np.zeros(n)
    for i in range(m):
        lhs_value = sum(value[i][j] * current_solution[site[i][j]] for j in range(k[i]))
        if lhs_value > constraint[i]:
            for j in range(k[i]):
                var_index = site[i][j]
                variable_contributions[var_index] += (value[i][j] * (current_solution[var_index] == 1))
    for index in range(n):
        improvement = objective_coefficient[index] - variable_contributions[index]
        neighbor_score[index] = improvement + (current_solution[index] * 0.5)
    temperature = np.random.uniform(0.1, 1.0)
    randomness = np.random.uniform(-temperature, temperature, size=n)
    neighbor_score += randomness
    return neighbor_score
```

#### Final Prompts

- (Cross) Develop a novel hybrid algorithm that combines local search and simulated annealing techniques to explore the solution space and minimize the objective function more effectively.
- (Cross) Design a novel optimization algorithm inspired by the existing methods, focusing on adaptive parameter tuning to enhance convergence toward better solutions.
- (Variation) Design a novel heuristic approach inspired by the principles of simulated annealing to optimize the following problem parameters.
- (Variation) Please identify the main algorithm parameters and assist me in creating a new algorithm that has a different parameter settings of the score function provided.



### B.7.3 EVOLUTIONARY RESULT OF MIXED INTEGER KNAPSACK SET PROBLEM

#### Heuristic (Obj Score: -3612.99096)

This novel algorithm enhances diversity in the solution search process by strategically selecting decision variables based on both their objective contributions and constraint interactions, while incorporating a degree of random exploration.

#### Code

```
import numpy as np
def select_neighborhood(n, m, k, site, value, constraint, initial_solution, current_solution,
    objective_coefficient):
    neighbor_score = np.zeros(n)
    contribution_scores = objective_coefficient * current_solution
    neighbor_score += contribution_scores
    for i in range(m):
        lhs_value = sum(value[i][j] * current_solution[site[i][j]] for j in range(k[i]))
        if lhs_value > constraint[i]:
            for j in range(k[i]):
                var_index = site[i][j]
                penalty = (lhs_value - constraint[i]) / max(1, np.sum(value[i]))
                neighbor_score[var_index] -= penalty * value[i][j] * np.random.uniform(0.8, 1.2)
    local_search_factor = (initial_solution - current_solution) ** 2
    neighbor_score += local_search_factor
    randomness = np.random.rand(n) * 0.1
    neighbor_score += randomness
    if np.max(neighbor_score) > 0:
        neighbor_score /= np.max(neighbor_score)
    return neighbor_score
```

#### Final Prompts

- (Cross) Design a hybrid heuristic algorithm that combines elements of genetic algorithms and simulated annealing to explore the solution space efficiently.
- (Cross) Develop a multi-phase heuristic optimization strategy that integrates particle swarm optimization with tabu search to dynamically adapt search parameters and enhance convergence rates.
- (Variation) Develop an algorithm that incorporates a novel optimization strategy, diverging from previous approaches, to enhance the objective function's outcome by exploring alternative parameter tuning techniques.
- (Variation) Please identify the main algorithm parameters and assist me in creating a new algorithm that has a different parameter settings of the score function provided.

## C CONVERGENCE ANALYSIS OF LLM-LNS

### C.1 EVOLUTIONARY PROGRESS IN COMBINATORIAL OPTIMISATION PROBLEM

Across both two combinatorial optimization problems Online Bin Packing and Traveling Salesman Problem, LLM-LNS consistently shows superior convergence and final solution quality compared to EOH.

In the Online Bin Packing problem shown in Figure 3, LLM-LNS shows better convergence behavior from the early stages. As the generations progress, LLM-LNS steadily improves and consistently outperforms EOH. The reduced variance in later generations highlights the stability of the LLM-LNS approach, which efficiently balances exploration and exploitation. Its dual-layer structure allows it to thoroughly explore the solution space, avoiding premature convergence and reaching a higher overall objective score. In contrast, EOH exhibits larger fluctuations and fails to achieve the same level of performance, indicating its limitations in maintaining robust progress during the evolutionary process.

