

LLMs for Mathematical Modeling: Towards Bridging the Gap between Natural and Mathematical Languages

Xuhan Huang¹, Qingning Shen¹, Yan Hu^{1*}, Anningzhe Gao², Benyou Wang¹

¹ The Chinese University of Hong Kong, Shenzhen

² Shenzhen Research Institute of Big Data

xuhanhuang, qingningshen@link.cuhk.edu.cn, huyan@cuhk.edu.cn

anningzhegao@gmail.com, wangbenyou@cuhk.edu.cn

Abstract

Large Language Models (LLMs) have demonstrated strong performance across various natural language processing tasks, yet their proficiency in mathematical reasoning remains a key challenge. Addressing the gap between natural and mathematical language requires advanced reasoning capabilities, approaching those of Artificial General Intelligence (AGI). However, the evaluation remains challenging, as perfectly representing reality is inherently elusive, and traditional methods like manual or direct comparison of mathematical statements (Ramamonjison et al., 2023) are insufficient for assessing true modeling ability. We propose a process-oriented framework to evaluate LLMs' ability to construct mathematical models, using solvers to compare outputs with ground truth. Introducing **Mamo**, a benchmark with 1,209 questions covering ordinary differential equations, linear programming, and mixed-integer linear programming, we enable automatic evaluation of modeling accuracy. The results show that existing LLMs struggle with complex mathematical modeling tasks, with larger models demonstrating superior performance, while open-source models remain competitive in simpler cases but still fall short of proprietary models in more challenging problems.

1 Introduction

Recent advancements in Large Language Models (LLMs) have garnered widespread interest, demonstrating remarkable capabilities across a broad spectrum of natural language processing tasks (Ouyang et al., 2022; OpenAI, 2023; Nijkamp et al., 2022; Tang et al., 2024). However, the domain of mathematical reasoning remains a critical aspect of the LLM evaluation (Wei et al., 2022). This focus gauges LLMs' proficiency in solving mathematical

challenges and reveals their underlying abilities in abstract reasoning and logic.

Natural language reflects the complexity and subtlety of human communication, while *mathematical language* provides accuracy and is ideal for representing the natural world (Giordano et al., 2013). Bridging their gap requires high-level cognitive skills akin to Artificial General Intelligence (AGI), a task at which LLMs show promise. Recent studies have already combined LLMs with solvers (Feng et al., 2023; Pan et al., 2023). In the field of optimization, AhmadiTeshnizi et al. (2023) introduced OptiMUS, an LLM-based agent that uses a solver to address optimization problems. The key task for LLM in OptiMUS involves first generating a mathematical formulation of a real-world problem, followed by writing Python code to invoke a solver. This implies the potential for testing LLM's modeling capabilities. However, the challenge lies in the evaluation of modeling, as perfectly representing reality may be inherently elusive.

To address this, we propose a new benchmarking strategy that evaluates mathematical modeling by utilizing solvers and answer verification. The core philosophy is to use solvers to resolve mathematical models and assess their accuracy by comparing the solver's output with the ground truth. Additionally, we introduce **Mamo**, a novel benchmark designed to assess the mathematical modeling capabilities of LLMs, leveraging solvers to shift the focus from conventional outcome-based evaluations to a more process-oriented approach. This strategy is particularly rational because it allows us to isolate and scrutinize the entire modeling process, from formulation to solution, providing a more holistic and rigorous evaluation of an LLM's mathematical reasoning abilities.

The contributions of this paper are as follows:

- We introduce an automatic evaluation framework to assess mathematical modeling

*Corresponding author: huyan@cuhk.edu.cn

through *solvers*, enabling exact answer matching.

- We have developed a new benchmark, named **Mamo**, specifically tailored for the evaluation of mathematical modeling capabilities. Mamo encompasses a wide array of modeling questions (with a total of 1,209 meticulously curated questions), including ordinary differential equations (ODEs) and optimization problems within linear programming and mixed-integer linear programming frameworks.

2 An Evaluation Framework for Mathematical Modeling

2.1 Definition: Mathematical Modeling

Mathematical modeling is widely considered an advanced skill typically possessed by experts. According to [Giordano et al. \(2013\)](#), mathematical models act as intermediaries that convert real-world problems into structured mathematical forms, facilitating the discovery of solutions. As illustrated in Figure 1, we aim to map a natural language question Q into a mathematical model M , which is expressed in mathematical language as follows:

$$f : Q \rightarrow M.$$

By solving this model, we can obtain an answer A to the original problem Q . Traditional mathematical modeling is a labor-intensive, highly specialized process that depends heavily on human expertise.

