

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 MINIOPT: REASONING TO MODEL AND SOLVE GENERAL OPTIMIZATION PROBLEMS WITH LIMITED RESOURCES

006 **Anonymous authors**

007 Paper under double-blind review

010 ABSTRACT

013 Modeling and solving optimization problems via large language models (LLMs)
014 has attracted increasing attention recently. Although both prompt-based and
015 learning-based methods have achieved progress, they remain limited by their re-
016 liance on large data volumes, high-quality annotations, expensive intermediate
017 step verification, and huge computational overhead. From a data privacy perspec-
018 tive, a low-cost localized deployment of small-scale LLMs is of significant value.
019 To train a small-scale LLM with excellent optimization generalization under lim-
020 ited resources, this paper proposes a reasoning to model and solve paradigm called
021 MiniOpt based on reinforcement learning (RL) with verifiable reward. To reduce
022 the demand for training data, MiniOpt adopts two-stage RL training. In the first
023 stage the model quickly learns the model-and-solve paradigm and in the second
024 stage it acquires strong optimization generalization ability. To reduce the cost of
025 verifying the response of LLMs, OptReward in MiniOpt verifies the complete-
026 ness of problem modeling and avoids the need for content validation. The above
027 techniques enable the training of small-scale LLMs with strong optimization gen-
028 eralization ability under limited resources, thereby resulting in low inference cost
029 for localized deployment and usage. Extensive experiments show that MiniOpt-
030 3B exhibits strong optimization generalization across various optimization types
031 and scenarios. For models with parameters fewer than 10B, MiniOpt-3B achieves
032 the highest average solving accuracy (SA). For models with more than 10B param-
033 eters, MiniOpt-3B still shows competitive performance. Notably, MiniOpt-3B in-
034 dicates superior SA on the hard OptMATH-Bench while only consuming 37.64%
035 of the average output tokens required by DeepSeek-R1. The code is available at
036 <https://anonymous.4open.science/r/MiniOpt-6194>.

037 1 INTRODUCTION

038 Optimization problems are ubiquitous in real-world scenarios, profoundly affecting diverse domains,
039 including industrial production and transportation planning (Song et al., 2023; Li et al., 2025b).
040 While traditional optimization solvers are efficient, their application heavily relies on expert knowl-
041 edge, requiring the manual conversion of problems described in natural language into precise math-
042 ematical formulations or code, which is a process that both time-consuming and non-generalizable.
043 The rise of LLMs has opened new pathways for the automated modeling and solving of optimization
044 problems using natural language descriptions (Deng et al., 2024; Sun et al., 2024; Li et al., 2025a),
045 bringing them closer to practical application scenarios. Representative work like LLMOPT (Jiang
046 et al., 2025) or Text2Zinc (Singirikonda et al., 2025) significantly advances this field by parsing
047 natural language descriptions into structured formulation, which are a unified general expression for
048 optimization problems, and subsequently generating solution codes efficiently.

049 However, deploying such LLM-based approaches faces three critical bottlenecks. First, ensuring
050 accurate text generations on a smaller parameter scale requires Supervised Fine-Tuning (SFT) with
051 a large amount of high-quality training data (Wu et al., 2025; Lu et al., 2025). However, obtaining
052 such data is very difficult, often requires a lot of time and effort, and is prone to errors (Xiao et al.,
053 2025b). Second, it is challenging to verify whether the result generated meets the requirements (Zhai
et al., 2025). The non-verifiability of this task also exposes the limitations of the learning-based

054 solving paradigm. It is a barrier to the accuracy of the final result. Finally, to ensure the reliability
 055 of the result, common methods include reflection or self-correction mechanisms (Jiang et al., 2025;
 056 Li et al., 2025c), or multi-agent chains (AhmadiTeshnizi et al., 2024). However, these methods
 057 significantly multiply the computational overhead. These disadvantages pose significant barriers to
 058 implementing LLMs with powerful optimization-solving capabilities in the business scenarios of
 059 small and medium-sized enterprises or even on mobile devices. At the same time, considering data
 060 privacy issues, it is significant to locally deploy small-scale LLMs with strong performance while
 061 reducing training costs, including the volume and quality of training data.

062 To address these problems, this paper proposes a framework called MiniOpt to train a small-scale
 063 LLM with excellent optimization generalization under limited resources. To reduce the demand for
 064 data volume, MiniOpt adopts a two-stage RL training pipeline. Before conducting RL training, an
 065 SFT warm-up is applied to enable the model to get effective rewards in the early stages of RL. Sub-
 066 sequentially, in the first RL stage, the model rapidly acquires the model-and-solve paradigm. In the
 067 second RL stage, training is conducted on a high-quality data subset obtained through type-scenario
 068 guided data selection, yielding strong optimization generalization (Jiang et al., 2025). The RL pro-
 069 cess is training using the proposed OptGRPO algorithm, which enhances data utilization efficiency
 070 and improves the model’s ability to solve complex problems. To reduce the validation cost, we in-
 071 troduce OptReward in MiniOpt, which will perform three tasks: validating data format, ensuring the
 072 completeness of optimization problem modeling, and verifying solution accuracy. This encourages
 073 the model to learn in accordance with the model-and-solve paradigm. The framework enables the
 074 training of small-parameter LLMs with strong optimization generalization capabilities under limited
 075 resources, making localized deployment of LLMs for optimization feasible, thereby resulting in low
 076 inference cost for localized deployment and usage. Moreover, due to the high solving efficiency of
 077 MiniOpt-3B, reducing token usage also leads to considerable economic benefits.

078 Building upon the aforementioned methodology, this paper conducts extensive experiments with
 079 MiniOpt-3B on 9 benchmarks across different optimization types and problem scenarios. The re-
 080 sults demonstrate its high optimization generalization capability. Compared to baselines below 10B
 081 parameters, MiniOpt-3B achieves the best performance. When evaluated against baselines exceed-
 082 ing 10B parameters, MiniOpt-3B exhibits only a 1.57% lower average solving accuracy (SA) than
 083 DeepSeek-R1, 0.37% higher than GPT-5 and outperforms LLMOPT-14B by 2.13% in average SA.
 084 Notably, MiniOpt Pareto dominates other baselines in terms of both parameter scale and the SA
 085 metric. On the hard OPTMATH-Bench, MiniOpt-3B has a higher SA than DeepSeek-R1 and only
 086 consumes about 37.64% of DeepSeek-R1’s token count. Furthermore, results from ablation studies
 087 indicate that RL training contributes most significantly to MiniOpt-3B.

088 The subsequent sections review the related work, introduce the MiniOpt framework, present exper-
 089 imental results and analysis, provide an in-depth discussion, and finally conclude the paper.

090 2 RELATED WORK

092 **LLMs for Modeling and Solving Optimization Problems.** For modeling and solving optimization
 093 problems with LLMs, there are already a variety of benchmarks (Huang et al., 2025b; AhmadiTesh-
 094 nizi et al., 2024; Yang et al., 2024b). Challenging benchmarks like OptMATH (Lu et al., 2025) and
 095 NL4Opt (Huang et al., 2025a) have led to numerous studies utilizing LLMs to solve optimization
 096 problems. Prompt-based approaches such as OptiMUS (AhmadiTeshnizi et al., 2024), CoE (Xiao
 097 et al., 2024) and LEAN-LLM-OPT (Liang et al., 2025) utilizes the powerful generation capability of
 098 LLMs to generate the solver code of the optimization problem through multi-stage pipeline, without
 099 performing any post-training. Learning-based methods enhance LLMs’ capabilities in modeling and
 100 solving mathematical problems. For example, LLaMoCo (Ma et al., 2024) proposes an SFT-based
 101 framework comprising a meticulously designed instruction set and a two-stage training methodol-
 102 ogy that incorporates contrastive learning warm-up followed by SFT. LLMOPT (Jiang et al., 2025)
 103 and NER4OPT (Dakle et al., 2023; Singirikonda et al., 2025) adopts a two-stage training process of
 104 modeling the optimization problems first and then solving them by generating solution code.

105 **Reinforcement Learning with Verifiable Reward.** While Reinforcement Learning from Human
 106 Feedback (RLHF) (Ouyang et al., 2022) plays a crucial role in post-training alignment, it suffers
 107 from high annotation costs and inherent human bias (Xiao et al., 2025b). Reinforcement Learning
 108 with Verifiable Reward (RLVR) (Lambert et al., 2024) leverages externally grounded, easily verifi-

108 able rewards (e.g., rule-based reward) to provide dense and structurally simple supervision (Wang
 109 et al., 2024; Xie et al., 2025; Gao et al., 2024). Its practicality is especially valuable in real-world
 110 black-box systems (Zhang et al., 2025; Xin et al., 2025; Pan et al., 2025), where verification is typ-
 111 ically feasible only at the output stage, making RLVR a broadly applicable framework for aligning
 112 LLMs. In the field LLMs for solving optimization problems, SIRL (Chen et al., 2025) and OR-
 113 R1 (Ding et al., 2025) significantly improves the model’s performance through the RLVR training
 114 paradigm.

3 METHODOLOGY: THE PROPOSED MINIOPT

3.1 OVERVIEW

This paper studies how to endow small-scale LLMs with strong optimization generalization under tight data and compute budgets. We introduce MiniOpt, whose framework is shown in Figure 1. MiniOpt introduces the reasoning process within the model-and-solve paradigm for optimization problems, the reward function and algorithm used for RL training, and the training pipeline for small-scale models. Through this framework, MiniOpt formulates the path from a natural language problem to an executable solver code as a single verifiable end-to-end task.

3.2 REASONING TO MODEL AND SOLVE PARADIGM

As shown in subfigure (a) in Figure 1, we introduce a reasoning to model and solve paradigm that turns a natural language optimization problem into a single verifiable objective. The paradigm is enforced by two compulsory and machine-parsable segments `<think>...</think>` and `<answer>...</answer>`.

The first segment, enclosed by `<think>...</think>`, carries all modeling contents. It performs an analysis of the problem statement, specifies a complete five-element formulation inspired by LLMOPT (Jiang et al., 2025) and determines the appropriate open-source solver. Specifically, the optimization problem can be described as minimizing the **objective** function $f(x)$ subject to the **constraints** $G(x) \leq c$, where the $x \in \mathcal{X} \subseteq \mathbb{R}^D$ is the D -dimensional decision **variables**, and $I = \{1, 2, \dots, D\}$ is the index **sets**. \mathcal{X} is the feasible region, and $c \in \mathbb{R}^m$ provides the upper bounds. The constants in $f(x)$ and $G(x)$ form the **parameters** set, which also includes c . The five-element formulation, $\mathcal{M} = (\text{Variables, Objective, Constraints, Sets, Parameters})$ map one-to-one to the components of an optimization problem. More details about the five-element formulation are provided in Appendix C.

During the reasoning process, to enhance the capability of LLM for solving optimization problems across diverse optimization types, it will analyze the problem during the thinking (i.e., reasoning) phase and select different open-source solvers for the given problem depending on the optimization types and the characteristics of the solvers, thereby ensuring an optimized match between the problem and the solver. Solver selection is guided by a prompt, which is introduced in Appendix M.

The second segment, enclosed by `<answer>...</answer>`, converts the five-element formulation into an executable Pyomo program that models the problem, invokes the solver, solves the instance, and prints the numerical answer. By constraining responses in this manner, we collapse the action space from free-form language to a programmable artifact whose intermediate structure and final result can be deterministically parsed and verified. The `<answer>` segment must implement the blueprint as a single python code fence containing a complete Pyomo script. Because these outputs are easy to verify, the OptReward in Section 3.4.1 can score format, five-element formulation, and accuracy based on rules via an automated procedure, providing low-cost supervision. During the training process, the prompt corresponding to this paradigm is shown in Appendix N.

3.3 TRAINING PIPELINE OF MINIOPT

Based on the paradigm described in Section 3.2 and the key techniques introduced in Section 3.4, we propose a training pipeline that enables MiniOpt to learn from limited resource constraints while achieving powerful solving performance and strong optimization generalization. This pipeline be-

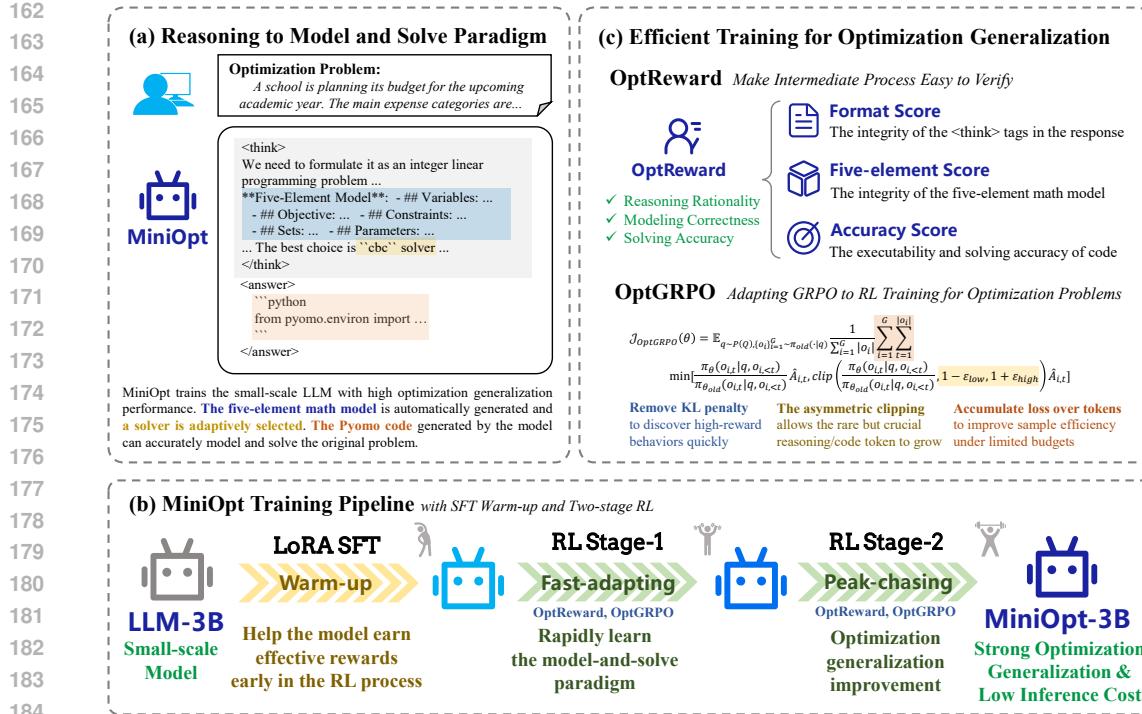


Figure 1: The framework of MiniOpt. Sub-figure (a) demonstrates the reasoning to model and solve paradigm of MiniOpt, encompassing problem modeling and solver adaptation during the thinking (i.e., reasoning) process of RL, and solution code generation in the response. Sub-figure (b) illustrates the training pipeline of MiniOpt, which involves sequential execution of SFT warm-up followed by two-stage RL. Sub-figure (c) presents the reward function OptReward and training algorithm OptGRPO used in MiniOpt’s RL training.

gins with a lightweight SFT warm-up followed by two-stage RL under OptReward. The illustration of the training pipeline is shown in subfigure (b) in Figure 1.