In the Traveling Salesman Problem shown in Figure 4, although LLM-LNS starts with a less favorable initial population compared to EOH, it quickly demonstrates its advantage. Initially, EOH performs better, but it stagnates after the first 8 generations, showing little improvement afterward. Meanwhile, LLM-LNS continues to refine its solutions and steadily decreases the objective score. This indicates that the dual-layer structure of LLM-LNS effectively prevents it from getting trapped in local optima, maintaining a high level of exploration even in later generations. By the end of the evolutionary process, LLM-LNS surpasses EOH, achieving better overall results.

In both problems, LLM-LNS's ability to maintain diversity early in the process, combined with its strong convergence in later stages, gives it a clear advantage over EOH. The dual-layer evolutionary strategy ensures that LLM-LNS avoids stagnation, allowing for continuous improvement and ultimately leading to superior performance in solving combinatorial optimization problems.

1350  
 1351  
 1352  
 1353  
 1354  
 1355  
 1356  
 1357  
 1358  
 1359  
 1360  
 1361  
 1362  
 1363  
 1364  
 1365  
 1366  
 1367  
 1368  
 1369  
 1370  
 1371  
 1372  
 1373  
 1374  
 1375  
 1376  
 1377  
 1378  
 1379  
 1380  
 1381  
 1382  
 1383  
 1384  
 1385  
 1386  
 1387  
 1388  
 1389  
 1390  
 1391  
 1392  
 1393  
 1394  
 1395  
 1396  
 1397  
 1398  
 1399  
 1400  
 1401  
 1402  
 1403

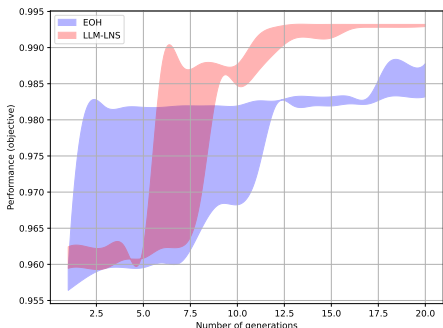


Figure 3: Evolutionary Progress of Heuristic Strategies in Online Bin Packing

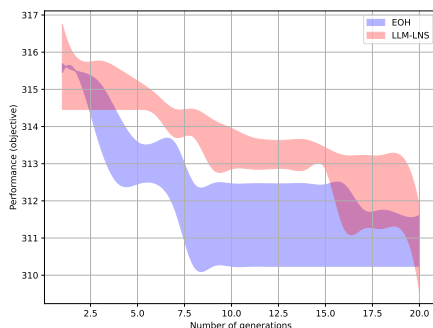


Figure 4: Evolutionary Progress of Heuristic Strategies in Traveling Salesman Problem




Figure 5: Convergence of Training and Testing Scores in 100-Generation of Online Bin Packing Problem.

### C.2 CONVERGENCE ANALYSIS OF GENERATIONS

In the Online Bin Packing Problem, we conducted 100 generations of iterative training using the proposed dual-layer strategy. Figure 5 shows the convergence trends for both the training and testing scores over these 100 generations. The results provide interesting insights into the behavior of our model during the evolutionary process, particularly in terms of how the training and testing losses evolve differently.

The training loss demonstrates a clear and consistent downward trend throughout the generations. Initially, the training score starts relatively high, but quickly drops within the first few generations. This rapid initial improvement indicates that the evolutionary algorithm is highly effective at optimizing the objective function within the training set. As the generations progress, the training score continues to decrease, eventually converging to a very low value. This steady decline suggests that the model is successfully adapting to the problem, continually refining its population and reducing the training objective. The absence of significant fluctuations in later generations implies that the model has reached a stable state, effectively minimizing the training loss with little variance.

