

Learn to Relax with Large Language Models: Solving Constraint Optimization Problems via Bidirectional Coevolution

Beidan Liu^{1,4}, Zhengqiu Zhu^{1,4*}, Chen Gao², Tianle Pu⁴, Yong Zhao^{1,4}, Wei Qi³, Quanjun Yin^{1,4}

¹National Key Laboratory of Digital Intelligent Modeling and Simulation, National University of Defense Technology, Changsha, China

²BNRist, Tsinghua University, Beijing, China ³Department of Industrial Engineering, Tsinghua University, Beijing, China

⁴College of Systems Engineering, National University of Defense Technology, Changsha, China

Email: {liubangdan, zhuzhengqiu12, zhaoyong15, put122, yin_quanjun}@nudt.edu.cn,
chgao96@mails.tsinghua.edu.cn, qiw@tsinghua.edu.cn

Abstract

Large Language Model (LLM)-based optimization has recently shown promise for autonomous problem solving, yet most approaches still cast LLMs as passive constraint checkers rather than proactive strategy designers, limiting their effectiveness on complex Constraint Optimization Problems (COPs). To address this, we present AutoCO, an end-to-end **Automated Constraint Optimization** method that tightly couples operations-research principles of constraint relaxation with LLM reasoning. A core innovation is a unified triple-representation that binds relaxation strategies, algorithmic principles, and executable codes. This design enables the LLM to synthesize, justify, and instantiate relaxation strategies that are both principled and executable. To navigate fragmented solution spaces, AutoCO employs a bidirectional global-local coevolution mechanism, synergistically coupling Monte Carlo Tree Search (MCTS) for global relaxation-trajectory exploration with Evolutionary Algorithms (EAs) for local solution intensification. This continuous exchange of priors and feedback explicitly balances diversification and intensification, thus preventing premature convergence. Extensive experiments on three challenging COP benchmarks validate AutoCO's consistent effectiveness and superior performance, especially in hard regimes where current methods degrade. Results highlight AutoCO as a principled and effective path toward proactive, verifiable LLM-driven optimization.

1 Introduction

Constraint optimization problems (COPs) are extensively studied in logistics (Yao et al., 2024), finance, and industrial planning (Mokhtarzadeh et al., 2021). Most real-world COPs are NP-hard, with complexity arising from hard constraints that create

conflicting decisions and global dependencies (Papadimitriou and Steiglitz, 1982). These constraints fragment the feasible region, complicating the solution process and underscoring the necessity for efficient and high-quality solutions.

Many methods have been proposed for solving COPs, such as branch-and-bound, heuristics, etc. While these traditional approaches perform well in specific contexts, they often rely heavily on expert experience for effective adjustments to various problems. Recently, Large Language Model (LLM)-based methods have gained attention for their potential to provide end-to-end solutions (Liu et al., 2026). These methods leverage pre-trained knowledge and specific evolutionary frameworks to design algorithms tailored for COPs (e.g., Reevo(Ye et al., 2024), EoH(Liu et al., 2025)).

However, existing LLM methods primarily rely on feedback from feasible solutions to guide algorithm design toward optimality(Romera-Paredes et al., 2024). As hard constraints increase, obtaining feasible solutions becomes more difficult, resulting in ineffective feedback. This can lead to unguided search processes that produce suboptimal designs. Furthermore, these methods prioritize end-to-end generation but lack detailed problem analysis and strategy design, creating a disconnect from human experiences.

To address COPs with hard constraints, researchers often utilize relaxation strategies, starting with a relaxed version of the original problem (Storn, 1999). After identifying partial feasible solutions, they gradually tighten the constraints to find a complete solution. This relaxation approach mitigates the fragmentation of the feasible region caused by hard constraints, facilitating rapid identification of feasible solutions. The concept is widely used in operations research, such as in Large Neighborhood Search, where relaxing constraints enables the systematic exploration of diverse regions in the search space (Pisinger and Ropke, 2019).

* Corresponding author.

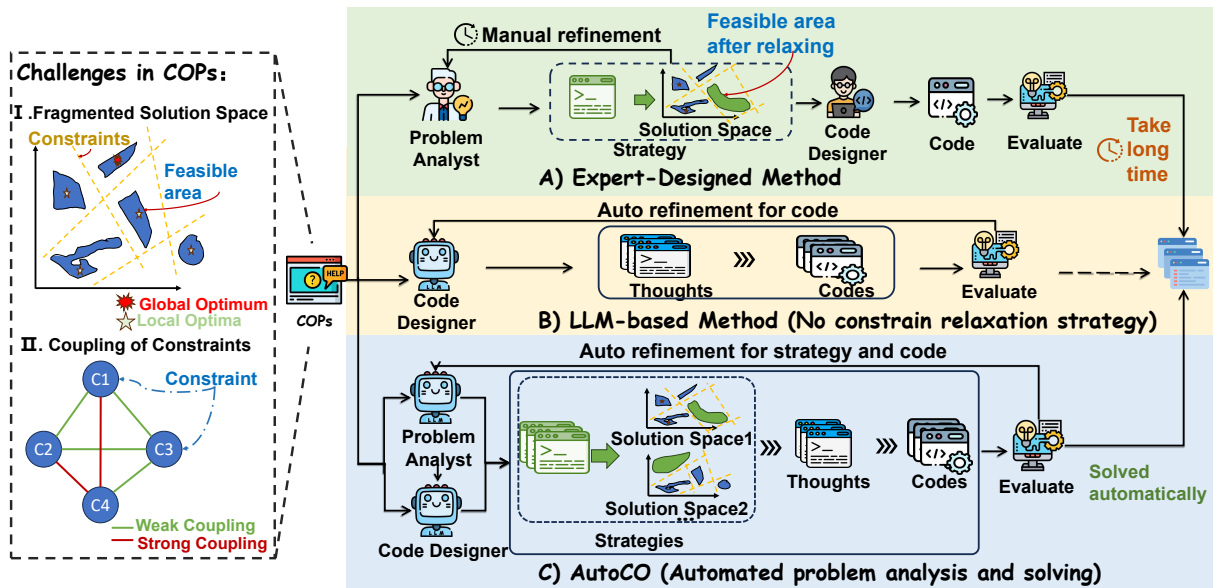


Figure 1: Comparisons of human/LLM-based solutions for COPs. A) Expert-designed method leverages human analysis to relax constraints for feasible solutions. B) Current LLM-based methods focus on code generation, lacking systematic problem analysis. C) Our AutoCO combines human-inspired relaxation strategies with automation to effectively discover feasible solutions.

Despite the advantages of this approach, existing relaxation strategies often depend on expert knowledge, making the formulation time-consuming. Integrating the world knowledge, reasoning, and strategy search capabilities of LLMs with these relaxation principles presents a promising opportunity, but introduces significant challenges: 1) *current LLM-based methods lack targeted mechanisms for exploring and relaxing constraints in constraint optimization landscapes, limiting their applicability to large-scale complex scenarios*; 2) *existing method focus on optimizing code and algorithmic concepts, neglecting the holistic representation of relaxation strategies, which hinders LLM’s ability to effectively optimize all three aspects together*; and 3) *the vast decision space makes it challenging to quickly identify the optimal relaxation strategy*.

To address these challenges, we introduce AutoCO (**A**utomated **C**onstraint **O**ptimization via LLM-driven bidirectional coevolution), an innovative solution that positions LLMs as proactive strategy designers for COPs (see Figure 1C). AutoCO enables LLMs to explicitly explore and optimize constraint relaxation strategies as integral components in the algorithm design process. The core innovation is a triple-representation scheme that maintains synchronized evolution across three coupled components: constraint relaxation strategies, algorithmic principles, and executable code. This

unified representation ensures that each generated solver explicitly employs its corresponding constraint relaxation strategy to solve the optimization problem. To efficiently coordinate this complex co-design process, we develop a bidirectional coevolution mechanism that combines local refinement via Evolutionary Algorithms (EA) with global exploration through Monte Carlo Tree Search (MCTS), enabling effective discovery of performant strategy-concepts-code triples. The contributions of this work are concluded as follows:

- We propose an end-to-end, three-phase solution (i.e., problem analysis, strategy search, and code execution) for solving COPs that empowers LLMs to master the complex art of constraint relaxation.
- We present a triple-representation scheme that maintains synchronized evolution of constraint strategies, algorithmic concepts, and executable code across abstraction levels. This unified encoding bridges the structural modeling gap in conventional approaches.
- We develop a bidirectional coevolution mechanism to address the vast strategy decision space. It delegates the fine-grained optimization of strategy-concepts-code triples to the local EA layer, while the global MCTS layer

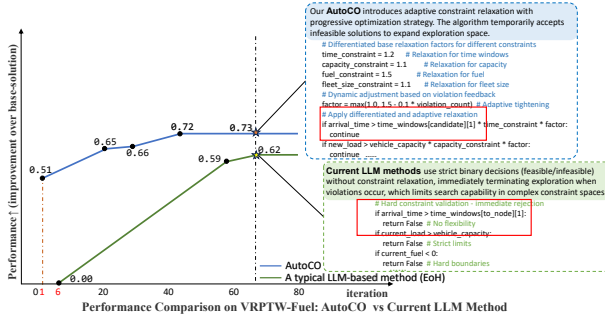


Figure 2: AutoCO (blue) vs. Current LLM method (green) performance on VRPTW-fuel problems over 100 iterations. The autonomous constraint relaxation temporarily expands feasible regions, enhancing optimization feedback.

explores the global strategy space to identify promising directions.

- We validate AutoCO’s effectiveness on three challenging COP benchmarks, showcasing its ability to tackle complex constraints, especially where Gurobi struggles, achieving a 24.7% average optimality gap reduction compared to the SOTA LLM-based methods.

2 Related Works

2.1 Traditional Methods for Constraint Relaxation

Various constraint relaxation methods have been developed to simplify complex operations research problems. Linear relaxation transforms integer programs into linear problems by relaxing integrality constraints (Clautiaux and Ljubić, 2025), while Lagrangian relaxation dualizes hard constraints into the objective via multipliers (Fisher, 1981). Recent works embed relaxation concepts directly into algorithm design: Kiessling temporarily relax constraint boundaries to boost performance (Kiessling et al., 2024), whereas Cao gradually relax constraints to improve search efficiency (Cao et al., 2017). However, these methods rely on expert knowledge, limiting their adaptability in different scenarios (Tseng and Bertsekas, 1987).

2.2 LLMs for Optimization Problems

LLMs are applied to optimization in diverse roles (Liu et al., 2026), serving as predictors for solution quality (Hao et al., 2024; Egami et al., 2023) or as feature extractors for domain knowledge (Wu et al., 2024). While demonstrating potential, these passive roles restrict reasoning and optimization.