3.3.1 WARM-UP BASED ON LIGHTWEIGHT SFT

Small scale LLMs struggle to produce an end-to-end, executable trajectory for complex optimization tasks, **thus if a model lacks fundamental problem-solving capabilities for operations optimization, directly applying reinforcement learning may lead to sparse rewards and result in unstable training.**

We therefore apply an SFT warm-up whose purpose is to provide a run-through capability, thereby allowing the model to obtain effective rewards in the early stages of RL training. This SFT warm-up does not adopt the reasoning to model and solve paradigm introduced in Section 3.2. Instead, it establishes a starting point upon which our two-stage RL can subsequently focus on paradigm acquisition and optimization generalization. The training data for warm-up are constructed from OptMATH-Train. For each instance, we prompt Qwen2.5-Coder-32B-Instruct (Hui et al., 2024) using the prompt in Appendix L to rewrite the original GurobiPy program into an equivalent Pyomo implementation and to select an appropriate open-source solver according to the detected structure. Every rewritten instance is executed and only those that compile, solve, and print the correct structure are retained. The remaining data serves as a candidate pool for the next stage of the two-stage RL training.

3.3.2 THE TWO-STAGE RL

After undergoing SFT warm-up, the models can conduct RL training more effectively. We employ a two-stage RL training under the same reward function (OptReward) for paradigm acquisition and optimization generalization, respectively. Both stages share the pipeline of parsing, executing,

216 and scoring in Section 3.2 and the OptReward of Sections 3.4.1, they differ in training data and
 217 hyperparameter setting. Algorithmically, both stages use the same OptGRPO (cf. paragraphs in
 218 Section 3.4.2).

219 Stage-1 aims to enable the model to acquire the reasoning to model and solve Paradigm: the model
 220 must generate a valid `<think>/<answer>` pair, produce the executable Pyomo code, and make
 221 a coherent solver choice. To this end we train on 1,585 relatively easy problems, so that most signal
 222 arises from the formatting and structural components of OptReward, rapidly improving executability
 223 and the ability to solve the optimization problem of natural language description.

224 Stage-2 focuses on optimization generalization once the paradigm is established. The training dis-
 225 tribution shifts to the problems with diverse optimization types and problem scenario, and the opti-
 226 mization emphasis moves to the accuracy score, encouraging refined modeling and solving behaviors
 227 (e.g., variable/constraint formulation and solver selection).

228 Based on the annotations of optimization types and problem scenario tags this paper assigned to
 229 each instance in OptMATH-Train, we sample a data subset from the candidate pool mentioned in
 230 Section 3.3.1 subject to two constraints: (i) type-uniform coverage, with exactly 600 instances per
 231 type, and (ii) within each type, the scenario frequencies match the distribution in the full dataset.
 232 The resulting data serves as the training data for the second stage of RL. While the dataset of the
 233 first stage is the union of the NL4Opt (AhmadiTeshnizi et al., 2024) and ICML Competition (Yang
 234 et al., 2024b) training splits. Detailed information on the construction of the training set is provided
 235 in the Appendix E.

236 Such a training strategy makes efficient use of limited data and reduces training costs, and ultimately
 237 allows for strong optimization generalization of LLMs with even small parameters and limited com-
 238 putational resources.

242 3.4 EFFICIENT TRAINING FOR OPTIMIZATION GENERALIZATION

244 Building upon the reasoning to model and solve paradigm mentioned in Section 3.2, we propose two
 245 key components for the RL training of MiniOpt as shown in subfigure(c) in Figure 1. An informative
 246 and easily verifiable reward function OptReward, and an improved algorithm OptGRPO builds upon
 247 GRPO (Shao et al., 2024).

251 3.4.1 OPTREWARD: VERIFIABLE REWARDS DESIGNED FOR MINIOP

253 Based on the reasoning to model and solve paradigm described in section 3.2, we propose OptRe-
 254 ward, including three automatically computed components: formatting correctness, structural suffi-
 255 ciency, and numerical accuracy. Each component is derived from deterministic parsing or execution
 256 of the output, It ensures the completeness of the modeling of problems but also enables verification
 257 to scale with machine time, yielding substantially lower verification costs.

258 **Format Score:** The format score S_{fmt} validates the response format. A response must contain
 259 exactly one `<think>...</think>` and one `<answer>...</answer>` in the correct order.
 260 If all conditions hold, we assign $S_{\text{fmt}} = +1$, otherwise $S_{\text{fmt}} = -1$. If the specified format is not
 261 present in the response, we deterministically set the remaining components to their error defaults,
 262 $S_{\text{five}} = -1$ and $S_{\text{acc}} = -2$, so that the total reward immediately reaches the global minimum.
 263 This forces the model to adopt the correct response format early in training and prevents expensive
 264 evaluation of malformed samples.

265 **Five-element Score:** To avoid ground-truth labeling of five-element content and the bias it may
 266 introduce, we use a presence-based rule aligned with the paradigm in Section 3.2. Conditional on
 267 valid formatting, we compute the five-element score S_{five} . In the `<think>` segment, the model
 268 response is expected to include five labeled summaries starting with “## Sets:”, “## Parameters:”,
 269 “## Variables:”, “## Objective:”, and “## Constraints:”. Each present summary adds 0.2 scores; if
 none is present we assign $S_{\text{five}} = -1$. This structure shaping keeps the modeling blueprint parsable.

$$S_{\text{five}} = \begin{cases} 0.2 \sum_{k=1}^5 I_k, & \text{if } \sum_{k=1}^5 I_k \geq 1, \\ -1, & \text{if } \sum_{k=1}^5 I_k = 0, \end{cases} \quad (1)$$

where $I_k = 1$ if the k -th required element e_k is present in the `<think>` segment, and $I_k = 0$ otherwise. $(e_1, \dots, e_5) = (\text{Sets}, \text{Parameters}, \text{Variables}, \text{Objective}, \text{Constraints})$.

Accuracy Score: The accuracy score is obtained by executing the Pyomo program contained in `<answer>`. Failure to extract a program or to complete execution yields $S_{\text{acc}} = -2$. When execution succeeds, we retrieve the optimal objective value \hat{f} from the model output and compare it with the ground-truth value f^* ; if they are equal, we assign $S_{\text{acc}} = 2$, otherwise $S_{\text{acc}} = -1.5$.

$$S_{\text{acc}} = \begin{cases} +2, & \text{if execution succeeds and } \hat{f} = f^*, \\ -1.5, & \text{if execution succeeds but } \hat{f} \neq f^*, \\ -2, & \text{if no executable code or execution fails.} \end{cases} \quad (2)$$

Combining the components with the formatting gate yields the total OptReward as follow:

$$R = \begin{cases} -4, & \text{if } S_{\text{fmt}} = -1, \\ S_{\text{fmt}} + S_{\text{five}} + S_{\text{acc}}, & \text{if } S_{\text{fmt}} = 1. \end{cases} \quad (3)$$

By OptReward, the format score enforces the strict `<think>/<answer>` format, the five-element score shapes a complete modeling blueprint in the think phase, and the accuracy score certifies correctness through equality of optimal objective values, enabling low-cost verifiable RL for problems in the field of optimization.

3.4.2 OPTGRPO: TRAINING SMALL-SCALE LLMs WITH LIMITED RESOURCES

GRPO (Shao et al., 2024) is an efficient RL algorithm, which replaces the critic model with a group baseline and updates the policy at the group level so that improves training stability and efficiency. The details of the algorithm are introduced in Appendix D. [In order to adapt to the designed OptReward and Two-Stage RL in this paper, taking inspiration from the previous work on the improvement of GRPO like DAPO](#) (Yu et al., 2025), we introduces three modifications, with the goal of achieving stronger optimization generalization within a limited budget.

First, building upon the strict gating already provided by OptReward through formatting score and five-element score, we set the coefficient on the KL penalty $\beta = 0$ to remove the KL penalty to encourage greater exploration by the model. This approach does not jeopardize stability while facilitating faster discovery of high-reward behaviors. Second, to prevent entropy collapse and allow rare but crucial reasoning/code tokens to grow, we replace symmetric clipping with an asymmetric interval $[1 - \varepsilon_{\text{low}}, 1 + \varepsilon_{\text{high}}]$ with a higher upper clipping threshold $\varepsilon_{\text{high}}$ than the lower threshold ε_{low} . This relaxes the trust region on probability increases while keeping a firm lower bound on decreases, which empirically improves executability in stage-1 RL training and supports generalization in stage-2 RL training. Following Specifically, we raise $\varepsilon_{\text{high}}$ to 0.28 during the training. Finally, to improve sample efficiency, critical for small-scale LLMs under limited data and verification budgets, instead of optimizing a sequence-averaged loss, we accumulate the loss over tokens and normalize by the total number of tokens in the group $\sum_i |o_i|$, which can yield denser learning signals for long outputs. The final loss of OptGRPO is as follows:

$$\begin{aligned} \mathcal{J}_{\text{OptGRPO}}(\theta) = & \mathbb{E}_{q \sim P(Q), \{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot | q)} \frac{1}{\sum_{i=1}^G |o_i|} \sum_{i=1}^G \sum_{t=1}^{|o_i|} \\ & \min \left[\frac{\pi_{\theta}(o_{i,t}|q, o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t}|q, o_{i,<t})} \hat{A}_{i,t}, \text{clip} \left(\frac{\pi_{\theta}(o_{i,t}|q, o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t}|q, o_{i,<t})}, 1 - \varepsilon_{\text{low}}, 1 + \varepsilon_{\text{high}} \right) \hat{A}_{i,t} \right]. \end{aligned} \quad (4)$$

324

4 EXPERIMENT

325
 326 We evaluate MiniOpt models on diverse optimization benchmarks spanning multiple types and sce-
 327 narios to assess whether small scale parameter LLMs (3B/7B/14B) can achieve strong optimization
 328 generalization ability. We report Solving Accuracy (SA) as primary metrics, SA presents the propor-
 329 tion of output solutions of executed code that are numerically equal to the optimal solution provided
 330 from the labels of benchmarks. We use Execution Rate (ER) as a supplement that shows the pro-
 331 portion of generated code samples that run successfully without errors. Comparisons cover general
 332 LLMs, general reasoning LLMs, prompting-based baselines and learning-based baselines. The ex-
 333 periments aim to answer four key questions below.

334 **(Q1) Optimization Generalization Ability at Small-Scale LLMs:** To what extent can MiniOpt
 335 at 3B/7B/14B achieve high SA across types and scenarios, and how does it compare with larger
 336 reasoning LLMs and prior learning-based approaches?

337 **(Q2) Pareto Front of Performance vs. Cost:** What is the limit of the scale of model parameters
 338 for achieving strong optimization generalization?

339 **(Q3) Importance of the Training Pipeline of MiniOpt:** How critical are the lightweight SFT
 340 warm-up and the two-stage RL to the performance of MiniOpt?

341 **(Q4) Importance of the Reasoning to Model and Solve Paradigm and OptReward:** How do the
 342 proposed paradigm and the corresponding Opt Reward contribute to boost SA and ER in solving
 343 optimization problems?

344 The four questions are answered sequentially in the following sections, which first provide a detailed
 345 introduction to the experiments and then present an analysis of the results.

346

4.1 EXPERIMENTAL SETUP

347 Since widely-used packages like Gurobi and COPT are close-source, generating code for them may
 348 entail licensing costs, and a model’s generalization capability across different modeling languages
 349 and solvers depends on the proportion of relevant corpora in its pre-training and post-training data.
 350 Therefore, following established practices (Jiang et al., 2025), we adopted the open-source, solver-
 351 agnostic Pyomo modeling language as the language for our training and inference solver code.

352 As for the solvers employed, the key consideration lies in their ability to select the appropriate
 353 solver for a specific optimization type. For instance, GLPK only supports linear programming and
 354 integer programming and is incapable of solving nonlinear problems. Therefore, this paper selects
 355 four types of solvers to cover the solving requirements of as many optimization problem types as
 356 possible, so as to automatically adapt to the problem types when generating the solving code.

357 The evaluation encompasses nine benchmarks about operations research optimization: NL4Opt
 358 (Ramamonjison et al., 2022), Mamo (Easy and Complex subsets) (Huang et al., 2025b), Indus-
 359 tryOR (Huang et al., 2025a), NLP4LP (AhmadiTeshnizi et al., 2024), ComplexOR (Xiao et al.,
 360 2024), OptMATH (Lu et al., 2025), OptiBench (Yang et al., 2025b), and ICML Competition (Yang
 361 et al., 2024b). We follow the same setting in LLMOPT (Jiang et al., 2025) to ensure consistency
 362 and comparability. For the newly included datasets, OptMATH-Bench and OptiBench, we adhere to
 363 their original data splits provided by the authors. For ICML Competition dataset, we use it exclu-
 364 sively as out-of-distribution test data to evaluate the generalization capability of our model.

365 To validate the correctness of the solutions obtained by MiniOpt and all baselines, this paper com-
 366 pares their log files generated during the execution process against the ground-truth solutions pro-
 367 vided by the benchmarks. A solution is deemed correct if the log file yields a non-empty match with
 368 the ground-truth; otherwise, it is considered unsuccessful.

369

4.2 ANALYSIS OF OPTIMIZATION GENERALIZATION

370 In this section, we compare MiniOpt with general LLMs (Qwen2.5 Series, DeepSeek-V3), general
 371 reasoning LLMs (Qwen3 Series, DeepSeek-R1, Gemini-2.5-Pro, GPT-5), prompt-based methods
 372 (Chain-of-Experts, OptiMUS, Reflexion), learning-based methods (OptMATH-7B, LLMOPT-14B,
 373 Step-OPt-8B/7B/3B), demonstrating the optimization generalization capability of MiniOpt. The

378
379
380
381
382Table 1: Comparison of the SA metric across 9 benchmarks with rankings (NL4Opt, ICML Competition, Mamo Easy, Mamo Complex, NLP4LP, ComplexOR, IndustryOR, OptiBench, OptMATH-Bench). **Bold** indicates 1st, wavy underline indicates 2nd, underline indicates 3rd. “Rank*” represents the result of sorting methods among parameter sizes below 10B.