On the other hand, the testing loss follows a somewhat different pattern. Initially, we observe a sharp decline in the testing score, which mirrors the behavior of the training score. However, after this initial drop, the testing score does not continue to improve as steadily as the training score. Instead, it stabilizes around a certain value and begins to exhibit small fluctuations. This behavior suggests that while the model is able to generalize to a degree, it encounters more variability in

Table 6: Comparison of objective values on large-scale MILP instances across different methods using SCIP as optimizer. For each instance, the best-performing objective value is highlighted in bold. The - symbol indicates that the method was unable to generate samples for any instance within 30,000 seconds, while \* indicates that the GNN&GBDT framework could not solve the MILP problem.

	SC <sub>1</sub>	SC <sub>2</sub>	MVC <sub>1</sub>	MVC <sub>2</sub>	MIS <sub>1</sub>	MIS <sub>2</sub>	MIKS <sub>1</sub>	MIKS <sub>2</sub>
Random-LNS	16164.2	171655.6	27049.6	277255.3	22892.9	222076.8	691.7	6870.1
ACP	17743.4	192791.2	27432.9	281862.4	23058.0	216008.8	29879.2	7913.5
CL-LNS	-	-	31285.0	-	15000.0	-	-	-
Gurobi	17934.5	320240.4	28151.3	283555.8	21789.0	216591.3	32960.0	329642.4
SCIP	25191.2	385708.4	31275.4	491042.9	18649.9	9104.3	29974.7	168289.9
GNN&GBDT	16728.8	261174.0	27107.9	271777.2	22795.7	227006.4	*	*
Light-MILPOPT	16147.2	166756.0	26956.8	269771.3	22963.6	230278.1	36125.5	<b>357483.8</b>
LLM-LNS(Ours)	<b>15950.2</b>	<b>161732.8</b>	<b>26763.4</b>	<b>268825.5</b>	<b>23137.19</b>	<b>230682.8</b>	<b>36147.7</b>	350468.7

the testing data compared to the training data. These fluctuations could be attributed to the inherent complexity or diversity of the unseen test instances, which the model has not been directly optimized for.

This phenomenon is reminiscent of the behavior observed during neural network training, where the training loss continues to decrease as the model becomes more specialized in fitting the training data, while the testing loss reaches a plateau and may exhibit some fluctuations. In this case, the testing loss reflects the model’s ability to generalize beyond the training set. The fact that the testing score does not continue to decrease beyond a certain point suggests that the model may have reached its limit in terms of generalization, possibly due to overfitting to the training data. However, the steady fluctuations in the testing score indicate that the model remains adaptable and does not suffer from severe overfitting, as there is no significant increase in the testing loss.

Overall, the divergence between the training and testing scores in later generations highlights the trade-off between optimization and generalization. While the dual-layer evolutionary strategy is highly effective at optimizing the training set, it must also balance the need for generalization to unseen data. The oscillation of the testing score around a stable value suggests that the model is reasonably robust but may benefit from additional techniques to further enhance its generalization performance, such as regularization or early stopping strategies in future iterations.

In summary, the convergence analysis of the 100-generation experiment reveals that while the training loss continues to decrease, the testing loss stabilizes with slight fluctuations. This behavior is indicative of a model that has successfully optimized for the training data while maintaining a reasonable level of generalization, akin to patterns observed in neural network training processes.

## D SUPPLEMENTARY EXPERIMENTS FOR LLM-LNS ON LARGE-SCALE MILP PROBLEMS

### D.1 PERFORMANCE OF LLM-LNS USING SCIP AS THE SUBSOLVER

In this supplementary set of experiments, we further evaluate the performance of LLM-LNS by incorporating SCIP as the subsolver for large-scale MILP problems. The results, summarized in Table 6, provide a comprehensive comparison across various methods using SCIP, offering deeper insights into the robustness and adaptability of LLM-LNS when faced with different solver strategies.