2.2 Gap between Natural and Mathematical Languages

Natural language captures the complexity and nuance of human thought, while *mathematical language* excels in precision and in modeling natural phenomena ([Giordano et al., 2013](#)). A single statement in natural language may correspond to different mathematical representations, and conversely, distinct natural language statements may map to the same mathematical formulation.

For example, consider the natural language statement: *The car starts with an initial velocity of zero.* This can correspond to at least two different mathematical representations:

Mathematical formulation 1: $y'(0) = 0$

Mathematical formulation 2: $v(0) = 0$.

Conversely, these two mathematical expressions may also correspond to natural language 2: *The population growth rate is initially zero.*

Bridging the gap between *natural language* and *mathematical language* requires advanced cognitive abilities, a task where LLMs have shown significant potential. However, evaluating the effectiveness of mathematical modeling remains challenging, as achieving a perfect representation of reality through models is often unattainable.

Q: A car with an initial velocity of 0 m/s, and acceleration 1 m/s ² , what is its displacement after 5 s?	
$M_1: y'(0) = 0$	$M_2: v(0) = 0$
$y''(t) = 1$	$v'(t) = 1$
$A = y(5)$	$A = \int_0^5 v(s) ds$

Figure 1: Q is the natural language problem, while M_1, M_2 are different mathematical models.

Challenges in Evaluation Due to the complex mapping between *natural language* and *mathematical language*, a real-world problem Q may correspond to multiple equivalent mathematical models, such as M_1, M_2, M_3, \dots . As illustrated in Figure 1, for a given problem Q , we have two distinct mathematical models, M_1 and M_2 , both of which correctly describe the same real-world problem and are considered equivalent in the context of Q . Traditionally, the evaluation of mathematical models, including determining their equivalence, has been performed by human experts, a process that is both labor-intensive and time-consuming. In optimization modeling area, [Ramamonjison et al. \(2023\)](#) directly compare mathematical statements by standardizing variable names and evaluating accuracy based on declaration matches. While this approach addresses differences in variable naming, the abstraction of variables from natural language context, which is critical in mathematical modeling, was not examined.

2.3 Evaluating Modeling Using Solvers and Answer Matching

The Evaluation Framework To facilitate the *automatic* evaluation of mathematical models, we propose an evaluation framework for mathematical modeling based on exact answer verification using solvers, as illustrated in Figure 2. A solver is an algorithmic tools designed to find solutions to various classes of problems. For instance, Gurobi is a well-known optimization solver ([Gurobi Optimization, LLC, 2024](#)), while Python packages ([Virtanen et al., 2020](#)) and MATLAB ([Inc., 2022](#)) are com-




Figure 2: The pipeline to use exact answer verification via an additional solver.

monly used for solving ODEs. Solvers are essential for operationalizing abstract models, denoted as \widehat{M} , generated by large language models (LLMs), translating them into concrete solutions \widehat{A} that can be analyzed and validated against real-world data. For a given real-world problem and its corresponding answer pair (Q, A) , LLMs generate a mathematical model, denoted \widehat{M} . The solver resolves this model, producing a solution \widehat{A} , which is then compared to the correct answer A . This framework directly evaluates the LLM’s problem-solving capability, with any inaccuracies in \widehat{M} typically leading to incorrect solutions \widehat{A} . By focusing on the final output, we avoid additional computational steps that could obscure the evaluation of the LLM’s performance.

The Philosophy The philosophy behind our evaluation framework is *goal-driven*, based on the idea that the ultimate aim of mathematical modeling is to produce accurate solutions to real-world problems. We prioritize the correctness of the solution \widehat{A} over the intermediate steps of model formulation. By comparing the solver’s output with the correct answer, we evaluate the effectiveness of the LLM’s generated model \widehat{M} . The assumption here is that if the model is well-constructed, it will lead to an accurate solution.

Despite its *goal-driven* structure, the framework retains a *process-oriented* emphasis on the modeling capability of LLMs. Unlike traditional benchmarks that treat answer correctness as an isolated metric, our approach explicitly evaluates the intermediate step of model formulation. The evaluative focus lies on whether the LLM constructs a mathematical model \widehat{M} that faithfully translates the real-world problem into mathematical language. Thus, while the solver’s answer validates the outcome, the core objective remains to assess the *modeling process*.

This dual perspective allows us to maintain a scalable yet principled evaluation framework.

While the solver’s output remains the validation mechanism, the primary objective is to assess the LLM’s ability to conceptualize and formulate mathematical problems rather than simply solving predefined mathematical tasks (more discussions are provided in Appendix M).