Research thus explores LLMs as active optimizers for solution generation (Yang et al., 2024) and even as algorithm designers (Romera-Paredes et al., 2024; Liu et al., 2024; Ye et al., 2024), often integrated with evolutionary computation. Nevertheless, existing LLM methods lack effective problem analysis for COPs, hampering their performance under hard constraints.

2.3 Synergistic Mechanisms of Evolutionary Algorithms and Monte Carlo Tree Search

EA and MCTS are widely used for optimization: EA explores globally via population evolution (Eiben and Smith, 2015), while MCTS structures search through trees (Browne et al., 2012). Their combination is studied in multi-objective optimization (Hong et al., 2024), game action evolution (Baier and Cowling, 2018), and space decomposition (Xia et al., 2022). However, these hybrids remain largely domain-specific, focusing on narrow classes like games or continuous spaces (Bai et al., 2022), limiting universal applicability.

3 Preliminaries

3.1 Problem Definition of COPs

COPs can be formally modeled as follows:

$$\begin{aligned}
 (P) : \min_{\mathbf{x}} \quad & f(\mathbf{x}) \\
 \text{s.t.} \quad & g_k(\mathbf{x}) \leq 0, \quad k = 1, \dots, m \\
 & h_j(\mathbf{x}) = 0, \quad j = 1, \dots, p \\
 & \mathbf{x} \in \mathcal{X},
 \end{aligned} \tag{1}$$

where $f(\mathbf{x})$ is the objective function, $g_k(\mathbf{x})$ are inequality constraints, $h_j(\mathbf{x})$ are equality constraints, and \mathbf{x} is the vector of decision variables. The set \mathcal{X} defines the feasible region for \mathbf{x} , encompassing both continuous and discrete representations.

3.2 Constraint Relaxation Strategy

To enable explicit exploration and optimization of constraint relaxation strategies, we formalize them as a structured search space based on prior research (Grossmann, 2002; Floudas, 1995). A constraint relaxation strategy σ is defined as a set of pairs, each specifying a constraint and its corresponding relaxation factor:

$$\sigma = \{(g_1, \delta_1), \dots, (g_m, \delta_m)\}, \tag{2}$$

where g_k is the k -th constraint and $\delta_k \in [\alpha_k, \beta_k]$ its relaxation factor, with α_k, β_k defining the lower and upper relaxation bounds.

The strategy space \mathcal{T} is defined as the set of all possible relaxation strategies:

$$\mathcal{T} = \{\sigma_1, \sigma_2, \dots, \sigma_N\}. \quad (3)$$

Applying a relaxation strategy σ to the original problem yields a relaxed problem:

$$\begin{aligned} (\text{P}_\sigma) : \min_{\mathbf{x}} \quad & f(\mathbf{x}) \\ \text{s.t.} \quad & g_k(\mathbf{x}) \leq \delta_k, \quad k = 1, \dots, m \\ & h_j(\mathbf{x}) = 0, \quad j = 1, \dots, p \\ & \mathbf{x} \in \mathcal{X}. \end{aligned} \quad (4)$$

This transformation broadens the feasible region ($\mathcal{F}_{\text{P}_\sigma} \supseteq \mathcal{F}_{\text{P}}$) when $\delta_k \geq 0$, facilitating initial solution discovery. Crucially, different relaxation strategies $\sigma \in \mathcal{T}$ produce distinct expanded regions and search trajectories, potentially influencing initial solution efficiency and final solution quality in \mathcal{F}_{P} .

4 Methodology

4.1 Overview

AutoCO leverages constraint relaxation strategies designed and optimized by LLMs to exhibit human-like capabilities in addressing constraint optimization problems, thereby tackling the challenges posed by increasing hard constraints in constraint optimization. As illustrated in Figure 3, AutoCO comprises three stages: problem analysis and initial strategy design, optimal strategy search, and code execution.

In the first stage, our approach parses problem descriptions to generate an initial set of constraint relaxation strategies. Each strategy modifies the constraints of the COPs to temporarily enhance the probability of searching feasible solutions, thereby fostering a more extensive exploration of potential solutions. **The optimal strategy search stage** employs a dual-layer cooperative optimization mechanism: the local layer utilizes EA for fine-grained optimization, while the global layer leverages MCTS to explore the strategy space thoroughly, ensuring comprehensive search coverage. **The code execution stage** evaluates generated algorithms on problem instances to provide feedback for further optimization (Details about this stage are provided in Appendix A).

4.2 LLM-Driven Constraint Relaxation Strategy Generation

To address the limitations of expert reliance and to incorporate the insight of constraint relaxation, we

propose a three-step relaxation method that transforms expert-dependent constraint handling into an LLM-driven automated approach, by combining domain expertise with LLM reasoning capabilities.

Constraint Importance Analysis: Constraint relaxation strategies grow exponentially with the increasing number of problem constraints. We utilize LLMs to parse problem text to identify all constraints $\mathcal{G} = \{g_1, g_2, \dots, g_m\}$ and strategically assign importance weights $w_i \in [0, 1]$ to each constraint g_i (with $w_i = 1$ for critical constraints and $w_i = 0$ for negligible ones), enabling a shift from exhaustive enumeration to intelligent, context-aware constraint analysis.

Constraint Relaxation Range Suggestion: Based on constraint analysis, we determine adaptive relaxation factor ranges $[\alpha_i, \beta_i]$ ($\alpha_i \leq 1 \leq \beta_i$) for each constraint g_i using LLMs. This range defines permissible relaxation factors that can be applied to each constraint: 1 corresponds to the original constraint boundary, $\alpha_i < 1$ permits constraint tightening, and $\beta_i > 1$ allows controlled constraint violation.

Constraint Relaxation Strategy Generation: The final step synthesizes insights from the weights $W = \{w_1, \dots, w_m\}$ and ranges $R = \{[\alpha_1, \beta_1], \dots, [\alpha_m, \beta_m]\}$, the LLM generates an initial strategy set $\Sigma = \{\sigma_1, \dots, \sigma_k\}$, where each strategy $\sigma_j = (\delta_{j1}, \dots, \delta_{jm})$ is an m -dimensional vector representing relaxation coefficients. Crucially, this sampling process adheres to constraint criticality structures, ensuring that the generated strategies are derived from the systematic analysis of constraint importance and relaxation potential, rather than produced randomly.

4.3 Triple Representation Scheme

Existing LLM-based optimization methods (Liu et al., 2024) use single or dual representations (algorithm and code) but neglect to generate and optimize constraint relaxation strategies. As hard constraints increase, these methods fail to provide timely feedback, resulting in high computational costs and inefficiencies.

To address this limitation, we design a novel triple individual representation by incorporating constraint relaxation strategies. Our comprehensive triple representation scheme $\mathcal{I}_j = \langle \sigma_j, A_j, C_j \rangle$ explicitly integrates constraint relaxation strategy σ_j , algorithmic concept A_j , and executable code

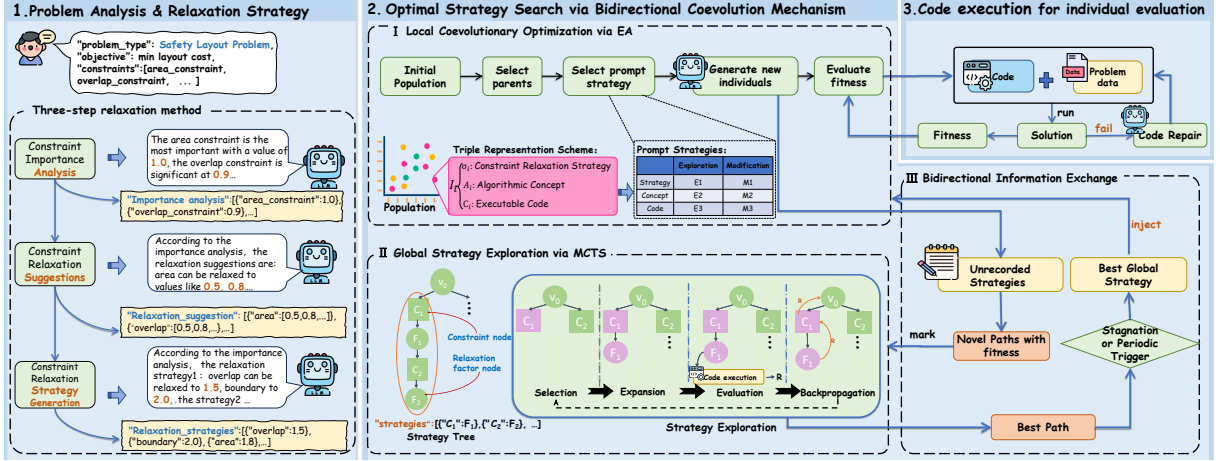


Figure 3: Architecture of AutoCO. Initially, we use LLMs to parse user-input problems and generate initial constraint relaxation strategies. Next, the bidirectional coevolution mechanism combining local EA and global MCTS explores and optimizes strategies and codes. Finally, we evaluate the generated algorithms on problem instances and provide individual fitness feedback.

C_j . This novel representation provides a systematic mechanism for strategy-code coevolution, transcending traditional boundaries between constraint handling, design, and implementation. By facilitating synchronized evolution across different abstraction levels, our approach enables rapid optimization and verification of relaxation strategies for COPs. An example scheme can be found in Appendix A.

4.4 Optimal Strategy Search via Bidirectional Coevolution Mechanism

To address complexities from hard constraints, particularly the rapid identification of effective strategies amid fragmented solution spaces and the extensive decision space of relaxation strategies, we devise a bidirectional coevolution mechanism combining a local EA with a global MCTS:

$$\mathcal{M}_{\text{exchange}} = \text{Bidirec_Update}(\mathcal{T}_{\text{MCTS}}, \mathcal{P}_{\text{EA}}), \quad (5)$$

where $\mathcal{M}_{\text{exchange}}$ represents the bidirectional exchange mechanism, $\mathcal{T}_{\text{MCTS}}$ is the MCTS search tree, and \mathcal{P}_{EA} is the EA population.

The mechanism leverages MCTS’s global perspective to help EA escape local optima, enabling effective navigation and optimization across diverse solution regions. Furthermore, it enhances resource efficiency by integrating EA’s exploration results into MCTS. It comprises three key components: (1) local coevolutionary optimization, (2) global strategy exploration, and (3) bidirectional information exchange.

Local Coevolutionary Optimization via EA

The EA performs local evolutionary refinement to

explore promising solutions. Parent selection uses tournament-based competition to balance selection pressure and diversity:

$$\mathcal{P}_{\text{EA}}^{t+1} = \text{Evolution}(\mathcal{P}_{\text{EA}}^t), \quad (6)$$

where $\mathcal{P}_{\text{EA}}^{t+1}$ represents the next-generation population. A key innovation is LLM-guided coordinated evolution, modifying constraint relaxation strategies, algorithmic concepts, and executable code while preserving coherence. Specialized prompts generate offspring that maintain logical consistency across the three representation levels. Details of these prompts are in Appendix B.