As seen in the results, LLM-LNS continues to demonstrate superior performance across most instances, consistently outperforming traditional LNS-based methods, learning-based frameworks such as GNN&GBDT, and even advanced solvers like Gurobi and SCIP. The highlighted bold values indicate that LLM-LNS achieves the best objective values in the majority of cases, reinforcing its scalability and effectiveness in large-scale MILP problems.

However, an interesting observation arises in the MIKS instances, where Light-MILPOpt outperforms LLM-LNS. This can be attributed to the unique challenges posed by MIKS in large-scale settings. Specifically, MIKS requires significantly more resources for neighborhood searches as the problem size increases, compared to smaller-scale instances. SCIP, as an optimizer, employs a different strategy for solving MIKS, which likely influences the performance of LLM-LNS when

Table 7: Comparison of standard deviation values on large-scale MILP instances across different methods using Gurobi as optimizer.

	SC <sub>1</sub>	SC <sub>2</sub>	MVC <sub>1</sub>	MVC <sub>2</sub>	MIS <sub>1</sub>	MIS <sub>2</sub>	MIKS <sub>1</sub>	MIKS <sub>2</sub>
Random-LNS	37.5	258.1	88.4	243.0	72.1	243.0	98.2	584.0
ACP	38.4	1039.3	71.6	403.5	60.3	928.8	118.2	649.2
CL-LNS	-	-	617.7	-	277.5	-	-	-
Gurobi	28.8	143.4	77.2	287.3	48.8	147.5	69.0	225.7
SCIP	13823.6	298211.7	107.3	262.0	57.5	85.8	73.2	242313.7
GNN&GBDT	360.1	3800.4	93.8	950.4	119.3	4738.8	*	*
Light-MILPOPT	1.0	145.7	79.4	209.4	52.1	133.1	41.7	272.5
LLM-LNS(Ours)	17.7	144.2	79.7	198.1	55.2	147.6	70.2	170.4

Table 8: Comparison of standard deviation values on large-scale MILP instances across different methods using SCIP as optimizer.

	SC <sub>1</sub>	SC <sub>2</sub>	MVC <sub>1</sub>	MVC <sub>2</sub>	MIS <sub>1</sub>	MIS <sub>2</sub>	MIKS <sub>1</sub>	MIKS <sub>2</sub>
Random-LNS	18.8	250.3	79.0	234.8	72.1	401.7	18.1	36.2
ACP	30.8	6338.3	77.2	217.6	60.3	946.4	1829.7	943.8
CL-LNS	-	-	617.7	-	277.5	-	-	-
Gurobi	28.8	143.4	77.2	287.3	48.8	147.5	69.0	225.7
SCIP	13823.6	298211.7	107.3	262.0	57.5	85.8	73.2	242313.7
GNN&GBDT	51.4	5587.6	91.4	474.0	80.0	660.4	*	*
Light-MILPOPT	37.7	693.4	77.3	216.9	51.6	151.7	80.0	1045.8
LLM-LNS(Ours)	20.4	169.5	82.6	188.7	54.3	75.9	68.7	1197.5

scaling to larger instances. In smaller-scale problems, LLM-LNS may have learned more aggressive strategies that are effective in those scenarios, but these strategies may lead to timeout issues in larger instances due to the increased computational complexity and extended iteration times required for SCIP. As a result, the overall improvement in performance is limited in these larger MIKS problems.

Despite these challenges, LLM-LNS still exhibits competitive performance in MIKS, managing to outperform many other methods, including Gurobi and traditional LNS strategies. The occasional time-out or reduced efficiency in MIKS does not overshadow the fact that LLM-LNS remains a robust and scalable solution across a wide range of large-scale MILP problems.

In conclusion, these supplementary experiments highlight the adaptability and robustness of LLM-LNS when using different subsolvers, including SCIP. Although challenges remain in specific problem instances like MIKS, LLM-LNS consistently delivers superior performance across most problem types, demonstrating its ability to generalize across solvers and problem scales. The results reinforce the notion that LLM-LNS effectively bridges the gap between traditional solvers and learning-based methods, offering a scalable solution for large-scale combinatorial optimization problems.