3 Mamo Benchmark

To implement this framework, we developed a new benchmark, **Mamo**, specifically for mathematical modeling.

3.1 Scope of Mamo Benchmark

Given the broad spectrum of mathematical modeling (Giordano et al., 2013) (see Appendix G) and the availability of specific solvers, our benchmark is meticulously designed to encompass domains of ODEs and linear programming problems. This section outlines the rationale behind our focus on these areas, driven by the criteria of *solver availability* and *result verifiability*.

Availability of Sophisticated Solvers: Advanced optimization solvers, such as COPT (Ge et al., 2023), along with Python libraries for solving ODEs, offer a strong foundation for testing the capabilities of LLMs. These tools enable a detailed and automated evaluation of the LLMs’ abilities to abstract, formulate, and accurately solve mathematical problems within these specific domains.

The Straightforward Verifiability of Results

While calculators address basic problems in these areas, the absence of a structured modeling approach hinders the assessment of LLMs’ ability to create solvable mathematical models.

Building on these considerations, we developed a benchmark focused on optimization and ODE, given the lack of universal PDE solvers. Compared to NLP4LP, introduced by (AhmadiTeshnizi et al., 2023) and focusing exclusively on optimization, our benchmark includes a larger number of

problems, with potential expansions into nonlinear optimization, probabilistic modeling, and PDE as new solvers emerge.

3.2 Data Selection

Our benchmark dataset comprises manually selected and GPT-generated questions, curated by a team with advanced mathematical education, including Ph.D. holders, and an average GPA of 3.89/4.0, ensuring strong analytical expertise. Details on the collectors' qualifications are provided in Appendix C.

Source Reliability For manual selection, we referenced textbooks and solutions in ODE and optimization (Thomas et al., 2016; Stewart et al., 2014; Lial et al., 2017; Braun, 1993; Boyce and DiPrima, 2012; Giordano et al., 2014; Zill, 2013; Bertsimas and Tsitsiklis, 1997; Hurlbert, 2010). We focused on deriving mathematical models and assigning realistic scenarios, ensuring answer accuracy. Each question was adapted to meet the benchmark criteria (Section 3.3) by adjusting scenarios or solutions as needed.

Question Formulation and Verification Data synthesis starts by creating mathematical constructs with random parameters, followed by answer computation to establish ground truth. GPT-4 then generates real-life scenarios, translating these models into context-rich problems, testing both mathematical reasoning and real-world contextualization. See the example in Appendix N.

3.3 Guidance of Data Collection

Suppose A is our answer to the question Q , $y(t)$ is the solution function in ODE problems

Prop. 1. $A = y(t_0)$ $A = a \times y(t_0) + b$

Prop. 2. $A = \max_{x,y \leq 0} (x + y) = 0.00$ $A = \sqrt{3}$

Prop. 3. Q : “in 3 significant figures”
 $A = 6.66$ $A = 6.7$

Prop. 4. Q : “In a company...”

Figure 3: Q is the natural language problem (may be ODE problem or optimization problem), where y is a solution function in ODE problems.

Our benchmark dataset follows specific criteria to ensure the problems are suitable for automated verification. Figure 3 provides examples of these criteria, with further explanations presented in Appendix T.

Property 1. Final-State Approach Solvability Problems should be structured so that a solver can directly compute the solution once the LLM formulates the mathematical model. For example, asking for the optimal value of an optimization problem is solvable using a 'final-state approach', whereas transforming an ODE solution requires extra steps and is not solvable this way.

Property 2. Unified and Numerical Answers

For ODEs, questions should seek the function's value at a specific time t_0 , and for optimization problems, the optimal value, ensuring a unified and numerical answer for clear comparison with the solver's output.

Property 3. Significant Figures and Precision

Each question will explicitly state the required significant figures or the level of numerical precision for the answer.

To assess the LLM's capability to abstract mathematical models from real-world scenarios, we introduce Property 4.

Property 4. Real-World Problem Context Questions should be framed as real-world problems without explicitly stating the underlying mathematical model.

3.4 Data Quality Checking

In the ODE section, we utilized GPT-4 to verify the solvability of generated questions as stated in Property 1, filtering out any invalid questions. A detailed evaluation identified issues such as missing information, incorrect answers, and unclear phrasing, leading to the removal or correction of problematic questions and resulting in a refined dataset. In the optimization section, specifically in the Easy_LP part, invalid questions were eliminated, and those with incorrect answers or misleading descriptions were corrected. For the Complex_LP section, questions involving incorrect mathematical models were revised (see an example in Appendix J). A summary of the review process is provided in Table 1.