Global Strategy Exploration via MCTS

The MCTS component explores the constraint relaxation strategy space \mathcal{T} for global guidance. To improve the recording and searching of these strategies, we modify the structure to alternate between constraint nodes D and relaxation factor nodes R , where each complete path defines a strategy σ .

We introduce corresponding UCT calculation formulas for these node types to balance exploration and exploitation during the selection phase:

$$\text{UCT}(n) = \begin{cases} \frac{Q(n)}{N(n)} + k \sqrt{\frac{\ln N(F)}{N(n)}}, & n = D_i \\ \frac{Q(n)}{N(n)} + k \sqrt{\frac{\ln N(D_i)}{N(n)}} w_i, & n = R_{ij} \end{cases} \quad (7)$$

where Q is the cumulative rewards for nodes; N is visit counts; F is the parent node; k is the exploration constant balancing exploration and exploitation; D_i is the i -th constraint node; and R_{ij} is the j -th relaxation factor of the i -th constraint.

During the evaluation phase, we use LLM to generate executable code C_{strategy} for new strategies, with the fitness from code execution serving as the reward R for backpropagation to update node values. Upon search completion, MCTS outputs the optimal strategy σ^* to guide the local EA.

Bidirectional Information Exchange The bidirectional exchange mechanism integrates EA and MCTS through strategic information transfer, enhancing computational efficiency and aiding in overcoming local optima.

$$\mathcal{I}_{\text{exchange}} = \begin{cases} S_{\text{MCTS}} \rightarrow \mathcal{P}_{\text{EA}} \\ I_{\text{EA}} \rightarrow \mathcal{T}_{\text{MCTS}} \end{cases}, \quad (8)$$

where S_{MCTS} represents promising strategies from MCTS, I_{EA} denotes feasible solutions from EA.

(1) EA→MCTS: Continuous feedback allows each feasible solution produced by EA to update MCTS tree statistics with evaluated constraint-relaxation strategies. This integration diminishes the computational burden on MCTS by utilizing pre-assessed strategies, thus avoiding redundant exploration of known areas.

(2) MCTS→EA: The injection triggers under two conditions: upon detecting EA stagnation (via fitness-based monitoring) or periodically, to sustain global search momentum. The innovation lies in conveying MCTS-derived promising strategies, steering EA toward unexplored areas with high potential. MCTS offers a global perspective by pinpointing strategy combinations that EA’s local search may miss, fostering escape from local optima through informed population diversification.

5 Experiment

5.1 Experimental Settings

Problems and datasets To validate AutoCO’s generalizability, we select three COPs that demonstrate different constraint characteristics: Vehicle Routing Problem with Time Windows (VRPTW), VRPTW with fuel (VRPTW-Fuel), and Safety Facility Layout (SFL).

Dataset: VRPTW and VRPTW-fuel (Solomon benchmark: S(25)/M(50)/L(100) customers (Solomon, 1987)), SFL (MINLPLib: 4/8/5(Dual) with safety regions (Bonami et al., 2008)).

Baseline Methods Baselines include exact solver (Gurobi), reinforcement learning (DeepACO (Ye et al., 2023; Berto et al., 2025)), metaheuristics (SA (Kirkpatrick et al., 1983), GA (Whitley, 1994),

Table 1: Key Challenging Constraints in COPs

Problem	Constraint	Description
SFL	Non-overlap: $(x_i + w_i \leq x_j) \vee (x_j + w_j \leq x_i) \vee \dots$ Safety inclusion: $(x_i, y_i, w_i, h_i) \subseteq a$	Nonlinear geometric exclusion Coupling of continuous placement with discrete regions
VRPTW	Time windows: $t_j \geq t_i + s_i + d_{ij} - M(1 - y_{ij})$	Early decisions constrain later feasibility
VRPTW-Fuel	Fuel: $F_r = \sum d_{ij}(c_d + c_l L_i) \leq F_{\text{cap}}$	Resource accumulation coupled with load and distance

PSO (Kennedy and Eberhart, 1995), MA (Moscato, 1989), DE (Mohamed et al., 2019)), and LLM methods (FunSearch (Romera-Paredes et al., 2024), EoH (Liu et al., 2024), ReEvo (Ye et al., 2024)).

Evaluation Metrics Optimality gap $\gamma = \frac{|f_{\text{best}} - f_{\text{opt}}|}{|f_{\text{opt}}|}$, runtime (T_{e2e}), feasible solution time (T_{eff}), and stagnation time (T_{stag}).

Implementation Settings DeepSeek-R1, population 45, 2-hour limit, 100 runs on Intel i5-13400F/RTX 4060 Ti.

Details of problem models, datasets, baseline methods, and implementation settings are provided in Appendix C.

5.2 Comprehensive Performance Comparison

To validate AutoCO’s generalizability and effectiveness in solving COPs, we adopt the optimality gap as the primary evaluation metric, comparing AutoCO and baseline methods against state-of-the-art (SOTA) solutions as shown in Table 2.

- **Comparison with Traditional Methods:**

Our method demonstrates competitive performance across various optimization domains. Specifically, in the VRPTW domain, AutoCO achieves optimality gaps of 0.53 (S), 0.45 (M), and 0.42 (L) on average. While slightly behind traditional methods in the first two instances, AutoCO shows improvement in larger problem scales, even surpassing methods like MA (0.50) and SA (0.47). Regarding VRPTW-fuel problem, which introduces additional constraints related to fuel consumption, AutoCO maintains optimality gaps of 0.31 (S), 0.00 (M), and 0.00 (L) while other methods such as DeepACO struggles, showing gaps of 1.12 (S) and 0.97 (M). Notably, Gurobi fails to find feasible solutions for larger instances

Table 2: Performance comparisons with the SOTA baselines (reflected by optimality gap with SOTA,↓).

Type	Method	VRPTW			VRPTW-fuel			SFL		
		S	M	L	S	M	L	4	8	5(Dual)
Exact Solver	Gurobi	0.00	0.00	0.00	0.00	-	-	0.00	0.00	0.00
Reinforcement Learning	DeepACO	<u>0.29</u> (0.04)	0.46(0.01)	0.50(0.09)	1.12(0.08)	0.97(0.09)	0.38(0.14)	0.04(0.02)	0.19(0.03)	0.18(0.02)
Metaheuristic	MA	0.45(0.11)	0.49(0.10)	0.50(0.50)	0.46(0.08)	0.54(0.10)	-	<u>0.02</u> (0.03)	0.15(0.03)	0.16(0.02)
	DE	0.37(0.15)	0.40(0.12)	0.57(0.10)	0.46(0.09)	0.89(0.10)	0.60(0.06)	0.03(0.03)	0.14(0.02)	0.14(0.02)
	SA	0.37(0.14)	0.36(0.13)	0.47(0.14)	0.74(0.09)	0.84(0.07)	0.72(0.60)	0.05(0.03)	<u>0.08</u> (0.03)	<u>0.08</u> (0.03)
	GA	0.43(0.13)	0.44(0.06)	0.55(0.14)	0.91(0.04)	0.81(0.09)	0.33(0.15)	0.05(0.03)	0.16(0.03)	0.16(0.03)
	PSO	0.42(0.14)	0.44(0.08)	0.47(0.03)	0.48(0.04)	0.82(0.11)	<u>0.31</u> (0.07)	0.09(0.02)	0.11(0.02)	0.11(0.03)
LLM-based	FunSearch	0.50(0.22)	0.80(0.36)	0.70(0.26)	1.30(0.85)	0.45(0.09)	-	0.28(0.06)	0.27(0.03)	0.28(0.05)
	EoH	0.36(0.16)	0.63(0.34)	0.52(0.30)	0.81(0.46)	<u>0.20</u> (0.22)	-	0.18(0.05)	0.24(0.03)	0.24(0.07)
	ReEvo	0.36(0.18)	<u>0.33</u> (0.09)	0.50(0.27)	0.36(0.25)	<u>0.20</u> (0.26)	-	0.13(0.03)	0.25(0.03)	0.26(0.04)
Ours	AutoCO	0.53(0.17)	0.45(0.12)	<u>0.42</u> (0.15)	<u>0.31</u> (0.07)	0.00 (0.08)	0.00 (0.02)	<u>0.02</u> (0.04)	0.17(0.03)	0.15(0.02)

Note: The performance on three problems is reported as mean(std). **Bold** values represent the optimal solution in each column, while underlined values indicate the second-best solution. '-' denotes no feasible solution found within the limited time. A lower optimality gap indicates superior performance, reflecting the method's ability to generate high-quality solutions closer to the optimal benchmark.

within the designated time, highlighting AutoCO's robustness in handling increased complexity and its capability to generate solutions without extensive manual fine-tuning.

- **Comparison with LLM-based Methods:** Compared to existing LLM-based methods, AutoCO demonstrates substantial advantages. In the VRPTW-fuel domain, AutoCO achieves optimality gaps of 0.31 (S) and 0.00 (M), significantly outperforming ReEvo's gaps of 0.36 (S) and 0.20 (M) respectively. Other LLM methods like FunSearch show a dramatic decline in performance under increased constraints, with a gap of 1.30 (S). Similar tendencies can be observed in other two constrained optimization problems. Our results indicate an average reduction in the optimality gap across problem sets, confirming AutoCO's consistent generation of high-quality solutions as problem complexity escalates, thereby validating the effectiveness of our approach in solving challenging COPs.

Our comprehensive analysis reveals that AutoCO significantly advances LLM capabilities in solving COPs. *This performance demonstrates AutoCO's potential to autonomously design strategies, reducing manual engineering and advancing end-to-end automation in COPs. More importantly, AutoCO provides effective initial feasible solutions, assisting traditional methods or exact solvers to search more efficiently for optimal solutions, avoiding waste and stagnation in infeasible regions.*




Figure 4: Time distribution of AutoCO components.

5.3 Time Efficiency Analysis

We assess AutoCO's computational performance on VRPTW through three metrics: end-to-end runtime (T_{e2e}), time to first feasible solution (T_{tff}), and performance stagnation time (T_{stag}). As shown in Table 3, our temporal analysis reveals nuances in the computational characteristics of our approach compared to three baseline methods. Figure 4 provides a breakdown of AutoCO's internal time distribution across different stages.