Category	Models / Methods	Avg.	Rank	Rank*	NL4Opt	ICML.C	Mamo.E	Mamo.C	NLP4LP	Com.OR	Indus.OR	OptiBench	OptMATH-Bench
General Models	Qwen2.5-3B-Instruct	9.98	22	11	19.13	18.78	17.18	2.37	18.60	0.00	2.00	11.74	0.00
	Qwen2.5-7B-Instruct	30.05	17	8	<u>53.48</u>	51.71	35.58	4.27	55.37	16.67	13.00	35.54	4.82
	Qwen2.5-14B-Instruct	43.59	12	-	67.39	63.17	80.21	15.64	67.36	22.22	22.00	41.65	12.65
	DeepSeek-V3	57.81	4	-	78.26	<u>77.56</u>	84.82	26.54	<u>79.34</u>	44.44	26.00	64.13	39.16
General Models (Thinking)	Qwen3-4B	10.25	21	10	16.52	17.56	13.80	6.64	15.29	5.56	2.00	11.90	3.01
	Qwen3-8B	19.77	19	9	30.87	29.51	23.93	9.95	34.71	11.11	6.00	28.26	3.61
	Qwen3-14B	22.54	18	-	24.35	22.68	36.20	11.37	19.42	38.89	15.00	22.31	12.65
	DeepSeek-R1	<u>58.51</u>	<u>3</u>	-	<u>83.91</u>	75.37	74.54	<u>39.81</u>	69.83	44.44	32.00	66.94	39.76
	Gemini-2.5-Pro	57.04	5	-	78.26	71.22	65.95	30.81	73.55	50.00	28.00	61.32	<u>54.21</u>
Prompt-based Methods	GPT-5	56.57	7	-	80.43	73.66	58.12	23.22	73.14	61.11	26.00	<u>64.63</u>	48.80
	Chain-of-Experts	41.03	14	-	66.52	56.59	63.65	22.75	59.09	33.33	19.00	45.29	3.01
	OptiMUS	18.76	20	-	13.48	33.17	37.27	11.85	18.18	16.67	8.00	26.61	3.61
Learning-based Models	Reflexion	41.28	13	-	56.52	52.20	84.82	18.01	53.72	38.89	19.00	41.16	7.23
	Step-OPT-LLaMA-3.2-3B	<u>39.61</u>	<u>15</u>	6	70.00	26.59	68.10	36.49	63.64	38.89	17.00	26.78	9.04
	Step-OPT-LLaMA-3-8B	<u>50.80</u>	10	4	<u>75.22</u>	<u>68.54</u>	<u>79.45</u>	50.71	64.88	27.78	27.00	49.75	13.86
	Step-OPT-Qwen2.5-3B	<u>36.15</u>	<u>16</u>	7	<u>41.30</u>	<u>38.54</u>	<u>75.31</u>	20.85	<u>53.31</u>	<u>27.78</u>	<u>21.00</u>	<u>40.00</u>	<u>7.23</u>
	Step-OPT-Qwen2.5-7B	<u>47.42</u>	11	5	<u>77.83</u>	<u>57.32</u>	<u>69.33</u>	<u>50.24</u>	48.35	38.89	27.00	48.76	9.04
	OptMATH-7B	52.37	9	3	78.70	66.83	84.20	<u>34.12</u>	68.60	33.33	19.00	52.23	34.34
Ours	LLMOPT-14B	54.81	8	-	80.28	<u>75.35</u>	<u>89.53</u>	44.08	73.42	35.29	<u>29.00</u>	53.83	12.50
	MiniOpt-3B	56.94	6	2	83.04	68.05	85.43	35.07	73.55	50.00	21.00	53.55	42.77
	MiniOpt-7B	<u>62.76</u>	<u>2</u>	1	<u>89.13</u>	<u>77.56</u>	<u>88.34</u>	38.39	<u>79.34</u>	55.56	26.00	59.34	51.20
396 397	MiniOpt-14B	<u>66.10</u>	<u>1</u>	-	<u>92.17</u>	<u>86.34</u>	90.80	33.65	<u>79.75</u>	61.11	27.00	67.44	56.63

398
399

information on these methods can be found in Appendix B. Tables 1 summarizes SA on nine benchmarks that span 7 optimization types and 22 scenarios, and the statistics on problem categories and scenarios are provided in Appendix A.2. The ER metrics for all the methods on the nine benchmarks and its analysis can be found in the Appendix F.

400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416

Overall Performance (Answer to Q1). Across all nine benchmarks, MiniOpt-7B achieved the strongest average performance among all baselines. Notably, MiniOpt-3B surpassed all prompt-based and learning-based methods. For instance, compared to LLMOPT-14B, the state-of-the-art learning-based model, MiniOpt-3B achieved an average SA that is 2.13% higher. When compared to DeepSeek-V3 and DeepSeek-R1, MiniOpt-3B scored 0.87% and 1.57% lower, respectively. Furthermore, as a higher-parameter variant of MiniOpt, MiniOpt-14B achieves the highest average SA of 66.10%, significantly raising the performance ceiling of MiniOpt in practical applications.

417
418
419
420
421
422
423

Besides, SIRL (Chen et al., 2025) is not included in Table 1 because the paper was accepted only shortly before our submission deadline. Although it proposes a reasoning model training pipeline, the important issue of improving the optimization generalization ability of small-scale models with limited data and computing resources has not been studied. According to the results presented in the SIRL paper (Chen et al., 2025), MiniOpt-7B outperforms the SIRL-Qwen2.5-7B model of similar size by 22.20% solving accuracy on the hard OptMATH-Bench benchmark.

424
425
426
427
428
429

Competitiveness of Small-Scale Models (Answer to Q1). MiniOpt remains competitive even at smaller scales. MiniOpt-3B reaches 56.94% SA on average, this performance already matching or exceeding several much larger reasoning models (e.g., the average SA higher than GPT-5 at 56.57%) and clearly outperforming general-purpose 14B post-trained models (e.g., +13.35% over Qwen2.5-14B-Instruct on average). Performance grows smoothly with scale under the same training pipeline, The average SA metric grows by +5.82% when parameters increase from 3B to 7B, and by +9.16% when further scaled to 14B.

430
431

Challenging Benchmarks (Answer to Q1). On the most demanding sets that require faithful modeling and solver usage, MiniOpt shows clear advantages. On the latest challenging benchmark OptMATH-Bench, MiniOpt-14B achieves 56.63% SA, outperforming Gemini-2.5-Pro (54.21%) and GPT-5 (48.80%). Even on extremely high-dimensional test sets such as Indus.OR, where small-parameter MiniOpt does not attain the highest SA due to limitations in instruction following, it still delivers competitive performance levels in both metrics.

Breadth across Types and Scenarios (Answer to Q1). To rigorously evaluate the optimization generalization, we analyze the SA of MiniOpt-7B and MiniOpt-3B across three optimization types and three application scenarios under two difficulty levels. At the medium-difficulty OptiBench,

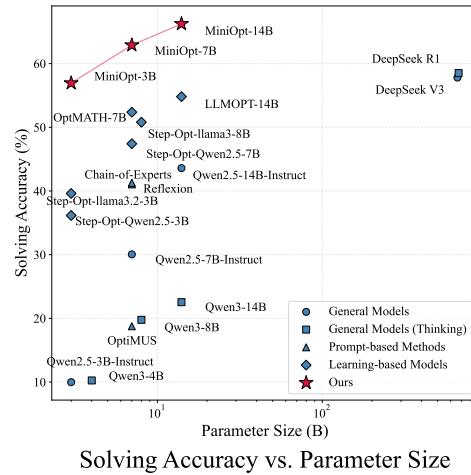
432 MiniOpt-7B achieved 59% (M), 74% (T), 68% (L), and 65% (LP), 63% (IP), 45% in Mixed-Integer
 433 Linear Programming (MILP). MiniOpt-3B scored 55% (M), 71% (T), 59% (L), and 60% (LP),
 434 58% (IP), 38% (MILP). On the hard-difficulty OptMATH-Bench, MiniOpt-7B achieved 58% (M),
 435 56% (T), 57% (L), and 55% (LP), 100% (IP), 56% (MILP). MiniOpt-3B attained 43% (M), 56%
 436 (T), 76% (L), and 73% (LP), 67% (IP), 47% (MILP). These results reflect consistent and scalable
 437 generalization across optimization types, problem scenarios, and difficulty levels. In order to in-
 438 vestigate whether the pipeline training in this paper would impair other abilities of the LLMs, we
 439 also discussed the generalization performance of MiniOpt-3B and 7B under different tasks except
 440 optimization problems. Detailed information is presented in Appendix J.

441 442 4.3 PARETO FRONT OF PERFORMANCE VS. COST

443 **Analysis of the Pareto Front (Answer to Q2).** Figure 2 and Figure 5 in Appendix H indicate that
 444 MiniOpt series (represented by the solid red line) establishes a new and superior Pareto front in the
 445 performance-versus-cost trade-off. Since the parameter size of GPT-5 and Gemini-2.5-Pro have not
 446 been disclosed, we do not label these two models in the figures). As the scale of the model increases,
 447 the average SA performance of MiniOpt also grows steadily. The MiniOpt-14B variant achieves the
 448 highest average SA of 66.10% among all models. It achieves a comprehensive performance lead
 449 while having substantially fewer parameters than top-tier general reasoning models such as DeepSeek-
 450 R1. **Compared to the similar modeling and solving** model Step-OPT, Step-OPT-LLaMA-3-8B achieves
 451 an average SA of 50.80% across 9 benchmarks, while the proposed MiniOPT attains an average SA
 452 of 62.76%. When the parameter scale of both models are reduced to 3B, Step-OPT-LLaMA-3.2-3B ex-
 453 hibit a performance drop of 11.19%. For Step-OPT-
 454 Qwen2.5, the SA degradation is 11.27%, **while MiniOPT only decreased by 5.82%**. This indicates that
 455 the key advantage of MiniOPT lies in its ability to
 456 maintain superior performance even with reduced parameter scale. From the perspective of capa-
 457 bility density (Xiao et al., 2025a), MiniOPT effectively achieves lower parameter requirements and
 458 inference costs while preserving comparable performance.

469 470 4.4 ABLATION STUDY

471 **Importance of the Reasoning to Model and Solve Paradigm and OptReward (Answer to Q4).**
 472 We ablate core components of MiniOpt-3B and report results of the SA metric in Table 2 and the ER
 473 metric in Table 6 in Appendix G. As evidenced in Table 2, each module of the proposed reasoning to
 474 model and solve paradigm demonstrates substantial contributions to modeling and solving optimiza-
 475 tion problems with smaller-scale models under limited training resources. Among these, RL pro-
 476 vides the most significant improvement, highlighting the importance of the proposed paradigm. The
 477 reasoning to model and solve paradigm and OptReward together turn free-form generation into a ver-
 478 ifiable formulation, which are easy to verify. Removing them and keeping only a final-answer signal
 479 (w/o OptReward) drops averages to 52.44% and 83.39% (Δ SA = -4.50%, Δ ER = -4.65%), respec-
 480 tively. The largest losses appear where the problems are challenging: Com.OR (Δ SA = -16.67%,
 481 Δ ER = -16.67%), Indus.OR (Δ SA = -3.00%, Δ ER = -6.00%), and OptMATH-Bench (Δ SA =
 482 -4.82%, Δ ER = -10.24%). These patterns align with the roles of the three reward components:
 483 the format score enforces the labels <think> and <answer> in the responses are complete; the
 484 five-element score shapes the intermediate blueprint, so the models learns to extract problem struc-
 485 ture before coding; the Accuracy Score certifies numerical correctness by executing the Pyomo code
 and comparing the returned optimum with the reference. In combination, this reasoning to model
 and solve paradigm, together with verifiable reward, steers learning toward structurally consistent



Solving Accuracy vs. Parameter Size

Figure 2: Comparison of average SA against model parameter scales for various methods. MiniOpt is the Pareto optimal among compared methods.

486

487

Table 2: Ablation study (MiniOpt-3B) on the SA metric across 9 benchmarks.

488

489

490

491

492

493

494

495

496

497

498

499

500

501

502 5 DISCUSSION

503

504

505

506

507

508

509

510

511

behaviors, yielding higher SA and ER across various optimization problems. For a detailed ablation analysis of the proposed model-and-solve paradigm, the training process, and the significance of OptReward and OptGRPO, please see Appendix G.

Exploration of Models with Smaller Parameter Scale. To further explore the performance of MiniOpt with smaller parameter scales, thereby better balancing the trade-off between parameters and performance, we deploy MiniOpt-1.5B with the same reasoning to model and solve paradigm (Section 3.2), OptReward (Section 3.4.1), and training pipeline (Section 3.3). The results shown in Table 7 and 8 in Appendix I demonstrate that when the model scale is reduced from 3B to 1.5B, MiniOpt achieves an average SA of 46.15% and an average ER of 80% across the 9 benchmarks, still surpassing all baselines except for learning-based methods, maintaining a strong competitiveness.

Dimensionality Scalability Issue. Our empirical results show that MiniOpt-3B can correctly model optimization problems with dimensionalities as high as 72 on the OptMATH-Bench task and 80 on the Mamo Complex task. This constitutes a remarkably strong performance for models of such limited scale, which highlights that the MiniOpt pipeline effectively enhances the model’s optimization modeling capacity and retains robust generalization even in high-dimensional settings. For problems with even higher dimensionality, the descriptions of numerical parameters typically become substantially longer. Such extended input sequences inherently pose a significant challenge to the comprehension capabilities of small-scale models. As corroborating evidence, larger models such as MiniOpt-7B successfully solve problems up to 109 dimensions on the OptMATH-Bench task.

Cost Savings of MiniOpt and Generality of Base Models. We find that MiniOpt-3B achieves a 3.01% higher SA than DeepSeek-R1 on the OptMATH-Bench while using 62.4% fewer average output tokens. For MiniOpt-7B, the corresponding improvements are 11.44% higher SA and a 39.6% reduction in average output tokens. More details and discussions are provided in Appendix K. This demonstrates that MiniOpt constitutes a general framework capable of enabling small-scale models to achieve strong optimization generalization under constrained data and computational resources. Furthermore, although Qwen2.5 series models are adopted as the base model in this paper, we have also conducted experiments with different base models and obtained closely aligned results.

528 6 CONCLUSION

529

530

531

532

533

534

535

536

537

538

539

This paper proposes a novel reasoning to model and solve paradigm and the small scale model, MiniOpt-3B, achieves higher performance with a small-scale parameter and limited resources. We explore the optimization generalization of the model in various types of optimization, problem scenarios, and high-variable dimensions. Empirical results demonstrate that MiniOpt exhibits strong generalization performance under these varying conditions. Furthermore, this study explores the minimum parameter scale required for MiniOpt to maintain competitive performance. Future work includes exploring efficient modeling and solving methods for optimization problems with high-dimensional variables or a large number of constraints.

540 ETHICS AND REPRODUCIBILITY STATEMENT
541542 **Ethics.** This work does not involve any human subjects, personal data, or sensitive information. All
543 the test datasets used are publicly available, and no proprietary or confidential information is used.
544545 **Reproducibility.** Experimental settings are described in Section 4.1 and datasets included
546 in Appendix E. The code is available at [https://anonymous.4open.science/r/](https://anonymous.4open.science/r/MiniOpt-6194)
547 MiniOpt-6194.548 LLM USAGE STATEMENT
549550 No LLMs were used in the research ideation and paper writing of this work.
551552 REFERENCES
553554 Ali AhmadiTeshnizi, Wenzhi Gao, and Madeleine Udell. OptiMUS: Scalable optimization mod-
555 eling with (MI)LP solvers and large language models. In *Advances in Forty-first International*
556 *Conference on Machine Learning*, Vienna, Austria, 2024.557 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared
558 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri,
559 Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan,
560 Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian,
561 Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios
562 Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex
563 Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders,
564 Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec
565 Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob Mc-
566 Grew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large
567 language models trained on code. *CoRR*, abs/2107.03374, 2021.568 Yitian Chen, Jingfan Xia, Siyu Shao, Dongdong Ge, and Yinyu Ye. Solver-Informed RL: Grounding
569 large language models for authentic optimization modeling. *CoRR*, abs/2505.11792, 2025.
570571 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
572 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John
573 Schulman. Training verifiers to solve math word problems. *CoRR*, abs/2110.14168, 2021.574 Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit S.
575 Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, Luke Marrs, Sam Petulla, Colin
576 Gaffney, Asaf Aharoni, Nathan Lintz, Tiago Cardal Pais, Henrik Jacobsson, Idan Szpektor, Nan-
577 Jiang Jiang, Krishna Haridasan, Ahmed Omran, Nikunj Saunshi, Dara Bahri, Gaurav Mishra, Eric
578 Chu, Toby Boyd, Brad Hekman, Aaron Parisi, Chaoyi Zhang, Kornraphop Kawintiranon, Tania
579 Bedrax-Weiss, Oliver Wang, Ya Xu, Ollie Purkiss, Uri Mendlovic, Ilai Deutel, Nam Nguyen,
580 Adam Langley, Flip Korn, Lucia Rossazza, Alexandre Ramé, Sagar Waghmare, Helen Miller,
581 Nathan Byrd, Ashrith Sheshan, Raia Hadsell Sangnie Bhardwaj, Pawel Janus, Tero Rissa, Dan
582 Horgan, Sharon Silver, Ayzaan Wahid, Sergey Brin, Yves Raimond, Klemen Kloboves, Cindy
583 Wang, Nitesh Bharadwaj Gundavarapu, Ilia Shumailov, Bo Wang, Mantas Pajarskas, Joe Hey-
584 ward, Martin Nikoltchev, Maciej Kula, Hao Zhou, Zachary Garrett, Sushant Kafle, Sercan Arik,
585 Ankita Goel, Mingyao Yang, Jiho Park, Koji Kojima, Parsa Mahmoudieh, Koray Kavukcuoglu,
586 Grace Chen, Doug Fritz, Anton Bulyenov, Sudeshna Roy, Dimitris Paparas, Hadar Shemtov, Bo-
587 Juen Chen, Robin Strudel, David Reitter, Aurko Roy, Andrey Vlasov, Changwan Ryu, Chas Le-
588 ichner, Haichuan Yang, Zelda Mariet, Denis Vnukov, Tim Sohn, Amy Stuart, Wei Liang, Minmin
589 Chen, Praynaa Rawlani, Christy Koh, JD Co-Reyes, Guangda Lai, Praseem Banzal, Dimitrios
590 Vytiniotis, Jieru Mei, and Mu Cai. Gemini 2.5: Pushing the frontier with advanced reasoning,
591 multimodality, long context, and next generation agentic capabilities. *CoRR*, abs/2507.06261,
592 2025.593 Parag Pravin Dakle, Serdar Kadioglu, Karthik Uppuluri, Regina Politi, Preethi Raghavan, SaiKr-
594 ishna Rallabandi, and Ravisutha Srinivasamurthy. Ner4Opt: Named entity recognition for opti-
595 mization modelling from natural language. In *Integration of Constraint Programming, Artificial*

594 *Intelligence, and Operations Research - 20th International Conference*, volume 13884, pp. 299–
 595 319, Nice, France, 2023.