For the cross-review process, 50 questions from the ODE dataset were selected for in-depth analysis by four independent reviewers¹. Reviewers were chosen based on their qualifications, including an average GPA of at least 3.5/4.0 in fundamental

¹The hourly wage is approximately 20 US dollars, with total compensation of around 473 US dollars.

Category	Initial Questions	Filtered	Corrected	Final Questions
ODE	383	37	41	346
Easy LP	688	36	8	652
Complex LP	211	0	40	211

Table 1: Summary of the review process.

Reviewer	Reviewer 1	Reviewer 2	Reviewer 3	Reviewer 4
Reviewer 1	-	0.67	0.71	0.56
Reviewer 2	0.67	-	0.65	0.50
Reviewer 3	0.71	0.65	-	0.52
Reviewer 4	0.56	0.50	0.52	-

Table 2: Inter-reviewer Cohen’s Kappa scores.

Reviewer	Accuracy Rate (%)
Reviewer 1	90.0
Reviewer 2	84.0
Reviewer 3	88.0
Reviewer 4	74.0

Table 3: Accuracy of reviewers.

math courses, completion or enrollment in the Ordinary Differential Equation course, and meeting English proficiency standards.

Reviewers evaluated the questions, classifying responses as either numerical answers or "error" if the question was problematic. The effectiveness of this review was assessed using Cohen’s Kappa to measure inter-rater reliability, along with individual accuracy rates. As shown in Table 2, the pairwise Cohen’s Kappa scores had an average of 0.60, reflecting moderate to substantial agreement among reviewers, with individual scores ranging from 0.50 to 0.71. Accuracy rates were moderately high, as detailed in Tables 3. For additional details on the review, see Appendix B.

3.5 Data Statistics

In the section on Ordinary Differential Equations, we present a total of 346 problems: 196 are based on first-order equations, 110 on second-order equations, and 40 on systems of equations, offering a comprehensive exploration across different levels. The optimization section is divided into two segments: Easy_LP, featuring 652 high school-level Mixed Integer Linear Programming (MILP) problems, and Complex_LP, comprising 211 undergraduate-level problems that integrate both LP and MILP. See the category statistics and the word cloud of the combined data in Appendix O.

4 Evaluation

4.1 Evaluation Protocol

4.1.1 Evaluation Details

Our evaluation protocol follows the methodology outlined in Section 2.3. First, the LLMs were prompted to generate formalized code (in each do-

main) to represent a mathematical model, denoted as \widehat{M} . This code was then executed by a solver, and the result, \widehat{A} , was compared with the ground truth answer, A .

For ODE problems, the LLM was prompted to generate Python code based on a natural language description. The generated code was executed, and the output was compared to the correct solution to assess its accuracy. For optimization problems, the LLM was asked to express the model in .lp format, after which an optimization solver COPT (Ge et al., 2023) was employed to find the optimal value for comparison with the correct answer. This process is independent of the specific solver used; for instance, MATLAB (Inc., 2022) can be used as an ODE solver, and Gurobi (Gurobi Optimization, LLC, 2024) can serve as an alternative optimization solver.

Benchmarking Models Our evaluation² includes both proprietary and open-source LLMs. The proprietary LLMs consist of the GPT-4 series (OpenAI, 2023), the Claude series³, and the Gemini series⁴. The open-source LLMs include the DeepSeek series (Shao et al., 2024; DeepSeek-AI, 2024), Llama-3.1 models (Dubey et al., 2024), Qwen-2.5 (Team, 2024), and Mixtral (Jiang et al., 2024).

Model Output We prompted the LLMs to generate Python code that called solvers (such as solve_ivp and odeint in SciPy, and dsolve in NumPy). In the optimization sections, the LLMs were asked to generate a standard .lp file. Compared to the testing methodology in NLP4LP (AhmadiTeshnizi et al., 2023), which involved prompt-

²See the settings in Appendix Q

³<https://www.anthropic.com/news/claude-3-family>

⁴<https://deepmind.google/technologies/gemini/>

ing the model to output Python code for invoking the optimization solver, the .lp format more closely aligned with the standard optimization form, making it more readable. Optimization solvers like COPT (Ge et al., 2023) could directly read and solve the .lp file. To better adhere to this format, we conducted tests using 3-shot learning. Detailed information about the number of few-shot samples used in our experiments and the results can be found in Appendix K and Appendix R. Additionally, at the start of the testing process, we performed an initial cleaning of the code, removing strings such as ````python```` that would have affected our testing.