- **Efficient Solution Generation and Stagnation Mitigation:** AutoCO demonstrates impressive efficiency in generating initial feasible solutions. As indicated in Table 3, AutoCO achieves T_{tff} values as low as 6.39 minutes for small instances and 11.65 minutes for large instances. This indicates that it generates feasible solutions significantly faster than baseline methods, with FunSearch ranging from 19.73 to 40.29 minutes, EoH ranging from 8.67 to 14.23 minutes, and ReEvo achieving 12.73 minutes for small instances and 12.54 minutes for large instances. Furthermore, AutoCO reduces T_{stag} to a minimum of

Table 3: Temporal performance characterization. (minutes)

Method	$T_{e2e}\downarrow$			$T_{tff}\downarrow$			$T_{stag}\downarrow$		
	S	M	L	S	M	L	S	M	L
FunSearch	53.03	49.54	137.93	40.29	31.73	19.73	8.71	10.27	28.49
EoH	35.35	32.11	94.64	14.21	8.67	14.23	14.79	21.33	19.53
ReEvo	30.21	29.93	86.97	12.73	7.02	12.54	9.27	19.98	24.46
AutoCO	65.07	104.11	108.87	6.39	9.31	11.65	25.53	5.99	11.77

5.99 minutes, compared to 8.71 to 28.49 minutes for FunSearch, 19.53 minutes for EoH, and 9.27 to 24.46 minutes for ReEvo, enhancing overall runtime efficiency.

- **Computational Time Distribution:** Figure 4 illustrates a comprehensive breakdown of AutoCO’s time allocation. The pie chart demonstrates that our computational effort is distributed efficiently among global exploration, solution generation, constraint verification, and evolutionary refinement, instead of being concentrated in any single stage. This balanced distribution underscores AutoCO’s holistic optimization strategy.

Although AutoCO’s end-to-end runtime appears longer (ranging from 65.07 to 108.87 minutes), it is crucial to note that *the structured time allocation allows for efficient feasible solution generation while managing low performance stagnation. This results in a more comprehensive exploration of the solution space. Compared to expert manual design, this runtime remains competitive and satisfactory.*

5.4 Ablation Study

We quantify the contributions of each component through the optimality gap increment (Δgap) relative to the complete AutoCO method. By selectively removing the constraint relaxation module (w/o σ), three-step generation strategy (w/o 3-steps), MCTS mechanism (w/o MCTS), and bidirectional information exchange (w/o Bidirectional), we systematically evaluate the individual contributions of each module in the AutoCO method. As shown in Table 4, the ablation experiments reveal the critical roles of these components: the constraint relaxation module causes an average performance degradation of +21.55%, indicating its crucial role in strategy formulation; the three-step generation strategy leads to a +14.10% performance decline, demonstrating the importance of structured generation; the MCTS mechanism shows a

+17.46% performance drop, highlighting the significance of global search in avoiding local optima; the bidirectional coevolution mechanism results in a +23.73% performance degradation, showcasing its substantial contribution to overall optimization. *The ablation experiments substantiate the meticulously designed modules and elucidate their synergistic optimization capabilities within AutoCO.*

Table 4: Ablation results of different variants. ($\Delta\text{gap}\%\downarrow$)

Variant	VRPTW	VRPTW-fuel	SFL	Avg.
Full AutoCO	0.0	0.0	0.0	0.0
w/o σ	+23.01	+24.98	+16.67	+21.55
w/o 3-steps	+14.75	+15.81	+11.74	+14.10
w/o MCTS	+15.42	+23.34	+13.62	+17.46
w/o Bidirectional	+25.07	+25.97	+20.15	+23.73

5.5 Relaxation Strategy Effectiveness Validation

To validate whether AutoCO’s autonomous relaxation strategies outperform expert-designed alternatives, we apply a base search method across different feasible regions on SFL-8 and SFL-5(Dual), examining three relaxation strategies: (1) Original Constraints (Unrelaxed), (2) AutoCO-Designed Strategy σ^* , and (3) Expert-Designed Strategy (details in Appendix C), and assess their success rates under varying computational budgets.

Figure 5 demonstrates performance across computational budgets. On SFL-8, original constraints plateau at 52%, while AutoCO strategy rapidly progresses from 0% to 80% success rate. AutoCO outperforms expert strategy, showing 1.1 times improvement at 2500 iterations. On SFL-5(Dual), which contains two safety regions, making the problem more challenging, performance gains are more modest, with AutoCO reaching 53% success. *AutoCO exhibits the most favorable ascending trend among all strategies, confirming that LLM-generated relaxation strategies effectively transform constraint navigation by constructing more accessible feasible regions than both rigid formulations and static expert designs.*

5.6 Optimization Dynamics Analysis of LLM-based methods

To systematically analyze the fundamental differences in optimization behaviors between traditional LLM-based methods and our AutoCO method, we conduct a comparative analysis across three SFL




Figure 5: Effectiveness of various constraint relaxation strategies under different computational budgets on SFL8 and SFL-5 (Dual).

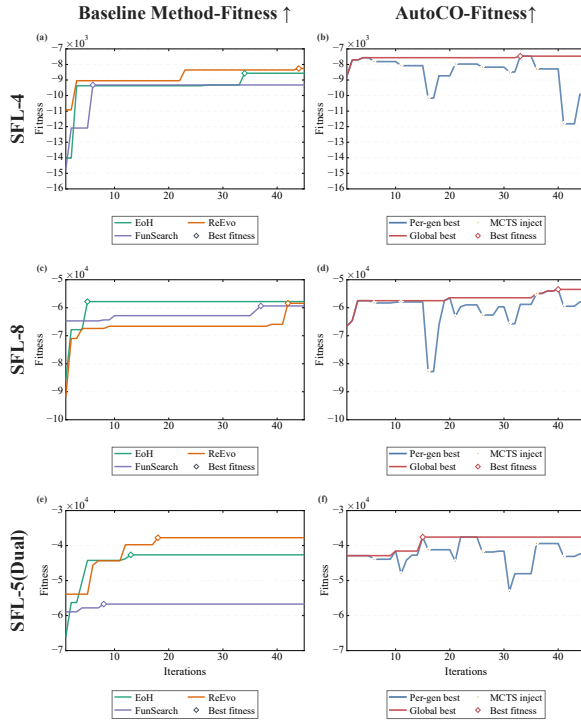


Figure 6: Convergence dynamics and fitness analysis of LLM-based methods across SFL datasets: current LLM-based methods (subplots a, c, e) and AutoCO (subplots b, d, f).

datasets using Figure 6, comparing baseline methods and AutoCO convergence dynamics.

Current LLM-based methods exhibit clear stagnation: After initial solutions, they focus exclusively on local optimization. In SFL-4, FunSearch converges at generation 5, while ReEvo and EoH attempt performance improvements by fine-tuning solutions between generations 35-40, remaining trapped in limited local search spaces.

AutoCO reveals distinctly different dynamics. Despite fitness stagnation, it maintains active search through bidirectional exchange and MCTS global strategies. As shown in subfigures b and d, AutoCO endures prolonged stagnation but continuously probes the solution space by injecting

global information, ultimately breaking through initial constraints and achieving significant fitness gains.

Empirical results clearly validate AutoCO’s key advantages: when traditional methods permanently stagnate after initial convergence, *AutoCO systematically escapes local optima by injecting global strategies from MCTS, temporarily accepting sub-optimal solutions to accumulate global information*. Although the experiment, constrained by time limitations, is set to 45 generations, the convergence plot demonstrates that AutoCO continues global exploration in later stages and achieves effective improvements, suggesting that with sufficient computational resources, AutoCO could potentially collect more comprehensive global information and achieve further performance enhancements.

6 Conclusion and Future Work

This paper introduces AutoCO, an end-to-end method for automated constraint optimization for COPs that transforms LLMs from passive validators into proactive strategy architects. Our key innovation lies in the systematic integration of three complementary components: (1) a structured three-step constraint relaxation method that codifies expert knowledge into LLM-interpretable principles, (2) a novel triple-representation scheme synchronizing constraint strategies, algorithmic concepts, and executable code across abstraction levels, and (3) a bidirectional coevolution mechanism that synergistically combines EA’s local refinement with MCTS’s global exploration to navigate solution spaces. Extensive experiments on VRPTW, SFL, and VRPTW-fuel verify AutoCO’s effectiveness and generalizability. While AutoCO represents a significant advance in autonomous optimization of COPs, future research could focus on: (1) extending multi-objective optimization with more sophisticated constraint coupling strategies, and (2) developing distributed coevolution frameworks for scalable optimization processes.

Limitations

This work primarily focuses on solving COPs through LLM-driven constraint relaxation and bidirectional coevolution. While the approach effectively handles complex constraints in domains like vehicle routing and facility layout, it overlooks several important aspects that warrant consideration.

First, the methodological scope, though validated on three challenging benchmarks, remains limited to static optimization problems with deterministic constraints, leaving stochastic and fully dynamic optimization environments unexplored. Additionally, as problem dimensionality increases, the strategy space grows exponentially, potentially affecting the scalability of the coevolution mechanism.

Second, while AutoCO demonstrates efficiency in generating initial feasible solutions, the computational overhead from maintaining multiple synchronized components, including LLM reasoning, evolutionary algorithms, and Monte Carlo tree search—results in longer end-to-end runtime compared to some baseline methods. This quality-efficiency trade-off may limit applicability in scenarios requiring real-time decision-making.

Furthermore, the method's performance is influenced by the reasoning and coding capabilities of LLMs. While the triple-representation scheme provides structured guidance for constraint relaxation and code generation, the effectiveness of the initial strategy design and evolutionary refinement may vary with the LLM's analytical consistency, particularly during the early exploration stages.

Finally, while the empirical results strongly support the effectiveness of the bidirectional coevolution mechanism, its theoretical convergence properties remain to be fully established.

Ethics Statement

The benchmarks and problem instances used in this work are all publicly available, ensuring no ethical concerns regarding data provenance. Beyond data sources, this work also considers the broader ethical implications of the AutoCO method.

First, regarding algorithmic robustness and reliability risks. As noted in the Limitations section, the method's performance is influenced by the reasoning capabilities of LLMs. Inconsistent analytical depth or programming proficiency in LLMs could affect the stability of constraint relaxation strategy generation, particularly during early explo-

ration phases. In safety-critical applications such as urban logistics or facility planning, suboptimal solutions generated due to such instabilities could lead to operational inefficiencies or resource misallocation. This risk necessitates maintaining human oversight when deploying such automated systems in real-world scenarios.

Second, regarding potential misuse and applicability boundaries. The method's capability to automatically design optimization strategies for complex constraint problems raises considerations about its appropriate application domains. While developed for beneficial purposes in logistics and planning, similar techniques could potentially be applied to contexts with significant societal implications, such as resource allocation in critical infrastructure or strategic planning in sensitive domains. Domain-specific risk assessments and ethical reviews are strongly recommended before deploying such automated optimization systems, particularly in high-stakes environments where suboptimal decisions could have substantial consequences.

Finally, regarding computational resource and environmental considerations. The bidirectional coevolution mechanism achieves strong performance at the cost of substantial computational resources due to its multiple synchronized components. To responsibly manage this demand, strict time limits are enforced during experiments, and algorithmic optimizations are implemented to maximize search efficiency within these constraints. While computationally intensive, this investment is justified by the framework's capability to automate complex optimization design tasks that traditionally require extensive human expertise.