596

597 DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Cheng-
 598 gang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang,
 599 Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting
 600 Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Haowei Zhang, Honghui
 601 Ding, Huajian Xin, Huazuo Gao, Hui Li, Hui Qu, J. L. Cai, Jian Liang, Jianzhong Guo, Jiaqi
 602 Ni, Jiashi Li, Jiawei Wang, Jin Chen, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li,
 603 Junxiao Song, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang,
 604 Lecong Zhang, Lei Xu, Leyi Xia, Liang Zhao, Litong Wang, Liyue Zhang, Meng Li, Miaojun
 605 Wang, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingming Li, Ning Tian, Panpan
 606 Huang, Peiyi Wang, Peng Zhang, Qiancheng Wang, Qihao Zhu, Qinyu Chen, Qiushi Du, R. J.
 607 Chen, R. L. Jin, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, Runxin Xu, Ruoyu Zhang,
 608 Ruyi Chen, S. S. Li, Shanghao Lu, Shangyan Zhou, Shanhua Chen, Shaoqing Wu, Shengfeng
 609 Ye, Shirong Ma, Shiyu Wang, Shuang Zhou, Shuiping Yu, Shunfeng Zhou, Shuting Pan, T. Wang,
 610 Tao Yun, Tian Pei, Tianyu Sun, W. L. Xiao, and Wangding Zeng. DeepSeek-V3 technical report.
CoRR, abs/2412.19437, 2024.

611

612 DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu,
 613 Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu,
 614 Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao
 615 Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan,
 616 Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao,
 617 Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding,
 618 Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang
 619 Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong,
 620 Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao,
 621 Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang,
 622 Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng
 623 Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen,
 624 R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhua Chen, Shengfeng Ye, Shiyu
 625 Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, and S. S. Li. DeepSeek-R1: Incentivizing
 626 reasoning capability in LLMs via reinforcement learning. *CoRR*, abs/2501.12948, 2025.

627

628 Haoxuan Deng, Bohao Zheng, Yirui Jiang, and Trung Hieu Tran. CAFA: Coding as auto-formulation
 629 can boost large language models in solving linear programming problem. In *The 4th Workshop on
 Mathematical Reasoning and AI at NeurIPS'24*, 2024. URL <https://openreview.net/forum?id=xC2xtBLmri>.

630

631 Zezhen Ding, Zhen Tan, Jiheng Zhang, and Tianlong Chen. Or-r1: Automating modeling and
 632 solving of operations research optimization problem via test-time reinforcement learning, 2025.
 633 URL <https://arxiv.org/abs/2511.09092>.

634

635 Jiaxuan Gao, Shusheng Xu, Wenjie Ye, Weilin Liu, Chuyi He, Wei Fu, Zhiyu Mei, Guangju Wang,
 636 and Yi Wu. On designing effective RL reward at training time for LLM reasoning. *CoRR*,
 637 abs/2410.15115, 2024.

638

639 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob
 640 Steinhardt. Measuring massive multitask language understanding. In *9th International Conference
 641 on Learning Representations*, 2021a.

642

643 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn
 644 Song, and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset.
 645 In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks
 646 Track (Round 2)*, 2021b.

647

648 Chenyu Huang, Zhengyang Tang, Shixi Hu, Ruqing Jiang, Xin Zheng, Dongdong Ge, Benyou
 649 Wang, and Zizhuo Wang. ORLM: A customizable framework in training large models for au-
 650 tomated optimization modeling. *Operations Research*, May 2025a. ISSN 1526-5463. doi:
 651 10.1287/opre.2024.1233.

648 Xuhan Huang, Qingning Shen, Yan Hu, Anningzhe Gao, and Benyou Wang. LLMs for mathematical
 649 modeling: Towards bridging the gap between natural and mathematical languages. In *Findings of*
 650 *the Association for Computational Linguistics 2025*, pp. 2678–2710, Albuquerque, New Mexico,
 651 2025b.

652 Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang,
 653 Bowen Yu, Kai Dang, An Yang, Rui Men, Fei Huang, Xingzhang Ren, Xuancheng Ren, Jingren
 654 Zhou, and Junyang Lin. Qwen2.5-Coder technical report. *CoRR*, abs/2409.12186, 2024.

655

656 Caigao Jiang, Xiang Shu, Hong Qian, Xingyu Lu, Jun Zhou, Aimin Zhou, and Yang Yu. LLMOPT:
 657 learning to define and solve general optimization problems from scratch. In *Advances in The*
 658 *Thirteenth International Conference on Learning Representations*, Singapore, 2025.

659

660 Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. Triviaqa: A large scale distantly
 661 supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meet-*
 662 *ing of the Association for Computational Linguistics*, 2017.

663

664 Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard H. Hovy. RACE: large-scale ReAd-
 665 ing comprehension dataset from examinations. In Martha Palmer, Rebecca Hwa, and Sebastian
 666 Riedel (eds.), *Proceedings of the 2017 Conference on Empirical Methods in Natural Language*
 667 *Processing*, 2017.

668

669 Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brah-
 670 man, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik,
 671 Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Chris Wilhelm,
 672 Luca Soldaini, Noah A. Smith, Yizhong Wang, Pradeep Dasigi, and Hannaneh Hajishirzi. TÜLU
 673 3: Pushing frontiers in open language model post-training. *CoRR*, abs/2411.15124, 2024.

674

675 Jinyuan Li, Yi Chu, Yiwen Sun, Mengchuan Zou, and Shaowei Cai. AutoPBO: LLM-powered
 676 optimization for local search PBO solvers. *CoRR*, abs/2509.04007, 2025a.

677

678 Kai Li, Fei Liu, Zhenkun Wang, Xialiang Tong, Xiongwei Han, Mingxuan Yuan, and Qingfu Zhang.
 679 Ars: Automatic routing solver with large language models. *CoRR*, abs/2502.15359, 2025b.

680

681 Xiaozhe Li, Jixuan Chen, Xinyu Fang, Shengyuan Ding, Haodong Duan, Qingwen Liu, and Kai
 682 Chen. OPT-BENCH: evaluating LLM agent on large-scale search spaces optimization problems.
 683 *CoRR*, abs/2506.10764, 2025c.

684

685 Kuo Liang, Yuhang Lu, Jianming Mao, Shuyi Sun, Chunwei Yang, Congcong Zeng, Xiao Jin,
 686 Hanzhang Qin, Ruihao Zhu, and Chung-Piaw Teo. Llm for large-scale optimization model auto-
 687 formulation: A lightweight few-shot learning approach, 2025. URL <https://dx.doi.org/10.2139/ssrn.5329027>.

688

689 Hongliang Lu, Zhonglin Xie, Yaoyu Wu, Can Ren, Yuxuan Chen, and Zaiwen Wen. OptMATH:
 690 A scalable bidirectional data synthesis framework for optimization modeling. In *Forty-second*
 691 *International Conference on Machine Learning*, 2025.

692

693 Zeyuan Ma, Hongshu Guo, Jiacheng Chen, Guojun Peng, Zhiguang Cao, Yining Ma, and Yue-Jiao
 694 Gong. LLaMoCo: Instruction tuning of large language models for optimization code generation.
 695 *CoRR*, abs/2403.01131, 2024.

696

697 OpenAI. GPT-5 system card. <https://cdn.openai.com/gpt-5-system-card.pdf>,
 698 2025.

699

700 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin,
 701 Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser
 702 Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan
 703 Leike, and Ryan Lowe. Training language models to follow instructions with human feedback.
 704 In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Infor-*
 705 *mation Processing Systems 2022*, New Orleans, LA, 2022.

706

707 Jiayi Pan, Xingyao Wang, Graham Neubig, Navdeep Jaithly, Heng Ji, Alane Suhr, and Yizhe Zhang.
 708 Training software engineering agents and verifiers with SWE-gym. In *ICLR 2025 Third Workshop*
 709 *on Deep Learning for Code*, 2025.

702 Rindranirina Ramamonjison, Timothy Yu, Raymond Li, Haley Li, Giuseppe Carenini, Bissan Ghad-
 703 dar, Shiqi He, Mahdi Mostajabdeh, Amin Banitalebi-Dehkordi, Zirui Zhou, and Yong Zhang.
 704 NL4Opt competition: Formulating optimization problems based on their natural language de-
 705 scriptions. In *Proceedings of the NeurIPS 2022 Competitions Track*, volume 220, pp. 189–203,
 706 2022.

707 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y. K. Li,
 708 Y. Wu, and Daya Guo. DeepSeekMath: Pushing the limits of mathematical reasoning in open
 709 language models. *CoRR*, abs/2402.03300, 2024.

710 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflex-
 711 ion: language agents with verbal reinforcement learning. In *Advances in Neural Information
 712 Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023*,
 713 New Orleans, LA, 2023.

714 Akash Singirikonda, Serdar Kadioglu, and Karthik Uppuluri. Text2Zinc: A cross-domain dataset
 715 for modeling optimization and satisfaction problems in MiniZinc. *CoRR*, abs/2503.10642, 2025.

716 Wen Song, Xinyang Chen, Qiqiang Li, and Zhiguang Cao. Flexible job-shop scheduling via graph
 717 neural network and deep reinforcement learning. *IEEE Transactions on Industrial Informatics*,
 718 19(2):1600–1610, 2023.

719 Yiwen Sun, Furong Ye, Xianyin Zhang, Shiyu Huang, Bingzhen Zhang, Ke Wei, and Shaowei
 720 Cai. AutoSAT: Automatically optimize SAT solvers via large language models. *CoRR*,
 721 abs/2402.10705, 2024.

722 Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang
 723 Sui. Math-Shepherd: Verify and reinforce LLMs step-by-step without human annotations. In *Pro-
 724 ceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume
 725 1: Long Papers)*, pp. 9426–9439, Bangkok, Thailand, 2024.

726 Yang Wu, Yifan Zhang, Yurong Wu, Yuran Wang, Junkai Zhang, and Jian Cheng. Training
 727 LLMs for optimization modeling via iterative data synthesis and structured validation. In Christos
 728 Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Findings
 729 of the Association for Computational Linguistics: EMNLP 2025*, pp. 12880–12896, Suzhou,
 730 China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-335-7.
 731 doi: 10.18653/v1/2025.findings-emnlp.691. URL [https://aclanthology.org/2025.
 732 findings-emnlp.691/](https://aclanthology.org/2025.findings-emnlp.691/).

733 Chaojun Xiao, Jie Cai, Weilin Zhao, Biyuan Lin, Guoyang Zeng, Jie Zhou, Zhi Zheng, Xu Han,
 734 Zhiyuan Liu, and Maosong Sun. Densing law of llms. *Nature Machine Intelligence*, 7(11):
 735 1823–1833, Nov 2025a. ISSN 2522-5839. doi: 10.1038/s42256-025-01137-0. URL <https://doi.org/10.1038/s42256-025-01137-0>.

736 Ziyang Xiao, Dongxiang Zhang, Yangjun Wu, Lilin Xu, Yuan Jessica Wang, Xiongwei Han, Xiaojin
 737 Fu, Tao Zhong, Jia Zeng, Mingli Song, and Gang Chen. Chain-of-Experts: When LLMs meet
 738 complex operations research problems. In *The Twelfth International Conference on Learning
 739 Representations*, Vienna, Austria, 2024.

740 Ziyang Xiao, Jingrong Xie, Lilin Xu, Shisi Guan, Jingyan Zhu, Xiongwei Han, Xiaojin Fu, Wing Yin
 741 Yu, Han Wu, Wei Shi, Qingcan Kang, Jiahui Duan, Tao Zhong, Mingxuan Yuan, Jia Zeng, Yuan
 742 Wang, Gang Chen, and Dongxiang Zhang. A survey of optimization modeling meets LLMs:
 743 Progress and future directions. In James Kwok (ed.), *Proceedings of the Thirty-Fourth Interna-
 744 tional Joint Conference on Artificial Intelligence, IJCAI-25*, pp. 10742–10750. International Joint
 745 Conferences on Artificial Intelligence Organization, 8 2025b. doi: 10.24963/ijcai.2025/1192.
 746 Survey Track.

747 Tian Xie, Zitian Gao, Qingnan Ren, Haoming Luo, Yuqian Hong, Bryan Dai, Joey Zhou, Kai Qiu,
 748 Zhirong Wu, and Chong Luo. Logic-RL: Unleashing LLM reasoning with rule-based reinforce-
 749 ment learning. *CoRR*, abs/2502.14768, 2025.

756 Huajian Xin, Z. Z. Ren, Junxiao Song, Zhihong Shao, Wanja Zhao, Haocheng Wang, Bo Liu,
 757 Liyue Zhang, Xuan Lu, Qiushi Du, Wenjun Gao, Haowei Zhang, Qihao Zhu, Dejian Yang, Zhibin
 758 Gou, Z. F. Wu, Fuli Luo, and Chong Ruan. DeepSeek-Prover-V1.5: Harnessing proof assistant
 759 feedback for reinforcement learning and monte-carlo tree search. In *Advances in the Thirteenth*
 760 *International Conference on Learning Representations*, Singapore, 2025.

761 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,
 762 Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin
 763 Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang,
 764 Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tingyu
 765 Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong
 766 Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report. *CoRR*, abs/2412.15115,
 767 2024a.

768 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang
 769 Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng
 770 Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang,
 771 Jian Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu,
 772 Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men,
 773 Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren,
 774 Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yingger Zhang,
 775 Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu.
 776 Qwen3 technical report. *CoRR*, abs/2505.09388, 2025a.

777 Zhicheng Yang, Yinya Huang, Wei Shi, Liang Feng, Linqi Song, Yiwei Wang, Xiaodan Liang, and
 778 Jing Tang. Benchmarking LLMs for optimization modeling and enhancing reasoning via reverse
 779 socratic synthesis. *CoRR*, abs/2407.09887, 2024b.