4.1.2 Evaluation Metrics

According to Property 2 and Property 3, both the ground truth answer A and the answer \hat{A} , derived from solving the LLM’s output mathematical model \hat{M} , are numerical values with specified significant figures. To account for minor discrepancies, we apply the following comparison rule:

General Comparison Criterion : Let n denote the number of decimal places in A ($n = 0$ if A is an integer). The answer \hat{A} is considered correct if:

$$\left| \frac{\hat{A} - A}{A} \right| < 10^{-4} \text{ or } |\hat{A} - A| < \min\{10^{-n}, 10^{-2}\}.$$

<p>Ex. 1. $A = 10^6$, $\hat{A} = 10^6 + 1 \xrightarrow{n=0} \left \frac{A-\hat{A}}{A} \right = 10^{-6} < 10^{-4}$ ✓</p> <p>Ex. 2. $A = 10^{-6}$, $\hat{A} = 1.1 \times 10^{-6} \xrightarrow{n=6} \left \frac{A-\hat{A}}{A} \right = 0.1 > 10^{-4}$, but $A - \hat{A} = 10^{-7} < 10^{-n}$ ✓</p> <p>Ex. 3. $A = 0.1$, $\hat{A} = 0.12 \xrightarrow{n=1} \left \frac{A-\hat{A}}{A} \right = 0.2 > 10^{-4}$, and $A - \hat{A} = 0.02 > \min\{10^{-n}, 10^{-2}\} = 0.01$ ✗</p>

Figure 4: Examples illustrating the application of the General Comparison Criterion

This rule accounts for precision errors and ensures an accurate evaluation of \hat{A} while allowing for slight computational variations. Figure 4 demonstrates the application of this criterion with examples, illustrating how discrepancies between the correct and generated answers are assessed. For further details, see Appendix P and the case study in Appendix E.

4.2 Benchmarking Results

Table 4 presented the evaluation results (accuracy rate) of different language models on the Mamo

Benchmark. The performance of these models was assessed both in their original form and when enhanced by the language model (i.e., code modifier) itself to rectify syntax errors (with improvements indicated by the subscript). Appendix A provided results when GPT-4-0613 was used as the code modifier, and Appendix F presented the error analysis.

Mamo is Challenging As the results in Table 4 indicate, even state-of-the-art models struggled with complex tasks such as second-order ODEs, systems of ODEs, and Complex_LP. This is due to the increasing number of variables and the more intricate relationships among these variables compared to easier tasks like Easy_LP.

Take-away 1. Existing LLMs struggle with complex tasks in mathematical modeling.

The Performance of o1 Interestingly, o1-preview, although it excelled in complex tasks such as second-order ODEs and Complex_LP, did not outperform (and was even worse than GPT-3.5-Turbo) in Easy_LP tasks, which consist of high-school-level linear programming modeling problems. As shown in Table 6, the execution success rates of o1-preview in each category were nearly 100% across all methods (raw, self-modified, or GPT-4-modified), indicating that this underperformance was not due to limitations in coding ability.

Take-away 2. o1 excelled in complex tasks but underperformed in easy tasks.

Open-source Models vs. Proprietary Models

One of the most notable aspects of this benchmark is the emergence of open-source models as serious competitors to closed-source models. The Llama-3.1-405B, in particular, achieved the best overall score across both evaluation categories. The Llama-3.1 series, Qwen series, and DeepSeek-v2.5, despite being open-source models, demonstrated remarkable performance, even surpassing highly regarded commercially developed models such as GPT-4 and Claude-3 in categories like second-order ODEs and Optimization.

Take-away 3. Open-source models are competitive in simpler tasks, but there is still a gap compared to closed-source models in more complex tasks.

Effectiveness of Scaling The evaluation demonstrates a correlation between model size and performance across almost all tested categories. As