Acknowledgments

This work is supported by the National Natural Science Foundation of China 72501291 and Youth Independent Innovation Fundation of NUDT ZK-2023-31.

References

- Fangyun Bai, Xinglong Ju, Shouyi Wang, Wenyong Zhou, and Feng Liu. 2022. [Wind farm layout optimization using adaptive evolutionary algorithm with Monte Carlo Tree Search reinforcement learning](#). *Energy Conversion and Management*, 252:115047.
- Hendrik Baier and Peter I. Cowling. 2018. [Evolutionary mcts for multi-action adversarial games](#). In *2018*

- IEEE Conference on Computational Intelligence and Games (CIG)*, pages 1–8.
- Federico Berto, Chuanbo Hua, Junyoung Park, Laurin Luttmann, Yining Ma, Fanchen Bu, Jiarui Wang, Haoran Ye, Minsu Kim, Sanghyeok Choi, Nayeli Gast Zepeda, André Hottung, Jianan Zhou, Jieyi Bi, Yu Hu, Fei Liu, Hyeonah Kim, Jiwoo Son, Haeyeon Kim, and 14 others. 2025. [RL4co: An extensive reinforcement learning for combinatorial optimization benchmark](#). In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2, KDD '25*, page 5278–5289. ACM.
- Pierre Bonami, Lorenz T. Biegler, Andrew R. Conn, Gérard Cornuéjols, Ignacio E. Grossmann, Carl D. Laird, Jon Lee, Andrea Lodi, François Margot, Nicolas Sawaya, and Andreas Wächter. 2008. [An algorithmic framework for convex mixed integer nonlinear programs](#). *Discrete Optimization*, 5(2):186–204.
- Cameron B. Browne, Edward Powley, Daniel Whitehouse, Simon M. Lucas, Peter I. Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez, Spyridon Samothrakis, and Simon Colton. 2012. [A survey of monte carlo tree search methods](#). *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1):1–43.
- Pan Cao, John Thompson, and H. Vincent Poor. 2017. [A sequential constraint relaxation algorithm for rank-one constrained problems](#). In *2017 25th European Signal Processing Conference (EUSIPCO)*, pages 1060–1064.
- François Clautiaux and Ivana Ljubić. 2025. [Last fifty years of integer linear programming: A focus on recent practical advances](#). *European Journal of Operational Research*, 324(3):707–731.
- Naoki Egami, Musashi Hinck, Brandon Stewart, and Hanying Wei. 2023. [Using imperfect surrogates for downstream inference: Design-based supervised learning for social science applications of large language models](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 68589–68601. Curran Associates, Inc.
- A. E. Eiben and J. E. Smith. 2015. *Introduction to Evolutionary Computing*, 2 edition. Natural Computing Series. Springer Berlin, Heidelberg. 55 b/w illustrations, 12 illustrations in colour.
- Marshall L. Fisher. 1981. [The lagrangian relaxation method for solving integer programming problems](#). *Management Science*, 27(1):1–18.
- Christodoulos A Floudas. 1995. *Nonlinear and Mixed-Integer Optimization: Fundamentals and Applications*. Oxford University Press.
- Ignacio E Grossmann. 2002. [Review of nonlinear mixed-integer and disjunctive programming techniques](#). *Optimization and Engineering*, 3:227–252.
- Hao Hao, Xiaoqun Zhang, and Aimin Zhou. 2024. [Large language models as surrogate models in evolutionary algorithms: A preliminary study](#). *Swarm and Evolutionary Computation*, 91:101741.
- Haokai Hong, Min Jiang, and Gary G. Yen. 2024. [Improving performance insensitivity of large-scale multiobjective optimization via monte carlo tree search](#). *IEEE Transactions on Cybernetics*, 54(3):1816–1827.
- J. Kennedy and R. Eberhart. 1995. [Particle swarm optimization](#). In *Proceedings of ICNN'95 - International Conference on Neural Networks*, volume 4, pages 1942–1948 vol.4.
- David Kiessling, Charlie Vanaret, Alejandro Astudillo, Wilm Decré, and Jan Swevers. 2024. [An almost feasible sequential linear programming algorithm](#). In *2024 European Control Conference (ECC)*, pages 2328–2335. IEEE.
- S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. 1983. [Optimization by simulated annealing](#). *Science*, 220(4598):671–680.
- Fei Liu, Xi Lin, Shunyu Yao, Zhenkun Wang, Xialiang Tong, Mingxuan Yuan, and Qingfu Zhang. 2025. [Large language model for multiobjective evolutionary optimization](#). In *International Conference on Evolutionary Multi-Criterion Optimization*, pages 178–191. Springer.
- Fei Liu, Xialiang Tong, Mingxuan Yuan, Xi Lin, Fu Luo, Zhenkun Wang, Zhichao Lu, and Qingfu Zhang. 2024. [Evolution of heuristics: Towards efficient automatic algorithm design using large language model](#). *Preprint*, arXiv:2401.02051.
- Fei Liu, Yiming Yao, Ping Guo, Zhiyuan Yang, Xi Lin, Zhe Zhao, Xialiang Tong, Kun Mao, Zhichao Lu, Zhenkun Wang, Mingxuan Yuan, and Qingfu Zhang. 2026. [A systematic survey on large language models for algorithm design](#). *ACM Computing Surveys*, 58(8):1–32.
- Ali W. Mohamed, Anas A. Hadi, and Kamal M. Jambi. 2019. [Novel mutation strategy for enhancing shade and lshade algorithms for global numerical optimization](#). *Swarm and Evolutionary Computation*, 50:100455.
- Mahdi Mokhtarzadeh, Reza Tavakkoli-Moghaddam, Chefi Triki, and Yaser Rahimi. 2021. [A hybrid of clustering and meta-heuristic algorithms to solve a p-mobile hub location-allocation problem with the depreciation cost of hub facilities](#). *Engineering Applications of Artificial Intelligence*, 98:104121.
- Pablo Moscato. 1989. On evolution, search, optimization, genetic algorithms and martial arts: Towards memetic algorithms. Technical Report C3P 826, Caltech Concurrent Computation Program 158-79, California Institute of Technology, Pasadena, CA.

- Christos H. Papadimitriou and Kenneth Steiglitz. 1982. *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.
- David Pisinger and Stefan Ropke. 2019. [Large neighborhood search](#). In Michel Gendreau and Jean-Yves Potvin, editors, *Handbook of Metaheuristics*, pages 99–127. Springer International Publishing, Cham.
- Bernardino Romera-Paredes, Mohammadamin Berekatain, Alexander Novikov, Matej Balog, M. Pawan Kumar, Emilien Dupont, Francisco J. R. Ruiz, Jordan S. Ellenberg, Pengming Wang, Omar Fawzi, Pushmeet Kohli, and Alhussein Fawzi. 2024. [Mathematical discoveries from program search with large language models](#). 625(7995):468–475.
- Marius M. Solomon. 1987. [Algorithms for the vehicle routing and scheduling problems with time window constraints](#). *Operations Research*, 35(2):254–265.
- Rainer Storn. 1999. [System design by constraint adaptation and differential evolution](#). *IEEE Transactions on Evolutionary Computation*, 3(1):22–34.
- Paul Tseng and Dimitri P. Bertsekas. 1987. [Relaxation methods for linear programs](#). *Mathematics of Operations Research*, 12(4):569–596.
- Nguyen Van Thieu and Seyedali Mirjalili. 2023. [Mealpy: An open-source library for latest metaheuristic algorithms in python](#). *Journal of Systems Architecture*, 139:102871.
- Darrell Whitley. 1994. [A genetic algorithm tutorial](#). *Statistics and Computing*, 4(2):65–85.
- Xingyu Wu, Yan Zhong, Jibin Wu, Bingbing Jiang, and Kay Chen Tan. 2024. [Large language model-enhanced algorithm selection: Towards comprehensive algorithm representation](#). In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence, IJCAI 2024*, pages 5235–5244.
- Hai Xia, Changhe Li, Sanyou Zeng, Qingshan Tan, Junchen Wang, and Shengxiang Yang. 2022. [Learning to search promising regions by a monte-carlo tree model](#). In *2022 IEEE Congress on Evolutionary Computation (CEC)*, pages 01–08.
- Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V Le, Denny Zhou, and Xinyun Chen. 2024. [Large language models as optimizers](#). In *The Twelfth International Conference on Learning Representations*.
- Canqi Yao, Shibo Chen, and Zaiyue Yang. 2024. [Cooperative operation of the fleet operator and incentive-aware customers in an on-demand delivery system: A bilevel approach](#). *IEEE Internet of Things Journal*, 11(6):9668–9680.
- Haoran Ye, Jiarui Wang, Zhiguang Cao, Federico Berto, Chuanbo Hua, Haeyeon Kim, Jinkyoo Park, and Guojie Song. 2024. [Reevo: Large language models as hyper-heuristics with reflective evolution](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 43571–43608. Curran Associates, Inc.
- Haoran Ye, Jiarui Wang, Zhiguang Cao, Helan Liang, and Yong Li. 2023. [Deepaco: Neural-enhanced ant systems for combinatorial optimization](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 43706–43728. Curran Associates, Inc.

A Details of The AutoCO Approach

In this part, we elaborate on the details of our proposed AotoCO (**A**utomated **C**onstraint **O**ptimization via LLM-driven bidirectional coevolution) as well as experimental results.

A.1 Triple Representation Scheme

In Constraint Optimization Problems (COPs), constraint handling is a critical and complex challenge. Existing LLM-based optimization methods (Liu et al., 2024) are often limited to dual representations of algorithms and code, struggling to effectively generate feasible solutions.

Multi-Drone Delivery Individual

<Constraint Relaxation Strategy >

```
1.drone_count_constraint: 5.0
2.battery_capacity_constraint: 2.0
3.payload_limit_constraint: 1.5
```

<Algorithmic Concept>

The algorithm design employs a sophisticated approach by introducing a dynamic weight scoring mechanism that comprehensively considers distance, profit, and constraint risk, with weights adaptively adjusted according to iteration progress. This is achieved through a constraint-aware probabilistic selection strategy that evaluates the risk of constraint violations to modulate selection probabilities, while simultaneously implementing an iterative evolution-based relaxation factor adjustment that dynamically modifies the degree of constraint relaxation based on current constraint violation status and solution progress.