780 Zhicheng Yang, Yiwei Wang, Yinya Huang, Zhijiang Guo, Wei Shi, Xiongwei Han, Liang Feng,
 781 Linqi Song, Xiaodan Liang, and Jing Tang. OptiBench meets ReSocratic: Measure and improve
 782 LLMs for optimization modeling. In *The Thirteenth International Conference on Learning Rep-*
 783 *resentations*, Singapore, 2025b.

784 Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong
 785 Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi
 786 Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiaze Chen, Jiangjie Chen, Chengyi
 787 Wang, Hongli Yu, Weinan Dai, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu, Wei-
 788 Ying Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingxuan Wang. DAPO: an
 789 open-source LLM reinforcement learning system at scale. *CoRR*, abs/2503.14476, 2025.

790 Haotian Zhai, Connor Lawless, Ellen Vitercik, and Liu Leqi. Equivamap: Leveraging LLMs for
 791 automatic equivalence checking of optimization formulations. In *2nd AI for Math Workshop @*
 792 *ICML 2025*, 2025. URL <https://openreview.net/forum?id=RvdjzN1ksm>.

793 Kechi Zhang, Ge Li, Yihong Dong, Jingjing Xu, Jun Zhang, Jing Su, Yongfei Liu, and Zhi Jin.
 794 CodeDPO: Aligning code models with self generated and verified source code. In *Advances*
 795 *in the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long*
 796 *Papers)*, pp. 15854–15871, Vienna, Austria, 2025.

800 A DATASETS

801 A.1 THE INTRODUCTION OF EVALUATION DATASETS

802 In this section, we provide an overview of the datasets used for performance evaluation in our
 803 experiments. These datasets cover a wide range of optimization types and scenarios, ensuring the
 804 robustness and generalization of our proposed method. In our practice, we use the version of the
 805 benchmark datasets above from <https://github.com/antgroup/LLMOPT>.

806 **NL4Opt** (Ramamonjison et al., 2022) dataset is curated from the NL4Opt Competition. For this
 807 benchmark, we used the test split containing 230 annotated linear programming word problems

810
811
812
813
814
815
816
817
818
819

Table 3: Statistics of the optimization problem datasets

Dataset Name	# of Data
NL4Opt	230
ICML.C	410
Mamo.E	652
Mamo.C	211
NLP4LP	242
Com.OR	18
Indus.OR	100
OptiBench	605
OptMATH-Bench	166

820
821
822 after manually removing 15 unsolvable problems from the original 245 problems. Each problem is
823 sourced from domains such as sales, advertising, and investment, ensuring a balanced representation.

824 **Mamo** (Huang et al., 2025b) dataset (optimization split of the original Mamo dataset) consists
825 of two parts: Easy_LP and Complex_LP. These two subsets provide 652 high-school-level and
826 211 undergraduate-level Linear Programming (LP) and Mixed-Integer Linear Programming (MILP)
827 problems, respectively.

828 **IndustryOR** (Huang et al., 2025a) is the first industrial dataset specifically designed for optimization
829 modeling. It incorporates data from 13 different industries and covers a variety of real-world scenar-
830 os. The dataset includes real operations research problems from eight different industries, covering
831 five types of optimization problems, and divided into three difficulty levels. The test dataset contains
832 100 instances with optimal solutions.

833 **NLP4LP** (AhmadiTeshnizi et al., 2024) dataset includes 242 feasible samples sourced from optimi-
834 zation textbooks and lecture notes. These problems cover areas such as facility location, net-
835 work flow, scheduling, and portfolio management. Each instance in NLP4LP includes a description,
836 sample parameter data file, and optimal value derived from textbook solutions or manual solving,
837 offering a range of complex optimization challenges of varying difficulty levels.

838 **ComplexOR** (Xiao et al., 2024) dataset is developed in collaboration with three specialists in op-
839 erations research. It contains 18 samples sourced from diverse references such as academic papers,
840 textbooks, and real-world industrial scenarios. These problems encompass a broad spectrum of
841 topics, including supply chain optimization, scheduling problems, and warehousing optimization,
842 providing comprehensive and complex optimization challenges.

843 **OptMATH-Bench** (Lu et al., 2025) is a large-scale, challenging benchmark specifically designed
844 to evaluate the optimization modeling capabilities of LLMs, encompassing diverse optimization
845 problem types across 10+ real-world application domains such as logistics, manufacturing,
846 transportation, and finance. The benchmark features significantly more complex problem descriptions
847 with an average length 2.9× longer than Mamo Easy, containing extended natural language contexts
848 and intricate constraints that pose greater challenges.

849 **OptiBench** (Yang et al., 2025b) is a comprehensive benchmark for evaluating large language mod-
850 els' end-to-end optimization problem-solving capabilities. The dataset contains 605 carefully cu-
851 rated optimization problems that span multiple optimization types and formats. OptiBench includes
852 problems of Linear Programming (LP), Integer Programming (IP), and Mixed-Integer Linear Pro-
853 gramming (MILP), encompassing a wide range of optimization complexities.

854 **ICML Competition** (Yang et al., 2024b) dataset comprises data from the ICML 2024 Challenges on
855 Automated Math Reasoning - Track 3: Automated Optimization curated from the competition's test
856 split. Since the original ground truth is not released by the organizers, all solutions in this dataset
857 are manually labeled. The dataset serves as a challenging benchmark for evaluating end-to-end
858 optimization reasoning and problem-solving capabilities of language models.

859
860 **A.2 THE DISTRIBUTION OF OPTIMIZATION TYPES AND PROBLEM SCENARIOS OF**
861 **BENCHMARKS**

862
863 To evaluate the generalization ability of the MiniOpt across different problem scenarios through ex-
periments, this paper has counted the number of optimization types and scenarios in 9 benchmarks.

864
865 The distribution histogram of optimization types in the nine benchmarks used in this paper is shown
866 in Figure 3, and the distribution histogram of problem scenarios is shown in Figure 4.
867

Type	Com.OR	ICML.C	Indus.OR	Mamo.C	Mamo.E	NL4Opt	NLP4LP	OptMAT H-Bench	Optibench
CO	4	8	12	39	0	0	0	19	11
IP	4	211	23	11	2	161	167	7	291
LP	9	175	31	123	650	69	75	16	214
MILP	1	16	30	38	0	0	0	91	44
MOP	0	0	4	0	0	0	0	0	0
NLP	0	0	0	0	0	0	0	0	45
SOCP	0	0	0	0	0	0	0	33	0
# Data	18	410	100	211	652	230	242	166	605

877
878 Figure 3: Histogram showing the distribution of optimization types across 9 benchmarks. We cate-
879 gorize the problems in the benchmarks into these types: Combinatorial Optimization (CO), Integer
880 Programming (IP), Linear Programming (LP), Mixed-Integer Linear Programming (MILP), Multi-
881 Objective Optimization Problems (MOP), Nonlinear Programming (NLP), Second-Order Cone Pro-
882 gramming (SOCP).
883

Scenario	Com.OR	ICML.C	Indus.OR	Mamo.C	Mamo.E	NL4Opt	NLP4LP	OptMAT H-Bench	Optibench
Agriculture	0	37	5	4	31	13	12	0	46
Aviation	3	2	2	2	0	2	1	12	2
Construction	0	1	1	1	42	1	1	0	8
Education	0	5	1	0	30	0	0	0	7
Energy	0	0	0	5	25	1	2	7	15
Environment	0	0	0	0	31	0	1	0	0
Finance	0	16	7	5	85	5	5	6	22
Healthcare	0	15	1	5	31	23	29	0	20
Logistics	1	27	8	32	25	21	23	22	42
Manufacturing	5	169	37	12	34	77	83	56	254
Marketing	1	9	0	0	43	2	2	0	11
Military	0	0	2	0	36	0	0	0	3
Public Utilities	0	9	1	4	7	2	2	2	19
Resources	0	15	0	12	9	11	8	3	17
Retail	0	16	4	2	31	8	9	1	19
Services	4	38	10	6	73	29	29	7	40
Sports	0	0	0	0	31	0	0	0	0
Supply Chain	2	12	7	48	31	5	3	11	11
Technology	0	0	1	3	0	0	0	1	3
Telecommunications	1	1	0	8	27	0	0	9	1
Transportation	1	31	8	57	45	23	26	9	40
Other	0	7	5	5	15	7	6	20	25
# Data	18	410	100	211	682	230	242	166	605

901
902 Figure 4: Histogram showing the distribution of optimization problem scenarios across 9 bench-
903 marks. We categorize the problems in the benchmarks into these scenarios: Agriculture, Avia-
904 tion, Construction, Education, Energy, Environment, Finance, Healthcare, Logistics, Manufac-
905 turing, Marketing, Military, Public Utilities, Resources, Retail, Services, Supply Chain, Technol-
906 ogy, Telecommunications, Transportation, Other.

A.3 TRAINING DATASETS FOR SFT WARM-UP AND TWO-STAGE RL

907
908 The data volume and data sources for each stage of MiniOpt’s training process are illustrated in
909 Table 4.
910

912 Table 4: The number of samples in the dataset for training. Note that all data mentioned in the table
913 comes entirely from the training split of the corresponding dataset.

Training Stage	Dataset Size	Data Source
SFT Warm-up	140K	OptMATH-Train
RL-Stage 1	1585	NL4Opt (Train) & ICML.C (Train)
RL-Stage 2	3000	OptMATH-Train

918 **B BASELINES**919 **B.1 GENERAL MODELS**

920 **Qwen2.5-3B/7B/14B** (Yang et al., 2024a). The models of the Qwen2.5 series are widely adopted
 921 as base models or baselines. The series showcases significant enhancements such as substantially
 922 improved knowledge, coding, and mathematical capabilities. Key features excel at instruction fol-
 923 lowing, processing long contexts up to 128K tokens, and robustly handling structured data like
 924 JSON.

925 **DeepSeek-V3** (DeepSeek-AI et al., 2024). DeepSeek-V3 introduces a sparse Mixture-of-Experts
 926 (MoE) model with 671B total parameters. It achieves training efficiency through Multi-head Latent
 927 Attention (MLA) architecture and an auxiliary-loss-free load balancing strategy. Pretrained on 14.8T
 928 tokens with an SFT and reinforcement learning (RL).

929 **B.2 GENERAL REASONING MODELS**

930 **Qwen3-4B/8B/14B** (Yang et al., 2025a). Qwen3 pioneers a unified architecture (0.6B to 235B) inte-
 931 grating thinking mode (complex reasoning) and non-thinking mode (rapid responses) with dynamic
 932 switching. Its thinking budget mechanism enables adaptive computational allocation. The series
 933 outperforms larger MoE models in tasks such as coding, mathematics, and agent application.

934 **DeepSeek-R1** (DeepSeek-AI et al., 2025). The DeepSeek-R1 is an enhanced model based on
 935 DeepSeek-R1-Zero presented in this work. It's a purely RL-driven reasoning model requiring no
 936 SFT pretraining. As the most representative model with thinking ability, DeepSeek-R1 is an impor-
 937 tant baseline for reasoning models.

938 **Gemini-2.5-Pro** (Comanici et al., 2025). Gemini-2.5-Pro is a powerful multimodal agent that has
 939 excellent programming / reasoning performance and enables the processing of long video content.
 940 The Gemini-2.5 family spans the full Pareto front of capability-cost optimization. Its integration of
 941 long-context understanding, multimodality, and reasoning unlocks novel agentic applications.

942 **GPT-5** (OpenAI, 2025). GPT-5 is the latest unified, router-mediated system that instantiates a spec-
 943 trum of language-model instances ranging from a high-throughput, low-latency model (gpt-5-main)
 944 to a deliberative, compute-intensive reasoning model (gpt-5-thinking). The router selects the appro-
 945 priate instantiation by conditioning on conversation type, task complexity, tool requirements, and
 946 explicit user directives, thereby optimizing both instruction adherence and inference efficiency.

947 **B.3 PROMPT-BASED METHODS**

948 **Reflection** (Shinn et al., 2023). Reflexion is an enhanced language agent framework utilizing feed-
 949 back mechanisms. It enables agents to excel at sequential decision-making tasks through task feed-
 950 back analysis and memory buffering without requiring weight updates. This framework accommo-
 951 dates diverse feedback signals and demonstrates effectiveness across programming, math problems
 952 and language reasoning domains.

953 **OptiMUS** (AhmadiTeshnizi et al., 2024). OptiMUS is a highly modular solver that leverages the
 954 text understanding and generating capabilities of LLMs. It constructs specialized agents for entity
 955 extraction, mathematical modeling, and code generation using concise prompts, while incorporating
 956 a reflection mechanism for iterative improvement.

957 **Chain-of-Experts** (Xiao et al., 2024). Chain-of-Experts is a multi-agent framework specifically
 958 designed for operations research optimization problems. The system features a central controller
 959 that coordinates an interaction sequence among specialized agents, including a term interpreter,
 960 modeling agent, and programming expert. Thus solving optimization problems through precise
 961 coordination of multiple modules.

962 **B.4 LEARNING-BASED MODELS**

963 **LLMOPT-14B** (Jiang et al., 2025). LLMOPT is a novel framework for optimization problem solv-
 964 ing that leverages LLMs. It begins by formulating a unified representation of optimization prob-

972 lems, thereby enhancing the model’s ability to generalize across diverse types of scenarios. Based
 973 on this unified description of five-element formulation, the framework generates the solving code.
 974 LLMOPT uses multi-instruction SFT and KTO alignment during training to enhance modeling ac-
 975 curacy and reduce model hallucinations.

976 **OptMATH-Qwen2.5-7B** (Lu et al., 2025). OptMATH-Qwen2.5-7B is trained end-to-end on the
 977 OptMATH-Train dataset, it generates both mathematical formulations and solver code from problem
 978 descriptions. The input consists of textual problem specifications, while the target output comprises
 979 concatenated sequences. Optimization follows the standard sequence-to-sequence loss function,
 980 enabling single-stage joint optimization of formulation and code generation.

981 **Step-OPT-8B** (Wu et al., 2025). Step-OPT is a model trained on a meticulously curated high-quality
 982 dataset. The training set is enhanced through Scope-Evolve and Complexity-Evolve techniques,
 983 which improve both the difficulty level and the coverage of application scenarios. Moreover, a
 984 multi-agent and stepwise verification mechanism is employed to enhance the quality of problems
 985 and solutions while eliminating errors present in the original dataset. Finally, supervised fine-tuning
 986 was performed on this high-quality training dataset resulting in the Step-OPT-8B model.

988 C THE DETAILS OF THE FIVE-ELEMENT FORMULATION

990 The five-element modeling formulation is a universal mathematical model for optimization prob-
 991 lems, which consists of five parts. We start from the following formulation:

$$994 \min_{\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^D} f(\mathbf{x}), \quad \text{s.t. } G(\mathbf{x}) \leq \mathbf{c}, \quad (5)$$

996 where $\mathbf{x} = (x_1, x_2, \dots, x_D)^\top$ is the D -dimensional decision variable, \mathcal{X} is the feasible region,
 997 $f : \mathcal{X} \rightarrow \mathbb{R}$ is the objective, $G(\mathbf{x}) : \mathbb{R}^D \rightarrow \mathbb{R}^m$ collects the constraints, and $\mathbf{c} \in \mathbb{R}^m$ provides
 998 the upper bounds. Among them, **Variables**, **Objective**, and **Constraints** correspond to \mathbf{x} , $f(\mathbf{x})$,
 999 and $G(\mathbf{x})$, while **Sets** and **Parameters** provide indices and numerical tables that instantiate and
 1000 vectorize f and G . Sets determine the dimensions and naming of decision and constraint families.
 1001 Parameters supply exogenous constants such as costs, coefficients, budgets, and demands. Vari-
 1002 ables specify domains and bounds(e.g., continuous, nonnegative, integer, or binary), which jointly
 1003 define the feasible region \mathcal{X} ; domain-type restrictions such as “positive integers” may equivalently
 1004 be encoded as explicit constraints, and our parser maps both styles to \mathcal{X} . Objective gives the mini-
 1005 mization or maximization expression, and Constraints provide named families of linear or nonlinear
 1006 equalities/inequalities composing G and the bound vector \mathbf{c} . This representation naturally spans
 1007 LP/IP/MILP: integrality arises through \mathcal{X} , linearity or nonlinearity is captured by the form of G ,
 1008 and multi-objective problems can be accommodated by extending $f(\mathbf{x})$ to a vector $F(\mathbf{x})$ with a
 1009 scalarization scheme. The think output ends with the five-element formulation, which serves as a
 1010 modeling blueprint.