Models	ODE			Optimization		Overall (%)
	First order (%)	Second order (%)	System (%)	Easy LP (%)	Complex LP (%)	
Proprietary Models						
o1-preview	<u>71.43</u> +0.00	<u>50.00</u> +0.00	<u>45.00</u> +0.00	80.21 +0.00	<u>36.02</u> +3.79	67.60 +0.66
GPT-4o	67.86 +0.51	36.36 +0.91	40.00 +0.00	87.27 +0.15	22.75 +8.53	66.67 +1.73
GPT-4-turbo	64.80 +1.02	32.73 +0.91	40.00 +0.00	87.88 +0.00	23.22 +6.64	66.34 +1.41
GPT-4	59.18 +1.02	28.18 +0.91	40.00 +2.50	86.50 +0.46	21.08 +3.09	63.81 +1.12
GPT-3.5-turbo	16.33 +6.12	7.27 +4.55	10.00 +5.00	81.29 +3.53	9.48 +0.00	49.13 +3.48
Claude-3-opus	55.61 +3.57	33.64 +1.81	<u>45.00</u> +2.50	83.74 +0.31	9.48 +0.47	60.38 +1.08
Claude-3-sonnet	52.04 +2.04	24.55 +0.90	27.50 +5.00	83.59 +0.61	24.64 +0.48	60.96 +0.99
Claude-3-haiku	41.33 +9.18	23.64 +0.91	17.50 +2.50	78.53 +7.36	17.54 +1.42	54.84 +5.87
Gemini-1-pro	25.00 +8.16	13.64 +0.00	2.50 +2.50	78.83 +1.08	14.69 +0.00	50.45 +1.99
Gemini-1.5-pro	64.80 +1.02	30.00 +0.00	<u>45.00</u> +5.00	85.28 +0.00	30.81 +5.21	66.09 +1.24
Open-source Models						
DeepSeek-v2.5	67.86 +0.00	36.36 +0.00	32.50 +0.00	86.20 +1.99	13.33 +0.00	65.27 +1.07
DeepSeek-math-7B-base	3.57 +0.00	2.73 +0.00	2.50 +0.00	49.39 +1.22	6.64 +0.00	28.70 +0.66
DeepSeek-math-7B-rl	1.02 +0.51	1.82 +0.00	5.00 +0.00	20.55 +0.62	0.00 +0.47	11.58 +0.50
Llama-3.1-8B-instruct	30.61 +1.53	4.55 +1.81	10.00 +2.50	74.39 +2.60	13.27 +0.00	48.14 +1.90
Llama-3.1-70B-instruct	62.24 +0.52	33.64 +5.45	32.50 +0.00	85.74 +0.36	22.75 +1.42	64.44 +1.02
Llama-3.1-405B-instruct	66.84 +1.02	40.00 +0.00	30.00 +0.00	86.20 +0.76	34.12 +4.74	<u>67.91</u> +1.40
Qwen-2.5-32B-instruct	60.71 +0.51	38.18 +0.00	35.00 +0.00	84.05 +0.00	12.32 +1.42	61.95 +0.33
Qwen-2.5-72B-instruct	60.71 +0.00	37.27 +0.91	32.50 +0.00	<u>89.42</u> +0.15	11.43 +1.90	64.53 +0.50
Mixtral-8x7B-instruct	23.47 +4.59	4.55 +1.81	10.00 +0.00	67.79 +2.15	9.95 +0.00	42.85 +2.06
Mixtral-8x22B-instruct	52.04 +1.53	11.82 +1.82	20.00 +0.00	81.90 +2.61	19.91 +1.42	57.82 +2.06

¹ Here GPT-4-turbo denotes GPT-4-turbo-2024-04-09; GPT-4o denotes GPT-4o-2024-05-13; GPT-4 refers to GPT-4-0613; GPT-3.5-turbo denotes GPT-3.5-turbo-0125.

Table 4: Evaluation Results on the Mamo Benchmark. The **subscript** denotes improvement using the original LLM as a **code modifier**, which corrects syntax errors without changing the underlying logic, differentiating modeling from formatting errors (see Section 4.3). The "Overall" score represents the weighted average of correct rates, weighted by question count. The highest score in each category is marked in **boldface**; an **underline** signifies the top score with the LLM itself as code modifier. The score after using LLM as code modifier is computed as the **sum** of the original score and the value in the subscript.