<Executable Code>

```
def select_next_point(self, current_drone_id=None, available_points=None, current_position=None,
                    current_load=None, solution=None, violations=None, iteration=None, max_iterations=None):
    relaxation_factors = {
        'drone_count_constraint': self.get_relaxation_factor('drone_count_constraint', 5.0),
        'battery_constraint': self.get_relaxation_factor('battery_constraint', 2.0),
        'payload_constraint': self.get_relaxation_factor('payload_constraint', 1.5)
    }
    if iteration is not None and max_iterations is not None:
        progress = iteration / max_iterations
        distance_weight = 0.6 * (1 - 0.5 * progress)
        profit_weight = 0.3 + 0.4 * progress
        constraint_weight = 0.1 + 0.3 * progress
    else:
        distance_weight = 0.5
        profit_weight = 0.3
        constraint_weight = 0.2
    if iteration is not None and max_iterations is not None:
        progress = iteration / max_iterations
        for constraint in relaxation_factors:
            violation_degree = violations.get(constraint, 0) / (iteration + 1)
            adaptive_factor = 1.0 - (progress ** 0.7) * (1 + violation_degree)
            relaxation_factors[constraint] = max(1.0, relaxation_factors[constraint] * adaptive_factor)
    .....
```

Figure 7: The triple individual representation example for the multi-drone delivery problem.

To overcome this limitation, we propose an innovative triple individual representation $\mathcal{I}_j = \langle \sigma_j, A_j, C_j \rangle$, focusing on systematically integrating constraint relaxation strategies. This representation, through explicitly introducing the constraint relaxation strategy σ_j , and containing algorithmic concept A_j and executable code C_j , provides support for collaborative evolution in subsequent optimization stages.

Figure 7 illustrates a triple individual representation example for the multi-drone delivery problem. As shown, the individual I_j explicitly integrates three key components: 1) Constraint Relaxation Strategy (σ_j): Contains LLM-analyzed relaxation factors for critical constraints in multi-drone delivery (e.g., drone count: 5.0, battery: 2.0). 2) Algorithmic Concept (A_j): Describes the algorithmic implementation logic designed by LLM for constraint relaxation strategies, such as a dy-

Algorithm 1: Bidirectional Coevolution Mechanism

Input: Initial population $P_0 = \{\mathcal{I}_i = \langle \sigma_i, A_i, C_i \rangle\}$, problem instances Π , MCTS tree \mathcal{T}
Output: \mathcal{I}^*

```

1: while not converged do
2:   // Local Evolutionary Optimization
3:    $P_{\text{off}} \leftarrow \emptyset$ 
4:   for  $\mathcal{I}_p \in \text{Select}(P_t)$  do ▷ Roulette:
      $P(i) \propto f(\mathcal{I}_i)$ 
5:      $\mathcal{I}_c \leftarrow \text{LLM\_Generate}(\mathcal{I}_p)$  ▷  $\langle \sigma, A, C \rangle$ 
6:      $f(\mathcal{I}_c) \leftarrow \text{Eval}(C_c, \Pi)$  ▷ Code Execution
7:      $P_{\text{off}} \leftarrow P_{\text{off}} \cup \{\mathcal{I}_c\}$ 
8:   end for
9:    $P_t \leftarrow \text{Top}_N(P_t \cup P_{\text{off}})$  ▷ Elitism
10:  // Global MCTS Exploration
11:  for  $k = 1$  to  $K$  do
12:     $v \leftarrow \text{TreePolicy}(\mathcal{T})$  ▷ UCT Selection
13:     $R \leftarrow f(\text{LLM\_Generate}(v))$  ▷ Simulate  $\sigma$ 
14:     $\text{Backup}(v, R)$  ▷  $Q_v += R, N_v += 1$ 
15:  end for
16:   $\sigma^* \leftarrow \arg \max_{\sigma} Q_{\text{leaf}} / N_{\text{leaf}}$ 
17:  // Bidirectional Transfer
18:  for  $\mathcal{I} \in P_t$  do
19:     $\text{UpdateTree}(\mathcal{T}, \sigma_{\mathcal{I}}, f(\mathcal{I}))$  ▷ Local→Global
20:  end for
21:  if  $\text{Stagnant}(P_t)$  then ▷  $S_{\text{total}} < \theta$ 
22:     $P_t \leftarrow P_t \cup \langle \sigma^*, A, C \rangle$  ▷ Global→Local
23:  end if
24: end while
25: return  $\arg \max f(\mathcal{I})$ 

```

amic weight scoring mechanism balancing distance, profit, and constraint risk. 3) Executable Code (C_j): The executable Python code designed by the LLM that integrates the relaxation strategy with the algorithmic concept.

A.2 Bidirectional Coevolution Mechanism

We devise a bidirectional coevolution mechanism combining a local Evolutionary Algorithm (EA) with a global Monte Carlo Tree Search (MCTS) to address constraint challenges in COPs. This mechanism enables effective navigation and optimization across diverse solution regions by leveraging MCTS's global perspective to help EA escape local optima while enhancing resource efficiency

through integrating EA’s exploration results into MCTS. It comprises three key components: (1) local coevolutionary optimization, (2) global strategy exploration, and (3) bidirectional information exchange. The detail of the mechanism is shown in Algorithm 1.

Algorithm 2: Code Execution and Individual Evaluation

Require: Strategy $\sigma \in \Sigma$, Problem Instances Π
Ensure: Fitness Value $f(\sigma) \in R^+$

$\triangleright \Sigma$: Strategy Space, P : Problem Instance Set $\triangleright R_{\max}$: Maximum Repair Attempts, Typically $R_{\max} = 3$

- 1: $r \leftarrow 0$ \triangleright Repair attempt counter
- 2: $R_{\max} \leftarrow 3$
- 3: **while** $r \leq R_{\max}$ **do**
- 4: $\xi \leftarrow \mathcal{E}(\sigma, \Pi)$ \triangleright Execute strategy σ on instances Π
- 5: $\mu \leftarrow \omega(\xi)$ \triangleright Compute solution metrics
- 6: $\Phi \leftarrow \chi(\xi)$ \triangleright Check constraint violations
- 7: **if** $\neg\Phi$ **then** \triangleright No constraint violations
- 8: **return** $f(\sigma) \leftarrow \lambda(\mu)$ \triangleright Calculate fitness from metrics
- 9: **else**
- 10: $\sigma \leftarrow \gamma(\sigma, \tau)$ \triangleright Repair strategy using LLM
- 11: $r \leftarrow r + 1$
- 12: **end if**
- 13: **end while**
- 14: **return** f_{penalty} \triangleright Penalty for repeated failures

A.3 Code execution

To ensure secure, reliable execution and robust evaluation of LLM-generated code C , AutoCO implements a lightweight isolated execution environment with comprehensive safety mechanisms. Resource exhaustion is prevented by enforcing strict execution time limits ($T_{\max} = 60\text{s}$) and memory constraints ($M_{\max} = 1024\text{MB}$).

The fitness evaluation process can be formulated as:

$$f(\mathcal{I}) = \frac{1}{|\Pi|} \sum_{i=1}^{|\Pi|} f_i(\mathcal{I}), \quad (9)$$

where $f(\mathcal{I})$ is the average fitness of an individual, $|\Pi|$ represents the total number of problem instances, and higher objective function values indicate better solutions.

The detailed process of code execution and individual evaluation is illustrated in Algorithm 2. Upon detection of syntax or runtime errors, an LLM-based iterative repair mechanism is triggered (maximum attempts R_{\max}). **If no feasible solution is found or execution fails after reaching the repair limit, the individual is assigned an infinite fitness value.** The detail is shown in Algorithm 2.

B Prompt Design

B.1 LLM-Driven Constraint Relaxation Strategy Generation

In the process of constraint relaxation strategy generation for COPs, we systematically design three key prompt stages, aiming to fully utilize the reasoning capabilities of LLMs to transform constraint handling from experience-driven to intelligence-driven methods. Each prompt was carefully crafted to guide LLM to conduct in-depth and systematic constraint analysis from different perspectives.

B.1.1 Constraint Importance Analysis Prompt

In the first stage, we guide the LLM to assess the importance of constraints through a carefully designed prompt. Figure 8 demonstrates the specific prompt design.

The core objective of this prompt is to transform complex constraint parsing into a structured importance assessment, enabling the LLM to intelligently identify and weigh the criticality of various constraints from the problem description.

B.1.2 Constraint Relaxation Range Suggestion Prompt

Based on the constraint importance analysis in the first stage, the second stage prompt (as shown in Figure 9) aims to determine reasonable relaxation ranges for each constraint.

This prompt guides the LLM to intelligently recommend minimum and maximum relaxation factors based on the constraint importance and problem characteristics, laying the foundation for subsequent strategy generation.

B.1.3 Constraint Relaxation Strategy Generation Prompt

Finally, the prompt shown in Figure 10 integrates the analysis results from the previous two stages to generate a series of diverse and meaningful constraint relaxation strategies.

Through this systematic prompt design, we successfully transformed the constraint handling

strategies and problem information, redesigning constraint relaxation strategies while maintaining the unchanged algorithm concept and code structure, thereby achieving flexible adjustment and optimization of constraint strategies (as shown in Figure 11).

Prompt for Mutation2

You are an expert in solving optimization problems, here are the detailed information of an unsolved nonlinear combinatorial optimization problem.

<Problem Information>

Focus on code implementation innovation. Currently, parent individual is available with the following details:

Parent :{Constraint Relaxation Strategy 1};
 {Algorithmic Concept1};
 {Executable Code 1}

Please do mutation operation according to the following steps:

First, read the parent's algorithmic concept and executable code, and understand the current code design logic.
Second, read the current constraint relaxation strategy, and clarify the constraint relaxation strategy that needs to be implemented.
Then, design a new constraint relaxation strategy.
Finally, maintain the current algorithmic concept and relaxation strategy unchanged, design an executable code completely different from the current one, and output the new individual's three-part information according to the following format requirements:

```

1.Constraint Relaxation Strategy :      """json
                                       {constraint 1:0.9, constraint 2:1.3...}
                                       """
2. Algorithmic Concept:                 """json
                                       [The detail of Algorithmic Concept]
                                       """
3. Executable Code :                   """python
                                       [Python code]
                                       """

```

Do not output anything other than the specified requirements.

Figure 12: The prompt detail of mutation strategy 2

Prompt for Mutation3

You are an expert in solving optimization problems, here are the detailed information of an unsolved nonlinear combinatorial optimization problem.

<Problem Information>

Focus on algorithmic concept implementation innovation. Currently, parent individual is available with the following details:

Parent :{Constraint Relaxation Strategy 1};
 {Algorithmic Concept1};
 {Executable Code 1}

Please do mutation operation according to the following steps:

First, read the parent's algorithmic concept and executable code, and understand the current code design logic.
Second, read the current constraint relaxation strategy, clarify the constraint relaxation strategy to be implemented and the current implementation approach for constraint relaxation in the code.
Finally, maintain the current code framework and relaxation strategy unchanged, design an algorithmic concept completely different from the current approach to handling relaxation strategy, and output the new individual's three-part information according to the following format requirements:

```

1.Constraint Relaxation Strategy :      """json
                                       {constraint 1:0.9, constraint 2:1.3...}
                                       """
2. Algorithmic Concept:                 """json
                                       [The detail of Algorithmic Concept]
                                       """
3. Executable Code :                   """python
                                       [Python code]
                                       """

```

Do not output anything other than the specified requirements.