1011 D DETAILS OF GRPO ALGORITHM

1013 This paper proposes OptGRPO based on the improvement of the GRPO algorithm (Shao et al.,
 1014 2024). For each query, GRPO sample G responses, compute the mean and standard deviation of
 1015 their scalar rewards, and form a group-normalized advantage $\hat{A}_i = \frac{(r_i - \mu)}{\sigma}$. Specifically, GRPO
 1016 optimizes the policy π_θ as follows:

$$1019 \mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{q \sim P(Q), \{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot|q)} \left[\sum_{i=1}^G \sum_{t=1}^{|o_i|} \frac{1}{|o_i|} \left(\min \left[\frac{\pi_\theta(o_{i,t}|q, o_{i,< t})}{\pi_{\theta_{\text{old}}}(o_{i,t}|q, o_{i,< t})} \hat{A}_{i,t}, \text{clip} \left(\frac{\pi_\theta(o_{i,t}|q, o_{i,< t})}{\pi_{\theta_{\text{old}}}(o_{i,t}|q, o_{i,< t})}, 1 - \varepsilon, 1 + \varepsilon \right) \hat{A}_{i,t} \right] \right. \right. \\ 1020 \left. \left. - \beta D_{\text{KL}}[\pi_\theta(\cdot|q) \parallel \pi_{\text{ref}}(\cdot|q)] \right) \right], \quad (6)$$

1026 where q denotes queries sampled from the input dataset and o denotes the model’s outputs. ε is the
 1027 clipping threshold, β is the coefficient on the KL penalty, and D_{KL} is the KL divergence between
 1028 the current policy π_θ and the reference policy π_{ref} .
 1029

1030 E THE PROCESSING PIPELINE OF TRAINING DATA 1031

1032 The processing pipeline for the SFT training set is detailed in Section 3.3.1. The resulting collection
 1033 of this pipeline is named OptMATH-Train-Pyomo, which contains approximately 140K samples.
 1034 The prompt template used in code conversion is introduced in Appendix L, and the prompt of the
 1035 solver adapter is introduced in Appendix M.
 1036

1037 In the first stage of MiniOpt’s RL training, we employ a set of relatively easy and small-scale optimi-
 1038 zation problems, for which the accuracy score is more readily maximized. This, in turn, incen-
 1039 tivizes the model to attain higher format score and five-element score, thereby accelerating mastery
 1040 of the reasoning to model and solve paradigm. Concretely, the dataset of the first stage is the union
 1041 of the NL4Opt (AhmadiTeshnizi et al., 2024) and ICML Competition (Yang et al., 2024b) training
 1042 splits, comprising 1585 problems. Each instance is presented as a natural-language prompt paired
 1043 with a reference answer and is fully compatible with our pipeline of parsing, executing, and scoring,
 1044 enabling straightforward computation of the OptReward.
 1045

1046 The second stage of RL targets optimization generalization (Jiang et al., 2025) under a limited train-
 1047 ing budget. The objective is to construct a training set that simultaneously covers diverse optimiza-
 1048 tion types and scenarios, while preserving the scenario proportions observed in the real distribution.
 1049 Starting from the OptMATH-Train pool containing 201K problems, we label each instance with
 1050 types and scenarios using the DeepSeek-V3 (DeepSeek-AI et al., 2024), with prompt templates
 1051 and data distributions provided in Appendix O. We then sample a 3000-instance stage-2 training set
 1052 subject to two constraints: (i) type-uniform coverage, with exactly 600 instances per type; and (ii)
 1053 within each type, the scenario frequencies match the distribution in the full pool. Formally, letting
 1054 $p(s)$ denote the overall scenario distribution in the complete pool, the target number of samples for
 1055 each scenario s under type t is allocated as $n_t(s) = \text{round}(600 \cdot p(s))$, where $\sum_s n_t(s) = 600$.
 1056

1057 To prevent leakage, problems used during the SFT warm-up (Section 3.3.1) are excluded from con-
 1058 sideration. The resulting 3K training set is uniform across types and aligned with the scenario dis-
 1059 tribution of the data pool. Compared with random or non-selected baselines, this selection improves
 1060 overall performance of MiniOpt under the same compute budget.
 1061

1059 F COMPARISON OF EXECUTION RATE ACROSS 9 BENCHMARKS 1060

1061 This section we use Execution Rate (ER), the proportion of generated code samples that run suc-
 1062 cessfully without errors. As shown in the Table 5, MiniOpt also exhibits superior text generation ca-
 1063 pabilities compared to baseline methods, which suggests its excellent code generation performance
 1064 given the problem modeling.
 1065

1066 G ABLATION STUDY OF MINIOPt ACROSS 9 BENCHMARKS 1067

1068 **Importance of the Training Pipeline of MiniOpt (Answer to Q3).** In the training pipeline of
 1069 MiniOpt, each module plays a distinct role in improving SA and ER. First, the lightweight SFT
 1070 warm-up provides a better starting point for RL training. Without it, averages for SA and ER fall
 1071 to 53.59% and 83.66%, respectively. The decreases are $\Delta \text{SA} = -3.35$ and $\Delta \text{ER} = -4.38$. Second, re-
 1072 moving all RL training collapses the average performance from SA of 56.94% and ER of 88.04%
 1073 to 39.25% and 75.36%, respectively. With the sharpest drops on challenging datasets like Com.OR
 1074 ($\Delta \text{SA} = -33.33\%$, $\Delta \text{ER} = -33.33\%$) and OptMATH-Bench ($\Delta \text{SA} = -22.29\%$, $\Delta \text{ER} = -34.33\%$)
 1075 where accurate modeling and solver selection are indispensable. Third, collapsing the two-stage
 1076 RL removes the progressive training that first consolidates the paradigm (stage-1) and then targets
 1077 generalization (stage-2), averages drop to the SA of 54.12% and ER of 84.38% ($\Delta \text{SA} = -2.82\%$,
 1078 $\Delta \text{ER} = -3.66\%$). Moreover, the data selection of stage-2 is crucial for sample efficiency that train-
 1079 ing on the full pool (w/o Data Selection) yield the average SA of 53.58% and the average ER of
 82.04% ($\Delta \text{SA} = -3.36\%$, $\Delta \text{ER} = -6.00\%$), while the random selected training data (w/ Random

1080
1081
1082
1083
1084Table 5: Comparison of the ER metric across 9 benchmarks with rankings (NL4Opt, ICML Competition, Mamo Easy, Mamo Complex, NLP4LP, ComplexOR, IndustryOR, OptiBench, OptMATH-Bench). **Bold** indicates 1st, wavy underline indicates 2nd, underline indicates 3rd. “Rank*” represents the result of sorting methods among parameter sizes below 10B.1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099

Category	Model / Method	Avg.	Rank	Rank*	NL4Opt	ICML.C	Mamo.E	Mamo.C	NLP4LP	Com.OR	Indus.OR	OptiBench	OptMATH-Bench
General Models	Qwen2.5-3B-Instruct	17.11	21	10	31.30	28.54	21.47	7.11	28.93	5.56	10.00	20.50	0.60
	Qwen2.5-7B-Instruct	41.69	18	8	66.09	68.54	40.34	19.43	71.90	16.67	30.00	53.22	9.04
	Qwen2.5-14B-Instruct	63.53	14	-	82.17	80.73	86.66	59.24	84.71	38.89	52.00	61.49	25.90
	DeepSeek-V3	83.50	7	-	97.83	<u>97.32</u>	96.63	71.56	97.93	72.22	70.00	84.79	63.25
(Thinking)	Qwen3-4B	14.15	22	11	19.57	20.00	15.18	18.48	18.18	11.11	5.00	16.20	3.61
	Qwen3-8B	25.55	20	9	36.96	38.29	26.53	18.01	39.26	22.22	10.00	35.04	3.61
	Qwen3-14B	31.17	19	-	28.70	25.12	40.18	30.81	22.73	61.11	27.00	27.44	17.47
	DeepSeek-R1	83.07	9	-	96.09	94.15	89.72	84.83	91.32	61.11	<u>82.00</u>	88.92	53.01
	Gemini-2.5-Pro	89.65	4	-	94.35	96.10	95.86	86.26	96.28	77.78	87.00	91.90	81.32
Prompt-based Methods	GPT-5	83.26	8	-	98.08	<u>97.32</u>	71.47	55.92	97.52	88.89	78.00	<u>92.07</u>	69.88
	Chain-of-Experts	61.72	15	-	79.57	73.41	72.85	56.87	77.27	55.56	54.00	65.45	20.48
	OptiMUS	52.13	17	-	45.22	75.85	73.77	43.60	44.63	44.44	48.00	65.95	27.71
Learning-based Models	Reflexion	80.42	10	-	91.74	91.46	97.55	67.77	95.45	83.33	66.00	81.65	48.80
	Step-OPT-LLaMA3.2-3B	67.44	13	6	98.70	48.05	90.64	73.93	99.17	66.67	60.00	48.76	21.08
	Step-OPT-LLaMA3-8B	78.29	11	4	96.96	93.90	98.93	<u>90.05</u>	94.63	61.11	67.00	73.72	28.31
	Step-OPT-Qwen2.5-3B	59.37	16	7	54.78	58.78	97.09	42.18	81.40	61.11	50.00	67.27	21.69
	Step-OPT-Qwen2.5-7B	71.76	12	5	94.78	77.56	87.73	84.83	68.18	72.22	64.00	71.24	25.30
	OptMATH-7B	85.07	6	3	99.13	95.85	98.47	90.05	99.17	66.67	69.00	82.81	64.46
	LLMOPT-14B	90.03	3	-	97.42	93.90	92.29	77.73	97.93	88.89	61.00	73.22	31.93
Ours	MiniOpt-3B	88.04	5	<u>2</u>	99.57	95.85	98.47	87.68	<u>99.59</u>	88.89	70.00	83.64	68.67
	MiniOpt-7B	90.61	2	<u>1</u>	99.57	98.05	98.93	<u>95.26</u>	100.00	88.89	74.00	84.30	76.51
	MiniOpt-14B	92.35	1	-	99.57	98.54	98.93	<u>97.16</u>	100.00	88.89	77.00	<u>89.09</u>	81.93

1100
1101
1102

Table 6: Ablation study (MiniOpt-3B) on the ER metric across 9 benchmarks.

1103
1104
1105
1106
1107
1108
1109
1110

Category	Model / Method	Avg.	NL4Opt	ICML.C	Mamo.E	Mamo.C	NLP4LP	Com.OR	Indus.OR	OptiBench	OptMATH-Bench
Ablations	MiniOpt-3B	88.04	99.57	95.85	98.47	87.68	99.59	88.89	70.00	83.64	68.67
	MiniOpt-3B w/o SFT Warm-up	83.66	99.13	96.59	98.16	84.36	98.35	66.67	68.00	84.46	57.23
	MiniOpt-3B w/o RL	75.36	96.09	92.68	97.09	69.19	97.93	55.56	62.00	73.39	34.34
	MiniOpt-3B w/o Two-stage RL	84.38	98.26	96.10	97.85	84.36	98.35	72.22	70.00	84.46	57.83
	MiniOpt-3B w/o Data Selection	82.04	99.57	95.61	96.78	86.73	98.76	55.56	66.00	82.15	57.23
	MiniOpt-3B w/ Random Selection	82.01	99.13	96.10	97.09	77.73	97.93	67.67	69.00	84.63	48.80
	MiniOpt-3B w/ OptReward	83.39	98.26	95.61	97.70	83.89	97.93	72.22	64.00	82.48	58.43
	MiniOpt-3B w/ GRPO	80.00	97.39	94.63	96.32	83.41	95.87	55.56	64.00	81.65	51.20
	MiniOPT-3b w/ DAPO	86.82	99.13	96.34	98.77	<u>85.31</u>	<u>99.17</u>	<u>83.33</u>	74.00	86.94	<u>58.43</u>

1111
1112
1113
1114
1115
1116
1117
1118
1119
1120

Selection) yield SA of 50.94% and ER of 82.01% on average (Δ SA = -6.00%, Δ ER = -6.03%), indicating that type-uniform, globally scenario-aligned sampling concentrates updates where they best improve cross-type, cross-scenario behavior. Finally, reverting our OptGRPO to the original GRPO (w/o OptGRPO) further decreases both metrics, averages SA of 48.33% and ER of 80.00%. The drop aligns with our algorithmic choices: removing KL part frees exploration for small-scale LLMs; Clip-Higher part prevents entropy collapse by allowing probability increases on rare but crucial reasoning/code tokens; Token-Level Loss part can enhance the impact of long output, which is conducive to the training of the reasoning model. Together these changes improve sample efficiency and training stability, which is critical for eliciting strong optimization generalization at small parameter scales.

1121
1122
1123
1124

H COMPARISON OF AVERAGE ER AGAINST MODEL PARAMETER SCALES FOR VARIOUS METHODS

1125
1126
1127

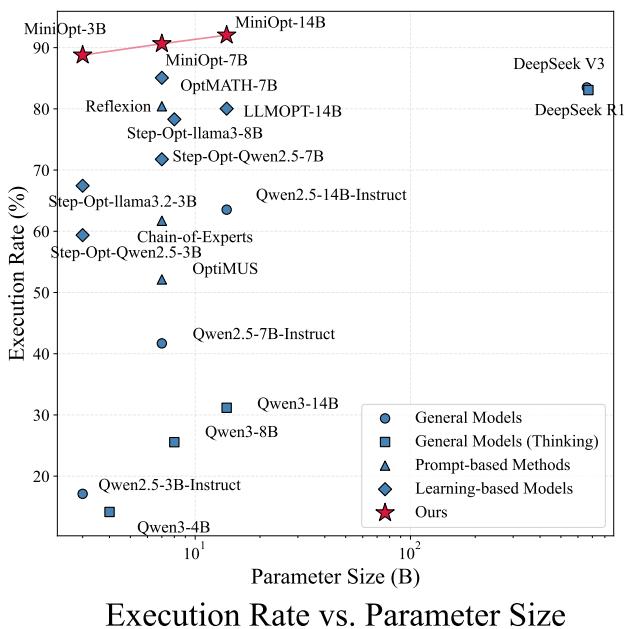
In this chapter, we present a comparative plot of the model parameters scale and the average ER across multiple methods, as illustrated in Figure 5.

1128
1129
1130
1131
1132

I COMPARISON OF RESULTS BETWEEN THE MINIMUM MODEL AND MODELS OF OTHER SCALES

1133

As shown in Tables 7 and 8, we compared the SA and ER of four different sizes of MiniOpt (1.5B, 3B, 7B, and 14B) across 9 benchmarks, respectively.