Models	ODE (%)			Optimization (%)		
	raw	self-modified	GPT-4-modified	raw	self-modified	GPT-4-modified
Proprietary Models						
o1-preview	99.71	100.00	100.00	97.80	99.65	98.73
GPT-4o	96.82	99.42	99.42	92.93	97.91	97.91
GPT-4-turbo	96.82	99.13	98.84	93.86	97.80	97.56
GPT-4	94.51	97.69	97.69	90.96	95.13	95.13
GPT-3.5-turbo	37.28	65.03	93.93	87.95	89.80	96.76
Claude-3-Sonnet	84.68	93.06	97.40	91.89	93.97	96.76
Claude-3-Haiku	78.03	92.77	98.27	82.61	92.58	96.29
Claude-3-Opus	88.73	96.82	97.98	90.38	93.40	95.94
Gemini-1.0-pro	56.36	69.94	88.73	88.06	89.11	94.21
Gemini-1.5-pro	93.64	98.27	98.84	95.48	98.73	96.87
Open-source Models						
DeepSeek-v2.5	97.69	99.13	99.13	90.36	93.83	97.55
DeepSeek-math-7B-base	28.90	29.77	79.77	55.50	58.40	75.67
DeepSeek-math-7B-rl	5.20	6.64	80.64	19.93	21.55	55.97
Llama-3.1-8b-instruct	57.51	72.97	92.20	80.99	88.30	94.55
Llama-3.1-70b-instruct	92.48	98.27	97.98	92.93	96.99	97.34
Llama-3.1-405b-instruct	98.84	99.71	99.71	92.58	98.03	95.48
Qwen-2.5-32B-instruct	97.69	98.84	99.13	93.17	96.29	94.67
Qwen-2.5-72B-instruct	97.40	97.69	97.98	90.80	94.06	91.15
Mixtral-8x7B-instruct	50.87	69.08	91.62	84.13	87.83	96.29
Mixtral-8x22B-instruct	80.92	88.73	96.24	87.72	93.51	96.29

Table 5: Execution Success Rates for Raw, Self-Modified, and GPT-4 Modified Methods, denoting scenarios without code modifiers, utilizing the LLM as a code modifier, and employing GPT-4-0613 as the code modifier. This table highlights the success rates of **file execution** (correct format) across these three scenarios, differing from the Table 4 which shows the **correctness rate**.




Figure 5: The testing pipeline involving code modifier.

indicated by the results of the Llama-3.1, Qwen, and Mixtral series, for models with the same architecture but different parameter scales (e.g., Llama-3.1-7B vs. Llama-3.1-405B), larger models tended to outperform their smaller counterparts. This observation highlights the effectiveness of scaling in the mathematical reasoning tasks of LLMs. However, the effectiveness of scaling may plateau on easier tasks, such as Easy_LP problems.

Take-away 4. *Larger LLMs perform better in mathematical modeling.*

4.3 On the Impact of Format Errors

Mitigation of Formatting Errors It is important to differentiate between errors in *modeling* and those related to coding or file *formatting* (See Appendix I). When the LLM output contains syntax errors in Python or .lp code, these errors must be corrected without altering the logic of the code to ensure that the evaluation focuses solely on modeling. To address syntax issues, we employ a **code modifier**, as the pipeline shown in Figure 5 (More details are discussed in Appendix D). In our experiments, we used the original language model and GPT-4-0613 as code modifiers, with the former selected for *reproducibility* and the latter for its superior coding *proficiency* and *stability*. A comparison of results from different code modifiers is provided in Section 4.2 and Appendix A, with prompts detailed in Appendix R.

Setting The testing pipeline is shown in Figure 5. As mentioned in Section 4.1.1, we selected the original LLM and GPT-4-0613 as code modifiers. The data presented in Table 5 provides a comparison of the performance (rate of correct format) of various models using different modification methods: i) *no modification*, ii) *self-modification*, and iii) *modification provided by GPT-4-0613*.

Observation Combining Table 4 and Table 5, the **improvement** can be attributed to misjudg-

ments caused by the LLM’s coding ability. For LLMs with a lower initial correct format rate (such as GPT-3.5-Turbo and Llama-3.1-8B in ODE tasks), the relative improvement is higher. This is because, in several cases, their output codes, initially judged as "Error" due to compile errors, actually represent correct mathematical models. It is also worth noting that the **improvement** has little effect on the ranking, indicating that the influence of different coding abilities on our benchmark is minimal.

Take-away 5. *Format errors have minimal impact on mathematical modeling.*

5 Related Work

Recent studies have advanced the use of large language models (LLMs) in mathematical problem-solving by introducing new datasets and solver methodologies. For example, Frieder et al. (2024) and Yuan et al. (2023) developed datasets to test LLMs on various mathematical problems and arithmetic expressions. Lewkowycz et al. (2022) focused on training models using natural language paired with LaTeX-formatted math from arXiv. To address complex computational tasks, Zhang et al. (2024) introduced the CARP dataset, while He et al. (2024) presented OlympiadBench, a bilingual benchmark for competition-level math reasoning. Liu et al. (2024) also developed MathBench, structured to assess LLMs’ theoretical knowledge. Additionally, Mirzadeh et al. (2024) proposed GSM-Symbolic, a benchmark that generates symbolic variants of GSM8K questions. Their study also introduced GSM-NoOp, a dataset containing irrelevant but seemingly related information to assess model robustness, demonstrating that while LLMs exhibit sensitivity to numerical variations, they remain resilient to superficial modifications.