Figure 13: The prompt detail of mutation strategy 3

Mutation Strategy 2: Code Implementation Prompt By delving into the subtle nuances of current algorithm concepts and constraint relaxation strategies, the prompt is dedicated to guiding the LLM to discover novel implementation pathways, constructing fundamentally different executable code while preserving the original strategic framework (as shown in Figure 12).

Mutation Strategy 3: Algorithmic concepts Prompt Grounded in the existing code framework and constraint relaxation strategy, the prompt is committed to breaking through the current algorithmic concept, introducing a conceptually distinct approach, and guiding the LLM to reimagine the core mechanisms of problem-solving and strategy processing (as shown in Figure 13).

B.2.2 Crossover Strategy Prompts

Crossover strategies generate more promising new individuals by integrating advantages in constraint relaxation strategies, algorithmic concepts, or implementation code across different parent algorithms. We design three crossover strategy prompts to embody the coevolutionary optimization.

Crossover Strategy 1: Constraint Relaxation Strategies Prompt This crossover prompt focuses on integrating and optimizing constraint relaxation strategies across parent algorithms. By systematically combining and analyzing constraint relaxation factors, exploring their interactions, and adjusting executable codes based on innovative relaxation approaches, this prompt aims to generate offspring with enhanced constraint processing capabilities, with primary emphasis on refining relaxation mechanisms (as shown in Figure 14).

Prompt for Crossover1

You are an expert in solving optimization problems, here are the detailed information of an unsolved nonlinear combinatorial optimization problem.

<Problem Information>

Focusing on combining the advantages of existing parent individuals and optimizing constraint relaxation strategies. Currently, parent individuals are available with the following details:

Parent1 :{Constraint Relaxation Strategy 1};
 {Algorithmic Concept1};
 {Executable Code 1}

Parent2 :{Constraint Relaxation Strategy 2};
 {Algorithmic Concept2};
 {Executable Code 2}

Please do crossover operation according to the following steps:

First, read the constraint relaxation strategies of two parent individuals.
Second, combine the constraint relaxation strategies of both parents, focusing on their constraint relaxation factors.
Then, adjust parent 1's executable codes based on the new relaxation strategy and parent 1's algorithmic concept.
Finally, the new individual will be composed of the innovative algorithmic concept, the adjusted code and relaxation strategies from parent 1, and output the new individual's three-part information according to the following format requirements:

```

1.Constraint Relaxation Strategy :      """json
                                       {constraint 1:0.9, constraint 2:1.3...}
                                       """
2. Algorithmic Concept:                 """json
                                       [The detail of Algorithmic Concept]
                                       """
3. Executable Code :                   """python
                                       [Python code]
                                       """

```

Do not output anything other than the specified requirements.

Figure 14: The prompt detail of crossover strategy 1

Crossover Strategy 2: Code Implementation Prompt The prompt aims to achieve algorithmic implementation optimization through strategic code merging and fine-tuning. By extracting advantageous code segments from parent algorithms,

maintaining original relaxation strategies, and ensuring code execution alignment, this prompt enables the creation of more adaptive and sophisticated algorithmic implementations.(as shown in Figure 15)

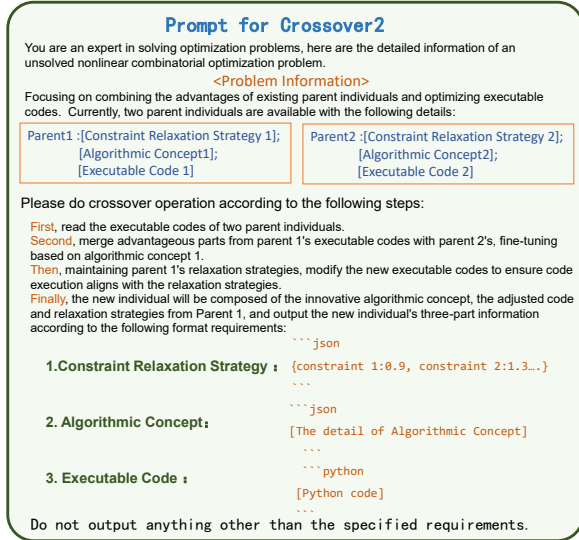


Figure 15: The prompt detail of crossover strategy 2

Crossover Strategy 3: Algorithmic concepts Prompt The prompt explores innovative integration of algorithmic core concepts while preserving the original code framework and relaxation strategies. By extracting advantageous conceptual elements from parent algorithms and carefully adjusting code details, this prompt facilitates the development of conceptually enriched algorithm variants that maintain the essential structural integrity of the original algorithms (as shown in Figure 16).

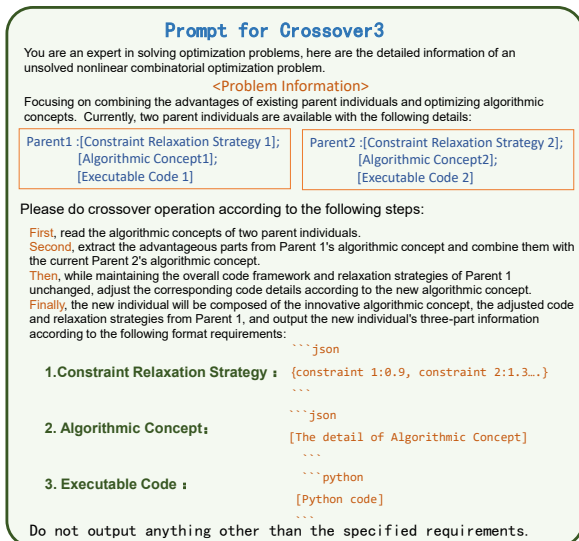


Figure 16: The prompt detail of crossover strategy 3

C Experimental Details

C.1 Benchmarks

To validate AutoCO's generalizability, we strategically select three COPs that systematically span the key challenge dimensions: Vehicle Routing Problem with Time Windows (VRPTW), Safety Facility Layout (SFL), and Vehicle Routing Problem with Time Windows and fuel (VRPTW-fuel).

Safety Facility Layout The SFL originates from the CMU-IBM Open Source MINLP Project: <https://egon.cheme.cmu.edu/ibm/page.htm>. The problem focuses on optimizing facility positioning within a given area, requiring simultaneous satisfaction of multiple complex spatial and safety constraints, and addressing significant interactions between constraints. Specifically, the Safety Facility Layout (SFL) problem minimizes the total cost ($C_{total} = \sum_{i < j} C_{ij} \cdot |c_i - c_j|_2 + \sum_i i C_i \cdot |c_i - c_{i0}|_2$), where ($\|\cdot\|_2$) represents Euclidean distance, subject to facility non-overlap restrictions (axis-aligned disjunction: $((x_i + w_i \leq x_j) \vee (x_j + w_j \leq x_i) \vee (y_i + h_i \leq y_j) \vee (y_j + h_j \leq y_i))$), positional range boundaries ($[\underline{x}_i, \bar{x}_i] \times [\underline{y}_i, \bar{y}_i]$), safety area limits ($|c_i - o_k|_2 \leq R_k$), and global boundary requirements. The problem introduces significant nonlinearity through Euclidean norms and circular safety inclusion constraints, creating complex, fragmented solution spaces, with test instances including SFL-4 (4 facilities, single safety zone), SFL-8 (8 facilities, single safety zone), and SFL-5 (5 facilities, dual safety zones), representing varying levels of geometric complexity and constraint interactions.

Vehicle Routing Problem with Time Windows

The VRPTW seeks a set of vehicle routes starting and ending at a single depot (node 0) to serve customers $N = \{1, \dots, n\}$ with travel distances d_{ij} , demands q_i , service times s_i , and time windows $[e_i, \ell_i]$. Each customer is visited exactly once; vehicle load may never exceed capacity Q ; early arrival allows waiting; and exceeding ℓ_i produces lateness. A solution is encoded as a giant-tour permutation decoded into a route family \mathcal{R} . For a route $r \in \mathcal{R}$ with sequence $(v_0 = 0, v_1, \dots, v_m, v_{m+1} = 0)$, its distance is $D_r = \sum_{k=0}^m d_{v_k v_{k+1}}$, $D = \sum_{r \in \mathcal{R}} D_r$. Arrival times propagate (with $t_{v_0} = 0, s_0 = 0$) by $t_{v_{k+1}} = \max(t_{v_k} + d_{v_k v_{k+1}}, e_{v_{k+1}}) + s_{v_{k+1}}$, $k = 0, \dots, m$. Lateness for customer i is $\max(0, t_i - \ell_i)$. A penalized objective consistent with our implementation is $C_{VRPTW} = D +$

$\alpha_{\text{time}} \sum_{i \in N} \max(0, t_i - \ell_i) + \alpha_{\text{fleet}} \max(0, |\mathcal{R}| - K_{\text{max}})$, where K_{max} bounds the fleet size. Constraints: (i) each $i \in N$ appears in exactly one route; (ii) every route starts/ends at the depot; (iii) cumulative load on any prefix does not exceed Q ; (iv) time-window upper bounds ℓ_i respected (or penalized if violated); (v) optional fleet-cardinality limit $|\mathcal{R}| \leq K_{\text{max}}$. Classical Solomon instances (Solomon, 1987) are challenging due to the interdependence of temporal feasibility and capacity-driven splitting.

Vehicle Routing Problem with Time Windows and fuel VRPTW-Fuel augments VRPTW by adding a per-route fuel capacity constraint. Retaining base notation, each arc (v_k, v_{k+1}) on a route carries load-dependent fuel: $f_{v_k v_{k+1}} = (\beta + \gamma L_{v_k}) d_{v_k v_{k+1}}$, with L_{v_k} denoting remaining load just before traversing the arc. Route fuel $F_r = \sum_{k=0}^m f_{v_k v_{k+1}}$ must satisfy $F_r \leq C_{\text{fuel}}$. Violations aggregate as $V_{\text{fuel}} = \sum_{r \in \mathcal{R}} \max(0, F_r - C_{\text{fuel}})$, yielding the penalized objective $C_{\text{VRPTW-Fuel}} = D + \alpha_{\text{time}} \sum_{i \in N} \max(0, t_i - \ell_i) + \alpha_{\text{fleet}} \max(0, |\mathcal{R}| - K_{\text{max}}) + \alpha_{\text{fuel}} V_{\text{fuel}}$. All classical constraints remain, with the added per-route fuel capacity. The affine load term $(\beta + \gamma L_{v_k})$ couples sequence, load depletion, and arc selection: front-loading high-demand customers reduces later fuel intensities but can degrade temporal feasibility or distance, sharpening trade-offs and fragmenting the search landscape. Ties among fully feasible solutions are further broken (implementation) by lower fleet size, then lower distance, then lower maximum single-route fuel consumption.