## D.2 COMPARISON OF STANDARD DEVIATION VALUES

The comparison of standard deviation (SD) values across different methods using both Gurobi and SCIP as sub-optimizers reveals several key insights into the stability of various approaches when solving large-scale MILP problems. Standard deviation reflects the consistency of the solutions; lower values indicate that the method is more stable and produces less variation in different runs.

As shown in Table 7, for the experiments using Gurobi, LLM-LNS consistently demonstrates low standard deviation values across most instances, indicating that it not only achieves superior objective values but does so with high stability. For example, in SC<sub>1</sub>, MVC<sub>2</sub>, and MIKS<sub>2</sub>, LLM-LNS has SD values of 17.7, 198.1, and 170.4, respectively, which are comparable to or lower than other methods. Light-MILPOpt also shows excellent stability in SC<sub>1</sub> and MIKS<sub>1</sub>, with SD values of 1.0 and 41.7, respectively, although its performance fluctuates more in other instances. In contrast, Random-LNS and ACP exhibit higher variability, especially in SC<sub>2</sub> and MIKS<sub>2</sub>, where ACP’s SD reaches as high as 1039.3 and 649.2, respectively, suggesting a lack of robustness in these instances. Gurobi itself also shows moderate consistency, while methods like CL-LNS fail to generate results for certain instances, indicating poor scalability for large problems.

As shown in Table 8, when SCIP is used as the optimizer, the trends remain somewhat similar. LLM-LNS continues to show stable performance, particularly in SC<sub>1</sub> and MVC<sub>2</sub>, with SD values of 20.4 and 188.7, respectively. However, SCIP itself exhibits extremely high variability in some instances,

Table 9: Comparison of error bar on large-scale MILP instances across different methods using Gurobi as optimizer.

	SC <sub>1</sub>	SC <sub>2</sub>	MVC <sub>1</sub>	MVC <sub>2</sub>	MIS <sub>1</sub>	MIS <sub>2</sub>	MIKS <sub>1</sub>	MIKS <sub>2</sub>
Random-LNS	65.4	318.3	142.1	350.8	104.4	333.6	158.9	808.8
ACP	56.8	1787.2	120.6	574.8	83.6	1233.0	173.7	742.7
CL-LNS	-	-	892.6	-	406.3	-	-	-
Gurobi	39.7	252.7	119.6	349.0	64.7	183.1	103.8	319.7
SCIP	25238.2	533457.2	165.2	402.1	96.9	103.6	94.6	433463.8
GNN&GBDT	511.3	5504.8	148.7	1522.6	160.1	7887.9	*	*
Light-MILPOPT	1.4	206.4	121.6	289.8	78.8	216.6	63.3	420.1
LLM-LNS(Ours)	27.9	187.9	125.4	289.8	82.2	199.3	111.7	259.2

Table 10: Comparison of error bar on large-scale MILP instances across different methods using SCIP as optimizer.

	SC <sub>1</sub>	SC <sub>2</sub>	MVC <sub>1</sub>	MVC <sub>2</sub>	MIS <sub>1</sub>	MIS <sub>2</sub>	MIKS <sub>1</sub>	MIKS <sub>2</sub>
Random-LNS	33.2	362.1	123.3	368.2	104.4	531.3	26.1	51.5
ACP	46.1	10845.3	106.0	324.1	83.6	1371.4	3253.2	1055.6
CL-LNS	-	-	892.6	-	406.3	-	-	-
Gurobi	39.7	252.7	119.6	349.0	64.7	183.1	103.8	319.7
SCIP	25238.2	533457.2	165.2	402.1	96.9	103.6	94.6	433463.8
GNN&GBDT	72.6	7349.2	147.2	678.8	100.4	1076.6	*	*
Light-MILPOPT	66.6	1223.3	118.5	305.6	79.1	239.4	124.2	1473.9
LLM-LNS(Ours)	31.7	231.2	131.9	266.7	68.9	94.7	105.9	1868.3

particularly in SC<sub>2</sub> and MIKS<sub>2</sub>, with SD values exceeding 298,000 and 242,000, respectively, which suggests that SCIP struggles with certain large-scale MILPs. This instability in SCIP could be due to its aggressive strategies or solver configurations being less suited to these specific problem instances. Light-MILPOpt again demonstrates relatively stable performance in most instances, although its SD increases significantly in some cases, such as MIKS<sub>2</sub>. GNN&GBDT and ACP also show considerable fluctuations, with ACP having an SD of 6338.3 in SC<sub>2</sub>, further highlighting its instability in large-scale settings.