Execution Rate vs. Parameter Size

Figure 5: Comparison of average ER against model parameter scales for various methods. MiniOpt is the Pareto optimal among compared methods.

Table 7: Comparison of the SA metric between MiniOpt-1.5B and larger scale counterparts across 9 benchmarks.

Solving Accuracy (SA)	Avg.	NL4Opt	ICML-C	Mamo.E	Mamo.C	NLP4LP	Com.OR	Indus.OR	OptiBench	OptiMATH-Bench
Dataset size		230	410	652	211	242	18	100	605	166
Ours	MiniOpt-1.5B	46.15	63.48	55.37	77.15	27.01	58.68	38.89	16.00	41.98
	MiniOpt-3B	56.94	83.04	68.05	85.43	35.07	73.55	50.00	21.00	53.55
	MiniOpt-7B	62.76	89.13	77.56	88.34	38.89	79.34	55.56	26.00	59.34
	MiniOpt-14B	66.10	92.17	86.34	90.80	33.65	79.75	61.11	27.00	67.44

J THE SEESAW ISSUE OF LLMs

We assess whether adapting our models to optimization modeling introduces a seesaw effect for small-parameter models. After evaluating pre-training and post-training scores on six widely used general-purpose benchmarks: MMLU (Hendrycks et al., 2021a), MATH (Hendrycks et al., 2021b), HumanEval (Chen et al., 2021), TriviaQA (Joshi et al., 2017), RACE (Lai et al., 2017), GSM8K (Cobbe et al., 2021), we observe that the average scores of MiniOpt-7B and MiniOpt-3B decrease by only 1.83% and 3.05%, respectively. Notably, the 3B model even exhibits a 0.7% improvement on GSM8K, a proxy for mathematical reasoning. These results indicate that our training paradigm for optimization modeling does not induce a pronounced seesaw effect. We attribute this to two factors. First, the two-stage RL framework coupled with a verifiable OptReward constrains learning to structural modeling correctness and executable solving efficacy, mitigating overfitting to superficial linguistic style or lengthy chain-of-thought and thereby substantially reducing the risks of catastrophic forgetting and cross-task seesaw effects. Second, the foundational competencies required for optimization modeling, such as mathematical understanding, symbolic reasoning, program synthesis, and execution, which highly overlap with those assessed by general benchmarks such as MMLU, MATH, GSM8K, and HumanEval. Consequently, targeted reinforcement in this domain does not overwrite existing representations, instead, it yields small positive transfer on tasks closely aligned with modeling and solving, as exemplified by GSM8K.

1188

1189
1190
Table 8: Comparison of the ER metric between MiniOpt-1.5B and larger scale counterparts across
9 benchmarks.

Execution Rate (ER)	Avg.	NL4Opt	ICML.C	Mamo.E	Mamo.C	NLP4LP	Com.OR	Indus.OR	OptiBench	OptMATH-Bench
Dataset size		230	410	652	211	242	18	100	605	166

	MiniOpt-1.5B	80.00	93.91	91.71	97.09	79.15	91.74	66.67	63.00	75.87	60.84
Ours	MiniOpt-3B	88.04	99.57	95.85	98.47	87.68	99.59	88.89	70.00	83.64	68.67
	MiniOpt-7B	90.61	99.57	98.05	98.93	95.26	100.00	88.89	74.00	84.30	76.51
	MiniOpt-14B	92.35	99.57	98.54	98.93	97.16	100.00	88.89	77.00	89.09	81.93

1196

1197

1198
1199
**K DETAILED DISCUSSION ON THE EFFICIENCY OF MINIOPT TO MODEL
1200 AND SOLVE OPTIMIZATION PROBLEMS**

1201

1202 To validate the high efficiency of the MiniOpt model during inference, we compared the average
1203 number of tokens generated in full responses on the OptMATH dataset by two small-parameter
1204 models of MiniOpt and DeepSeek-R1. As shown in Table 9, MiniOpt achieves higher ER and SA
1205 while using fewer average output tokens. This advantage stems from MiniOpt’s internalized unified
1206 reasoning to model and solve paradigm tailored for optimization modeling problems, which guides
1207 the model along a more efficient reasoning path during inference.

1208

1209

1210
Table 9: Comparison of average output token consumption in OptMATH-Bench.

Model	Avg. Token Count	SA (%)	ER (%)
MiniOpt-3B	1911.72	42.77	68.67
MiniOpt-7B	3068.99	51.20	76.51
DeepSeek-R1	5078.90	39.76	53.01

1215

1216

1217
L SYSTEM PROMPT FOR CODE CONVERSION FROM GUROBIPY TO PYOMO

1218

1219

1220 This section provides the system prompt for the large language model to convert GurobiPy code in
1221 OptMATH-Train into Pyomo code, where the content within `[·]` and `{·}` will be replaced with the
1222 corresponding parts.

1223

1224

1225
PROMPT TEMPLATE FOR CODE CONVERSION

1226

1227

You are an expert in optimization problems. Your task is to convert
the given gurobipy code into pyomo code.

1228

1229

****Instructions:****

1230

1. Don’t give any explanation, just provide the converted pyomo
code in the following format:

1231

`'''python
[pyomo code here]`

1232

2. Please note that the following solvers are available for use: ‘
glpk’, ‘cbc’, ‘ipopt’, ‘scip’. Other solvers should not be
utilized.

1233

3. Please add ‘from pyomo.environ import *’ at the beginning of
your code.

1234

4. Please print the optimal objective value at the end of the code.

1235

****Gurobipy code:****

1236

{gurobipy}

1237

1238

1239

1240

1241

1242
1243

M PROMPT FOR THE SOLVER SELECTION

1244
1245
1246

This section provides the system prompt to guide MiniOpt models in autonomously selecting solvers after modeling optimization problems.

1247

PROMPT FOR SOLVER SELECTION

1248

1249

Solver Selection Guide:

1250

- ``glpk``: Best for small-to-medium linear problems (LP).
- ``cbc``: Recommended for mixed-integer linear programming (MILP) and larger linear problems. Handles binary/integer variables well.
- ``ipopt``: Use for nonlinear problems (NLP) with continuous variables. Does NOT support discrete variables.
- ``scip``: Most versatile - handles mixed-integer nonlinear problems (MINLP), large-scale problems, and complex constraints.

1257

Select solver based on:

1258

1. Variable types (continuous vs integer/binary)
2. Linearity of objective/constraints
3. Problem scale (small: glpk/cbc, large: scip/ipopt)
4. Nonlinearity presence (use ipopt/scip)

1262

1263

N SYSTEM PROMPT FOR RL TRAINING

1264

This section provides the system prompt used by MiniOpt models during reinforcement learning (RL) training.

1265
1266
1267
1268

1269

SYSTEM PROMPT FOR RL TRAINING

1270

1271

You are a helpful assistant. The assistant first thinks about the reasoning process in the mind and then provides the user with the answer. The reasoning process and answer are enclosed within `<think> </think>` and `<answer> </answer>` tags, respectively, i.e., `<think>` reasoning process here `</think><answer>` answer here `</answer>`, please make sure to answer according to the above format. Now the user asks you to solve an optimization reasoning problem, you should:

1278

1. Detailed reasoning about the problem within `<think> </think>` tags.
2. Write the corresponding five-element model (derived from your analysis).
3. Determine the mathematical properties of problem and select an appropriate solver from 'glpk', 'cbc', 'ipopt', 'scip'.
4. Recheck and correct if necessary at the end of the `<think> </think>` section.
 - Verify the five-element model fully captures the problem's requirements.
 - Confirm no constraints/variables are missing or over-simplified.
 - Ensure the solver choice aligns with the problem's mathematical properties.
5. Provide the corresponding Pyomo code based on checked five-element model within `<answer> </answer>` tags.

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295

In mathematics, optimization problem can be modeled as the following expression $\min_{\mathbf{x}} \{f(\mathbf{x})\} \text{ s.t. } g(\mathbf{x}) \leq \mathbf{c}$, where $\mathbf{x} = (x_1, x_2, \dots, x_n)$

```

1296
1297     \ldots,  $x_D$ )^\top is the  $D$ -dimensional decision variable, $  

1298     \mathcal{X} \subset \mathbb{R}^D$ is the feasible domain,  

1299      $f: \mathcal{X} \rightarrow \mathbb{R}$  is the objective  

1300     function and the goal is to find the minima of  $f$ ,  $G(\mathbf{x})$  are the constraints of  

1301      $\mathbf{c}(\mathbf{x})$ .  

1302
1303 The above definition can be mapped to a five-element consisting of  

1304     'Variables, Objective, Constraints, Sets, Parameters'.  

1305     Variables indicate what  $\mathbf{x}$  is, Objective  

1306     describes the form of the objective function  $f(\mathbf{x})$ , and Constraints indicates the constraints  $G(\mathbf{x})$  and  $\mathbf{c}(\mathbf{x})$ . These three can abstract the  

1307     optimization problem. Sets and Parameters are their specific  

1308     explanations: Sets describe and explain the subscripts of the  

1309     vectors or matrices in them, and Parameters supplement their  

1310     specific values.  

1311
1312 You need to give a detailed reasoning process for the problem first  

1313     , and then write the corresponding five-element model based on  

1314     the problem description and information provided by user.  

1315
1316 Please complete the following template to model the optimization  

1317     problem into five-element:  

1318
1319 <think>  

1320 Your reasoning process here...  

1321
1322 ## Sets:  

1323 [You need to fill in]  

1324
1325 ## Parameters:  

1326 [You need to fill in]  

1327
1328 ## Variables:  

1329 [You need to fill in]  

1330
1331 ## Objective:  

1332 [You need to fill in]  

1333
1334 ## Constraints:  

1335 [You need to fill in]  

1336 </think>  

1337
1338 In Pyomo, all constraints must be formulated using ' $\leq$ ', ' $\geq$ ', or  

1339     ' $=$ '. If you need to use ' $>$ ' or ' $<$ ', you can introduce a very  

1340     small value to transform the inequality. Please note that the  

1341     following solvers are available for use: 'glpk', 'cbc', 'ipopt',  

1342     'scip'. Other solvers should not be utilized.  

1343
1344 **Solver Selection Guide:**  

1345 - 'glpk': Best for small-to-medium linear problems (LP).  

1346 - 'cbc': Recommended for mixed-integer linear programming (MILP)  

1347     and larger linear problems. Handles binary/integer variables  

1348     well.  

1349 - 'ipopt': Use for nonlinear problems (NLP) with continuous  

1350     variables. Does NOT support discrete variables.  

1351 - 'scip': Most versatile - handles mixed-integer nonlinear  

1352     problems (MINLP), large-scale problems, and complex constraints.  

1353
1354 **Select solver based on:**  

1355 1. Variable types (continuous vs integer/binary)  

1356 2. Linearity of objective/constraints

```

```

1350
1351     3. Problem scale (small: glpk/cbc, large: scip/ipopt)
1352     4. Nonlinearity presence (use ipopt/scip)
1353
1354     Please select an appropriate solver based on the type and quantity
1355     of variables, objectives, and constraints. After thinking, when
1356     you finally reach the five-element model, you should give the
1357     corresponding Pyomo code within the <answer> </answer> tags, i.e
1358     ., <answer> '''python\n code here''' </answer>. The user will
1359     extract the complete code you provide through the regular
1360     expression r'''python\n(.*)?''' in the <answer> </answer> tags
1361     . The execution result of the code should include the optimal
1362     solution and the objective value. The optimal objective value
1363     will be extracted automatically from your last printed result.
1364
1365
1366
1367
1368
1369
1370

```

O LABELING PROMPT AND THE DATA DISTRIBUTIONS OF OPTMATH-TRAIN

This section provides the system prompts for the large language model to label the type and scenario of problems in OptMATH-Train, where `{.}` will be replaced with the corresponding content. After labeling, the distribution of scenarios in OptMATH-Train is displayed in Figure 6, and the distribution of types in OptMATH-Train is displayed in Figure 7