C.2 Experiment Implementation

All experiments use DeepSeek-R1 LLM (population size: 45, time limit: 2 hours). 100 independent runs per instance on Intel® Core™ i5-13400F/NVIDIA RTX 4060 Ti platform. Non-LLM baselines run as single-threaded CPU executables.

C.3 Baseline and Baseline Setting

To comprehensively evaluate AutoCO’s performance, we selected diverse baseline methods spanning exact solvers, Reinforcement Learning, metaheuristics, and LLM-based approaches.

C.3.1 Exact Solver

Gurobi As a leading commercial mixed-integer programming solver, Gurobi (see <https://www.gurobi.com/>)

utilizes advanced mathematical programming techniques, such as Branch and Bound, Cut Generation, and preprocessing techniques. Its ability to precisely solve complex optimization problems and numerical stability make it the gold standard for exact solutions.

C.3.2 Reinforcement Learning

DeepACO Introducing a neural-enhanced metaheuristic framework that generalizes and amplifies traditional Ant Colony Optimization (ACO) through deep reinforcement learning (Ye et al., 2023). This approach employs a graph neural network to encode combinatorial problem instances, generating advanced heuristic priors that significantly accelerate the solution construction process. The framework is trained via proximal policy optimization, enabling adaptive pheromone initialization and dynamic parameter adjustment, which collectively enhance search efficiency and solution quality across diverse combinatorial problems.

C.3.3 Metaheuristic Algorithms

Simulated Annealing (SA) Inspired by the physical annealing process, SA introduces a temperature parameter to control the probability of accepting suboptimal solutions (Kirkpatrick et al., 1983). At high temperatures, the algorithm is more inclined to accept suboptimal solutions, helping to escape local optima; as the temperature decreases, the algorithm gradually converges to the optimal solution.

Genetic Algorithm (GA) Mimicking natural selection and evolutionary processes, GA optimizes solutions through population-level selection, crossover, and mutation operations (Whitley, 1994). Each generation selects the most optimal individuals through fitness evaluation, simulating the natural selection mechanism of biological evolution.

Particle Swarm Optimization (PSO) Inspired by collective behaviors of bird flocks and fish schools, each "particle" in PSO represents a potential solution, exploring the solution space by mutual following and learning (Kennedy and Eberhart, 1995). Particles are influenced not only by their own historical optimal solutions but also guided by the global optimal solution.

Memetic Algorithm (MA) Combining genetic algorithms’ evolutionary mechanisms with local search intensification, MA implements a hybrid

optimization paradigm (Moscato, 1989). This algorithm maintains population diversity through global evolutionary operations while employing problem-specific local search techniques to refine individual solutions. This dual-optimization strategy achieves an effective balance between exploration and exploitation, frequently demonstrating superior performance in complex combinatorial optimization landscapes.

Differential Evolution (DE) Operating through vector population mutation and recombination, DE implements a robust evolutionary strategy for continuous optimization domains (Mohamed et al., 2019). This approach generates new candidate solutions by applying weighted differential mutations between population vectors, followed by crossover operations that recombine parental and trial components. Its mutation mechanism provides self-organizing coordinate system rotation and translation invariance, ensuring robust performance across diverse optimization scenarios.

C.3.4 Large Language Model-based Methods

FunSearch A sophisticated optimization methodology leveraging LLMs and a novel island-based evolutionary search paradigm (Romera-Paredes et al., 2024). By decomposing the search process into multiple autonomous computational domains exploring solution spaces in parallel, this approach facilitates robust algorithmic discovery. Each computational island generates candidate solutions through strategic generation and evaluation mechanisms, with intermittent knowledge transfer preventing local optimization convergence.

EOH An advanced computational framework for automated heuristic generation integrating LLMs with evolutionary computation (Liu et al., 2024). The methodology concurrently evolves algorithmic conceptualization and implementation, emulating expert-level heuristic design strategies. Employing a quintuple of sophisticated prompting strategies across exploration and modification taxonomies, it systematically synthesizes diverse and computationally efficient heuristic algorithms.

ReEvo A pioneering hyper-heuristic optimization framework utilizing LLMs-driven evolutionary search with introspective refinement mechanisms, ReEvo (Ye et al., 2024). By synthesizing evolutionary computation principles with the reflective capabilities of generative models, this approach en-

ables progressive algorithmic enhancement through multi-dimensional strategic and tactical refinement processes.

C.4 Configuration of MetaHeuristic Baselines

All meta-heuristic baselines (GA, MA, DE, SA, PSO) are implemented using the **mealpy (3.0.3)** library to ensure a standardized and reproducible comparison (Van Thieu and Mirjalili, 2023). A unified experimental protocol is strictly followed, which includes an identical termination budget (epoch=80, population=45, run=100).

To guarantee the statistical robustness and determinism of our results, we employ a deterministic set of 20 seeds defined *a priori* by an arithmetic progression: $S = \{s_0 + \Delta \cdot i \mid i = 0, \dots, 19\}$ with $s_0 = 17$ and $\Delta = 73$. This rule-based construction prevents cherry-picking while preserving broad coverage of pseudorandom initializations. All (algorithm, instance) pairs use every seed in S ; no seed is added, removed, or replaced after observing any results.

To ensure fairness and mitigate parameter sensitivity, a single set of hyperparameters is fixed for each algorithm and applied across all problems, which avoids iterative per-instance tuning. Key deviations from the library’s default parameters, determined through a preliminary calibration, are summarized in Table 5.

Table 5: Core hyperparameters (Defaults vs. Adopted).

Algorithm	Parameter	Default	Adopted
GA	p_c (crossover rate)	0.95	0.85
	p_m (mutation rate)	0.025	0.05
PSO	c_1 (cognitive)	2.05	1.8
	c_2 (social)	2.05	1.8
	w (inertia)	0.4	0.6
DE	w_f (mutation factor)	0.1	0.6
	c_r (crossover rate)	0.9	0.9
	strategy	0	5
MA	p_c (crossover rate)	0.85	0.85
	p_m (mutation rate)	0.15	0.15
	p_{local} (local search probability)	0.5	0.5
	max_local_gens (local search depth)	10	10
SA	T_0 (initial temperature)	100	500
	$step_size$ (perturbation step size)	0.1	0.2

This standardized setup, with its deterministic seed strategy and fixed hyper-parameters, ensures that performance differences are attributable to the core search strategies rather than implementation or tuning biases.

C.5 Relaxation Strategy Effectiveness Validation

The solution space in COPs often exhibits highly fragmented characteristics, with the feasible domain partitioned into isolated connected components. This experiment defines a foundational search algorithm to explore the solution space under different relaxation strategies, aiming to validate the topological structure reconstruction effects of various relaxation approaches. Specifically, we evaluate the performance differences between AutoCO design, expert-defined, and non-relaxation baseline strategies in alleviating solution space fragmentation.

C.5.1 Base Random Walk Search Algorithm

To eliminate the inherent search algorithm’s influence on solution space exploration, we design a heuristic-free baseline, namely random walk search algorithm (the details are shown in Algorithm 3).

Table 6: Constraint relaxation strategies for SFL8.

Strategy	Overlap	Position Range	Safety Area	Global Boundary
AutoCO-Designed	1.5	1.8	1.2	1.3
Expert-Designed	2.0	1.1	1.3	2.2
No Relaxed	1.0	1.0	1.0	1.0

Table 7: Constraint relaxation strategies for SFL5 (Dual).

Strategy	Overlap	Position Range	Safety Area	Global Boundary
AutoCO-Designed	2.0	2.6	1.5	2.2
Expert-Designed	1.8	2.0	1.4	2.5
No Relaxed	1.0	1.0	1.0	1.0

C.5.2 Experimental Configuration

The experimental configuration is structured to ensure comprehensive and reliable results: 1) We generate a uniformly random pre-generated initial solution set to maintain baseline consistency across all strategies; 2) By using an identical random seed for solution generation, we ensure equivalent and controlled initial conditions for each comparative approach; 3) We consistently evaluate feasibility using original strict constraints, implementing a 30-second search limit per iteration to standardize computational resources; 4) To establish statistical robustness, we conduct 1000 trials for each configuration, systematically exploring computational budgets from 500 to 2500 steps, which enables a rigorous and multifaceted assessment of the proposed relaxation strategies’ performance and reliability.

C.5.3 Relaxation Strategy Configurations

We evaluate relaxation strategies across two datasets: SFL8 and SFL5 (dual), comparing AutoCO, expert-designed, and safe region relaxation approaches. Tables 6 and 7 illustrate the detailed relaxation strategies for each dataset.

Algorithm 3: Base Random Walk Search Algorithm

Require: Initial solution $S_0 \in S$, Relaxation parameter $\Theta \in R^+$, Exploration steps $T \in N$

Ensure: Feasible solution discovery count η , Optimal solution $S^* \in S$, Optimal cost $C^* \in R^+$

▷ S : Solution Space, Θ : Constraint Relaxation Threshold

- 1: Initialize: $S \leftarrow S_0, \eta \leftarrow 0, S^* \leftarrow \emptyset, C^* \leftarrow +\infty$
- 2: **for** $t = 1$ **to** T **do** ▷ Check current solution feasibility under strict constraints
 - 3: **if** $\phi_{\text{strict}}(S)$ **then**
 - 4: $\eta \leftarrow \eta + 1$
 - 5: $C \leftarrow \omega(S)$ ▷ Calculate solution cost
 - 6: **if** $C < C^*$ **then**
 - 7: $S^* \leftarrow S$
 - 8: $C^* \leftarrow C$
 - 9: **end if**
 - 10: **end if**
 - ▷ Generate candidate move via atomic exploration
 - 11: $S' \leftarrow \xi(S)$ ▷ ξ : Atomic move function
 - ▷ Accept move based on relaxed constraints
 - 12: **if** $\phi_{\text{relaxed}}(S', \Theta)$ **then**
 - 13: $S \leftarrow S'$
 - 14: **end if**
 - 15: **end for**
 - 16: **return** η, S^*, C^*
 - ▷ Atomic Move Function $\xi(S)$
 - 17: **function** $\xi(S)$
 - 18: $r \leftarrow \text{RandomRectangle}()$ ▷ Randomly select rectangle
 - 19: $\delta \leftarrow \text{RandomDirection}()$ ▷ Random move direction
 - 20: $S_{\text{new}} \leftarrow \text{Copy}(S)$
 - 21: $S_{\text{new}}[r].x \leftarrow S[r].x + \delta.dx$
 - 22: $S_{\text{new}}[r].y \leftarrow S[r].y + \delta.dy$
 - 23: **return** S_{new}
 - 24: **end function**