In summary, LLM-LNS not only consistently outperforms other methods in terms of objective values but also maintains strong stability across a wide range of instances, particularly when compared to methods like Random-LNS, ACP, and SCIP. This robustness makes LLM-LNS a strong candidate for solving large-scale MILP problems effectively and consistently.


### D.3 COMPARISON OF ERROR BAR

The error bar comparison across different methods using Gurobi and SCIP as optimizers provides insights into the variability and confidence in solutions across large-scale MILP instances. Error bars quantify the uncertainty or inconsistency in the results, with smaller values indicating more reliable and consistent performance.

As shown in Table 9, for methods using Gurobi, LLM-LNS again demonstrates strong reliability with relatively small error bars across most instances. For example, in SC<sub>1</sub>, MVC<sub>2</sub>, and MIKS<sub>2</sub>, LLM-LNS has error bars of 27.9, 289.8, and 259.2, respectively. These values are noticeably smaller than those for methods like Random-LNS and ACP, which exhibit much larger error bars, reflecting greater instability. Light-MILPOpt also shows excellent performance with particularly low error bars in SC<sub>1</sub> (1.4) and MIKS<sub>1</sub> (63.3), but its error increases significantly in some other instances. Notably, SCIP exhibits extremely large error bars in several instances, such as SC<sub>2</sub> and MIKS<sub>2</sub>, where the error bars exceed 533,000 and 433,000, respectively, indicating significant inconsistency in its performance on these large-scale problems. GNN&GBDT also shows high error bars, suggesting that its performance is less reliable across different runs.


As shown in Table 10, when using SCIP as the optimizer, LLM-LNS continues to demonstrate relatively low error bars, particularly in SC<sub>1</sub>, MVC<sub>2</sub>, and MIKS<sub>1</sub>, where the values are 31.7, 266.7, and 105.9, respectively. These results are significantly more stable compared to methods like ACP and GNN&GBDT, which show very high error bars in instances like SC<sub>2</sub> (error bar of 10845.3 for ACP) and MIKS<sub>2</sub>. SCIP itself again shows extremely high error bars for instances such as SC<sub>2</sub> and MIKS<sub>2</sub>, further highlighting its instability in handling large-scale problems. Light-MILPOpt

1566  
1567  
1568  
1569  
1570  
1571  
1572




1573 Figure 6: Time-objective value graphs of medium-scale problems using Gurobi:  $SC_1$ ,  $MVC_1$ ,  $IS_1$ ,  
1574 and  $MKS_1$ .

1575  
1576  
1577  
1578  
1579  
1580  
1581  
1582




1583 Figure 7: Time-objective value graphs of large-scale problems using Gurobi:  $SC_2$ ,  $MVC_2$ ,  $IS_2$ , and  
1584  $MKS_2$ .

1585  
1586  
1587  
1588  
1589  
1590  
1591  
1592



1593 Figure 8: Time-objective value graphs of medium-scale problems using SCIP:  $SC_1$ ,  $MVC_1$ ,  $IS_1$ ,  
1594 and  $MKS_1$ .