SYSTEM PROMPT FOR TYPE LABELING

```

1373     Please classify the following optimization problem into one of
1374     these technical types based on the mathematical formulation and
1375     decision variables, not just surface-level descriptions:
1376
1377     1. Linear Programming (LP): Problems with linear objective function
1378        and linear constraints, all continuous variables
1379     2. Integer Programming (IP): Problems with linear or nonlinear
1380        components where ALL variables are discrete/integer
1381     3. Mixed Integer Linear Programming (MILP): Problems with linear
1382        components containing BOTH continuous and discrete variables
1383     4. Nonlinear Programming (NLP): Problems with nonlinear objective
1384        function and/or nonlinear constraints (variables may be
1385        continuous/discrete)
1386     5. Combinatorial Optimization (CO): Problems focused on selecting/
1387        discrete structures (graphs, permutations, sets) with typically
1388        binary variables
1389     6. Multi-objective Programming (MOP): Problems explicitly
1390        optimizing multiple conflicting objectives simultaneously
1391     7. Second-Order Cone Programming (SOCP): Problems with a linear
1392        objective function, linear constraints, and second-order cone
1393        constraints (e.g.,  $\|Ax + b\| \leq c^T x + d$ )
1394
1395     # Problem:
1396     {{Question}}
1397
1398     # Output
1399     Analyze the mathematical structure step by step and classify its
1400     type. Finally, output the type abbreviation in the following
1401     format:
1402     Type: Abbreviation of the type
1403
1404     Note:
1405     - Focus on the fundamental mathematical formulation, not
1406       application domain
1407     - Check variable types (continuous/discrete/binary) and objective/
1408       constraint linearity
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267
2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321
2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2339
2340
2341
2342
2343
2344
2345
2346
2347
2348
2349
2350
2351
2352
2353
2354
2355
2356
2357
2358
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375
2376
2377
2378
2379
2380
2381
2382
2383
2384
2385
2386
2387
2388
2389
2390
2391
2392
2393
2394
2395
2396
2397
2398
2399
2400
2401
2402
2403
2404
2405
2406
2407
2408
2409
2410
2411
2412
2413
2414
2415
2416
2417
2418
2419
2420
2421
2422
2423
2424
2425
2426
2427
2428
2429
2430
2431
2432
2433
2434
2435
2436
2437
2438
2439
2440
2441
2442
2443
2444
2445
2446
2447
2448
2449
2450
2451
2452
2453
2454
2455
2456
2457
2458
2459
2460
2461
2462
2463
2464
2465
2466
2467
2468
2469
2470
2471
2472
2473
2474
2475
2476
2477
2478
2479
2480
2481
2482
2483
2484
2485
2486
2487
2488
2489
2490
2491
2492
2493
2494
2495
2496
2497
2498
2499
2500
2501
2502
2503
2504
2505
2506
2507
2508
2509
2510
2511
2512
2513
2514
2515
2516
2517
2518
2519
2520
2521
2522
2523
2524
2525
2526
2527
2528
2529
2530
2531
2532
2533
2534
2535
2536
2537
2538
2539
2540
2541
2542
2543
2544
2545
2546
2547
2548
2549
2550
2551
2552
2553
2554
2555
2556
2557
2558
2559
2560
2561
2562
2563
2564
2565
2566
2567
2568
2569
2570
2571
2572
2573
2574
2575
2576
2577
2578
2579
2580
2581
2582
2583
2584
2585
2586
2587
2588
2589
2590
2591
2592
2593
2594
2595
2596
2597
2598
2599
2600
2601
2602
2603
2604
2605
2606
2607
2608
2609
2610
2611
2612
2613
2614
2615
2616
2617
2618
2619
2620
2621
2622
2623
2624
2625
2626
2627
2628
2629
2630
2631
2632
2633
2634
2635
2636
2637
2638
2639
2640
2641
2642
2643
2644
2645
2646
2647
2648
2649
2650
2651
2652
2653
2654
2655
2656
2657
2658
2659
2660
2661
2662
2663
2664
2665
2666
2667
2668
2669
2670
2671
2672
2673
2674
2675
2676
2677
2678
2679
2680
2681
2682
2683
2684
2685
2686
2687
2688
2689
2690
2691
2692
2693
2694
2695
2696
2697
2698
2699
2700
2701
2702
2703
2704
2705
2706
2707
2708
2709
2710
2711
2712
2713
2714
2715
2716
2717
2718
2719
2720
2721
2722
2723
2724
2725
2726
2727
2728
2729
2730
2731
2732
2733
2734
2735
2736
2737
2738
2739
2740
2741
2742
2743
2744
2745
2746
2747
2748
2749
2750
2751
2752
2753
2754
2755
2756
2757
2758
2759
2760
2761
2762
2763
2764
2765
2766
2767
2768
2769
2770
2771
2772
2773
2774
2775
2776
2777
2778
2779
2780
2781
2782
2783
2784
2785
2786
2787
2788
2789
2790
2791
2792
2793
2794
2795
2796
2797
2798
2799
2800
2801
2802
2803
2804
2805
2806
2807
2808
2809
2810
2811
2812
2813
2814
2815
2816
2817
2818
2819
2820
2821
2822
2823
2824
2825
2826
2827
2828
2829
2830
2831
2832
2833
2834
2835
2836
2837
2838
2839
2840
2841
2842
2843
2844
2845
2846
2847
2848
2849
2850
2851
2852
2853
2854
2855
2856
2857
2858
2859
2860
2861
2862
2863
2864
2865
2866
2867
2868
2869
2870
2871
2872
2873
2874
2875
2876
2877
2878
2879
2880
2881
2882
2883
2884
2885
2886
2887
2888
2889
2890
2891
2892
2893
2894
2895
2896
2897
2898
2899
2900
2901
2902
2903
2904
2905
2906
2907
2908
2909
2910
2911
2912
2913
2914
2915
2916
2917
2918
2919
2920
2921
2922
2923
2924
2925
2926
2927
2928
2929
2930
2931
2932
2933
2934
2935
2936
2937
2938
2939
2940
2941
2942
2943
2944
2945
2946
2947
2948
2949
2950
2951
2952
2953
2954
2955
2956
2957
2958
2959
2960
2961
2962
2963
2964
2965
2966
2967
2968
2969
2970
2971
2972
2973
2974
2975
2976
2977
2978
2979
2980
2981
2982
2983
2984
2985
2986
2987
2988
2989
2990
2991
2992
2993
2994
2995
2996
2997
2998
2999
3000
3001
3002
3003
3004
3005
3006
3007
3008
3009
3010
3011
3012
3013
3014
3015
3016
3017
3018
3019
3020
3021
3022
3023
3024
3025
3026
3027
3028
3029
3030
3031
3032
3033
3034
3035
3036
3037
3038
3039
3040
3041
3042
3043
3044
3045
3046
3047
3048
3049
3050
3051
3052
3053
3054
3055
3056
3057
3058
3059
3060
3061
3062
3063
3064
3065
3066
3067
3068
3069
3070
3071
3072
3073
3074
3075
3076
3077
3078
3079
3080
3081
3082
3083
3084
3085
3086
3087
3088
3089
3090
3091
3092
3093
3094
3095
3096
3097
3098
3099
3100
3101
3102
3103
3104
3105
3106
3107
3108
3109
3110
3111
3112
3113
3114
3115
3116
3117
3118
3119
3120
3121
3122
3123
3124
3125
3126
3127
3128
3129
3130
3131
3132
3133
3134
3135
3136
3137
3138
3139
3140
3141
3142
3143
3144
3145
3146
3147
3148
3149
3150
3151
3152
3153
3154
3155
3156
3157
3158
3159
3160
3161
3162
3163
3164
3165
3166
3167
3168
3169
3170
3171
3172
3173
3174
3175
3176
3177
3178
3179
3180
3181
3182
3183
3184
3185
3186
3187
3188
3189
3190
3191
3192
3193
3194
3195
3196
3197
3198
3199
3200
3201
3202
3203
3204
3205
3206
3207
3208
3209
3210
3211
3212
3213
3214
3215
3216
3217
3218
3219
3220
3221
3222
3223
3224
3225
3226
3227
3228
3229
3230
3231
3232
3233
3234
3235
3236
3237
3238
3239
3240
3241
3242
3243
3244
3245
3246
3247
3248
3249
3250
3251
3252
3253
3254
3255
3256
3257
3258
3259
3260
3261
3262
3263
3264
3265
3266
3267
3268
3269
3270
3271
3272
3273
3274
3275
3276
3277
3278
3279
3280
3281
3282
3283
3284
3285
3286
3287
3288
3289
3290
3291
3292
3293
3294
3295
3296
3297
3298
3299
3300
3301
3302
3303
3304
3305
3306
3307
3308
3309
3310
3311
3312
3313
3314
3315
3316
3317
3318
3319
3320
3321
3322
3323
3324
3325
3326
3327
3328
3329
3330
3331
3332
3333
3334
3335
3336
3337
3338
3339
3340
3341
3342
3343
3344
3345
3346
3347
3348
3349
3350
3351
3352
3353
3354
3355
3356
3357
3358
3359
3360
3361
3362
3363
3364
3365
3366
3367
3368
3369
3370
3371
3372
3373
3374
3375
3376
3377
3378
3379
3380
3381
3382
3383
3384
3385
3386
3387
3388
3389
3390
3391
3392
3393
3394
3395
3396
3397
3398
3399
3400
3401
3402
3403
3404
3405
3406
3407
3408
3409
3410
3411
3412
3413
3414
3415
3416
3417
3418
3419
3420
3421
3422
3423
3424
3425
3426
3427
3428
3429
3430
3431
3432
3433
3434
3435
3436
3437
3438
3439
3440
3441
3442
3443
3444
3445
3446
3447
3448
3449
3450
3451
3452
3453
3454
3455
3456
3457
3458
3459
3460
3461
3462
3463
3464
3465
3466
3467
34
```

1404
 1405 - For MOP, there must be explicit multiple objectives
 1406 - For pure discrete problems with special structures (e.g. graphs),
 1407 prefer CO over IP

1408

1409

1410

1411 **SYSTEM PROMPT FOR SCENARIO LABELING**

1412

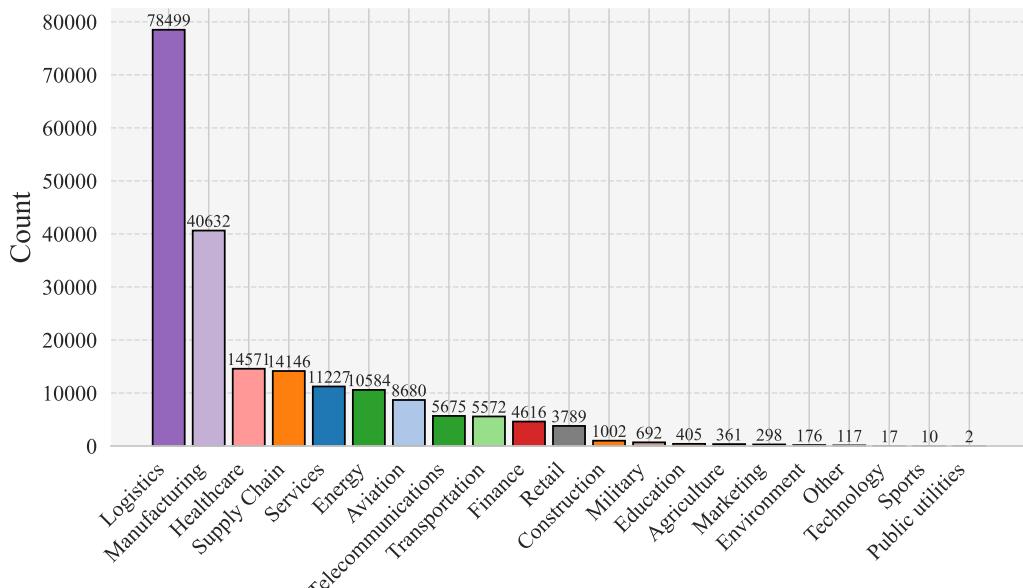
1413 Please classify the following optimization problem into one of
 1414 these application domains based on the core decision-making
 1415 context and primary business function, not just keywords
 1416 mentioned in the problem:

- 1417 1. Supply Chain: Decisions about inventory management, distribution
 1418 network, warehousing operations
- 1419 2. Finance: Decisions about portfolio management, investments, risk
 1420 management, financial planning
- 1421 3. Manufacturing: Decisions about production processes, quality
 1422 control, factory operations
- 1423 4. Transportation: Decisions about routing, vehicle scheduling,
 1424 fleet management, traffic flow, carrier selection
- 1425 5. Healthcare: Decisions about medical staff scheduling, patient
 1426 flow, hospital resources
- 1427 6. Energy: Decisions about power generation, energy conservation,
 1428 grid distribution
- 1429 7. Technology: Decisions about network design, data center
 1430 operations, cloud resources
- 1431 8. Retail: Decisions about store operations, pricing, inventory,
 1432 equipment, store layout
- 1433 9. Agriculture: Decisions about farming operations, crop planning,
 1434 irrigation
- 1435 10. Logistics: Decisions about delivery operations, warehouse
 1436 management, distribution
- 1437 11. Resources: Decisions about raw materials, equipment allocation,
 1438 material management
- 1439 12. Marketing: Decisions about campaign planning, budget allocation
 1440 , target selection
- 1441 13. Education: Decisions about course scheduling, resource
 1442 allocation in schools
- 1443 14. Environment: Decisions about environmental protection,
 1444 emissions control, conservation
- 1445 15. Construction: Decisions about project planning, construction
 1446 resource allocation
- 1447 16. Military: Decisions about military operations, deployment,
 1448 supply management
- 1449 17. Sports: Decisions about game scheduling, team formation,
 1450 strategy
- 1451 18. Telecommunications: Decisions about network coverage, bandwidth
 1452 allocation
- 1453 19. Aviation: Decisions about flight scheduling, crew assignment,
 1454 airport operations
- 1455 20. Services: Decisions about service operations, staff scheduling,
 1456 capacity management
- 1457 21. Public utilities: Decisions about utility services,
 1458 infrastructure management, service delivery
- 1459 22. Other: Problems that don't clearly fit into above categories

1454 # Problem:
 1455 {{Question}}

1456 # Output
 1457

1458
 1459 Let's think step by step, give the analysis of the problem and
 1460 classify it into one of the above application domains. Finally,
 1461 output the name of the domain in the following format:
 1462 Category: Name of the Domain
 1463
 1464 Note:
 1465 - Focus on the fundamental business function and decision-making
 1466 context
 1467 - Don't be misled by secondary keywords or background story
 1468 - Consider who is making the decision and what is their primary
 1469 business purpose



1493
 1494 Figure 6: Proportion of every scenarios of instances in OptMATH-Train (201K).
 1495
 1496
 1497
 1498
 1499
 1500
 1501
 1502
 1503
 1504
 1505
 1506
 1507
 1508
 1509
 1510
 1511

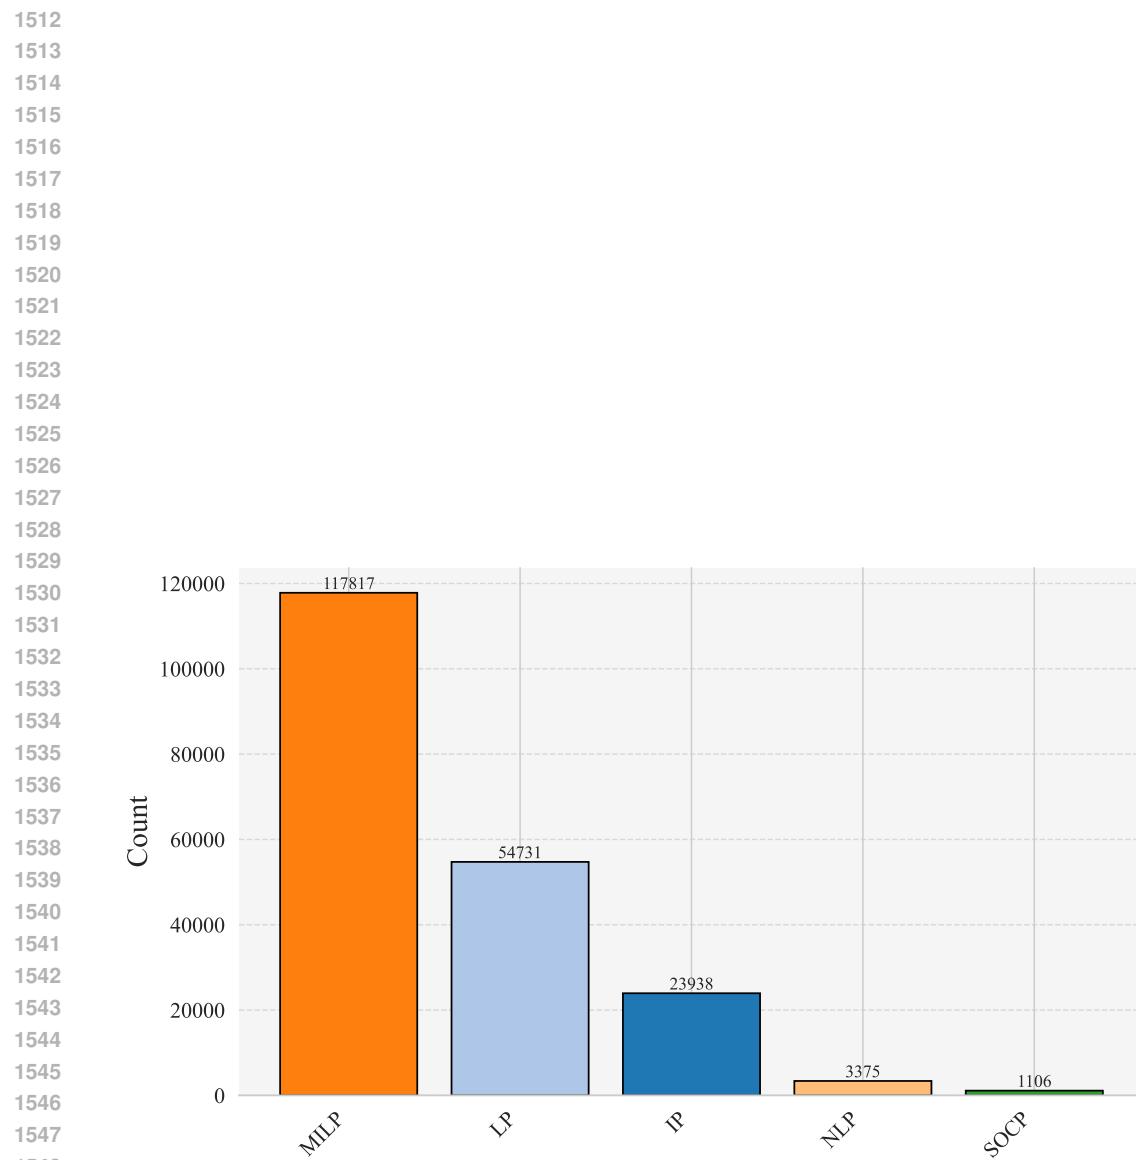


Figure 7: Proportion of every problem types of instances in OptMATH-Train (201K).