1595  
1596  
1597  
1598  
1599  
1600  
1601  
1602



1603 Figure 9: Time-objective value graphs of large-scale problems using SCIP:  $SC_2$ ,  $MVC_2$ ,  $IS_2$ , and  
1604  $MKS_2$ .

1605  
1606

1607 performs well in some instances but also shows considerable variation in others, with error bars as  
1608 high as 1473.9 in  $MKS_2$ .

1609  
1610  
1611  
1612

1609 Overall, LLM-LNS consistently demonstrates lower error bars across both optimizers, Gurobi and  
1610 SCIP, indicating that it provides more reliable and consistent solutions for large-scale MILP prob-  
1611 lems. This makes it a strong candidate for scenarios where both solution quality and stability are  
1612 critical.

1613  
1614  
1615

#### 1614 D.4 CONVERGENCE ANALYSIS

1616  
1617  
1618  
1619

1616 In this section, we analyze the convergence performance of our proposed approach, our proposed  
1617 LLM-LNS, in comparison to several baseline methods for solving large-scale MILP problems, in-  
1618 cluding Random-LNS, ACP, Gurobi, GNN&GBDT, and Light-MILPOPT. The experimental results  
1619 are shown in Figures 6 through 9, which include instances of four different problem types: Set  
Covering (SC), Maximum Vertex Covering (MVC), Independent Set (IS), and Mixed Integer Knap-

sack Set (MIKS). We evaluate both medium-scale and large-scale instances using two solvers as sub-optimizer, Gurobi and SCIP.

The analysis of the convergence curves reveals several important observations:

- **Faster Initial Convergence:** For nearly all problem instances, the **LLM-LNS** approach demonstrates a significantly faster initial convergence compared to the baseline methods. The objective value drops sharply within the first few time steps, indicating that our method can quickly identify high-quality solutions. In contrast, methods like **Random-LNS** and **ACP** exhibit slower initial convergence, requiring more time to achieve similar reductions in the objective value.
- **Superior Final Objective Value:** Across both medium- and large-scale problem instances, our proposed LLM-LNS consistently achieves lower final objective values compared to the other methods. This is particularly evident in the large-scale instances (e.g.,  $SC_2$ ,  $MVC_2$ ,  $IS_2$ , and  $MIKS_2$ ), where the superiority of our method becomes more pronounced. While methods such as Random-LNS and ACP plateau early, often with suboptimal solutions, our proposed LLM-LNS continues to improve the solution even after other methods have stagnated.
- **Stable Convergence Behavior:** The convergence curves of our proposed LLM-LNS exhibit smooth and gradual decreases in the objective value, indicating stable optimization behavior. In contrast, some of the baseline methods, especially Random-LNS and GNN&GBDT, show more erratic convergence patterns, characterized by large and sudden jumps in the objective value. This suggests that our method is more robust and avoids the instability that can arise in heuristic-based search strategies.
- **Scalability:** The performance gap between our proposed LLM-LNS and the baseline methods becomes even more pronounced in large-scale problem instances. For example, in the large-scale  $MIKS_2$  and  $SC_2$  instances, our proposed LLM-LNS outperforms all other methods by a significant margin, converging to a much lower objective value within a shorter time frame. This demonstrates the scalability of our method, as it remains effective even as the problem size increases, whereas the performance of other methods, such as LightMILPOPT and ACP, degrades considerably.
- **Comparison with Exact Solvers:** When compared to the exact solver Gurobi, our proposed LLM-LNS shows comparable or even superior performance, particularly in terms of convergence speed. While Gurobi tends to find solutions that improve gradually over time, our proposed LLM-LNS reaches competitive solutions much faster, which is crucial in time-constrained scenarios. This highlights the practical advantage of our method in scenarios where computational resources or time are limited.

In summary, the experimental results demonstrate that **LLM-LNS** has clear advantages in terms of convergence speed, final solution quality, and robustness compared to both heuristic-based and exact optimization methods. Our approach is particularly well-suited for large-scale MILP problems, where it consistently outperforms the baseline methods by a significant margin.