In solver development, efforts have focused on integrating LLMs into novel problem-solving ap-

proaches. [Yang et al. \(2023\)](#) explored LLMs as direct solvers, highlighting their autonomous problem-solving potential. [He-Yueya et al. \(2023\)](#) introduced a hybrid method that combines LLMs with external symbolic solvers for equation solving, while [Pan et al. \(2023\)](#) developed LOGIC-LM, a framework blending LLMs with symbolic solvers to tackle logical problems. [AhmadiTeshnizi et al. \(2023\)](#) contributed to both dataset development and solver integration by creating NLP4LP, a benchmark for LP and MILP problems, and OptiMUS, an LLM-based agent for optimization problem solving. In line with these developments, [Lyu et al. \(2023\)](#) introduced Faithful-CoT, a two-stage framework that translates natural language into symbolic reasoning chains (represented as code) before solving the problem using corresponding solvers. This approach aligns closely with our benchmark's objective of assessing the faithful translation of problems into structured mathematical formulations.

6 Conclusion and Future Directions

Mathematical modeling is a crucial yet challenging aspect of mathematical reasoning. Evaluating LLMs' ability to handle modeling is essential, as their performance will continue to improve over time, making specialized benchmarks in specific areas necessary to accurately assess and push their capabilities. Due to the complexity of mathematical models, assessment is difficult. To address this, we introduce the Mamo benchmark, which incorporates solvers to validate answers and assess the correctness of the modeling process. Mamo demonstrates that LLMs can be effectively evaluated on complex reasoning tasks, where the correct answer may vary but remains equivalent in principle. This approach paves the way for further research into LLMs' reasoning capabilities.

Limitations

While the current benchmarking methodology necessitates LLMs to generate Python code for ODEs and .lp files for optimization problems, it opens up exciting avenues for future research. The process, while effective, may inadvertently influence the LLMs' focus towards formalization over the conceptual modeling, potentially affecting the depth of mathematical reasoning. This opens the door for benchmarks that better distinguish between an LLM's modeling skills and formalization abilities. Future work could focus on assessing LLMs' mathematical modeling directly, bypassing formalization for a more nuanced evaluation.

Ethics Statement:

The benchmark is designed with a stringent ethical framework to ensure that all problems are socially responsible and do not perpetuate or encourage harmful biases or stereotypes. Questions are meticulously reviewed to avoid any content that could lead to the dissemination of misinformation, support unethical practices, or cause harm to individuals or groups. This includes but is not limited to issues of privacy, security, and the potential for misuse of the model. Furthermore, the benchmark abstains from any problem that could indirectly endorse unethical behaviors or decisions in real-world scenarios, upholding the highest standards of academic integrity and social responsibility.

Acknowledgment

This work was supported by the Shenzhen Science and Technology Program (JCYJ20220818103001002), Shenzhen Doctoral Startup Funding (RCBS20221008093330065), Tianyuan Fund for Mathematics of National Natural Science Foundation of China (NSFC) (12326608), Shenzhen Key Laboratory of Cross-Modal Cognitive Computing (grant number ZDSYS20230626091302006), and Shenzhen Stability Science Program 2023.

References

- Ali AhmadiTeshnizi, Wenzhi Gao, and Madeleine Udell. 2023. Optimus: Optimization modeling using mip solvers and large language models. *arXiv preprint arXiv:2310.06116*.
- Dimitris Bertsimas and John N. Tsitsiklis. 1997. *Introduction to linear optimization*. Athena Scientific.
- William E. Boyce and Richard C. DiPrima. 2012. *Elementary differential equations and boundary value problems*, 10 edition. Wiley.
- Martin Braun. 1993. *Differential equations and their applications an introduction to applied mathematics*, 4 edition. Springer.
- DeepSeek-AI. 2024. Deepseek-v2: A strong, economical, and efficient mixture-of-experts language model. *Preprint*, arXiv:2405.04434.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Al-lonsius, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Lauren Rantala-Yearly, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan

Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Rapparthi, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaoqing Ellen Tan, Xinfeng Xie, Xuchao Jia, Xuwei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aaron Grattafiori, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alex Vaughan, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Franco, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Barambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, Danny Wyatt, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Firat Ozgenel, Francesco Caggioni, Francisco Guzmán, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Govind Thattai, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Japhet Asher, Jean-Baptiste Gaya,

Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Karthik Prasad, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kun Huang, Kunal Chawla, Kushal Lakhotia, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Maria Tsim-poukelli, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usumier, Nikolay Pavlovich Laptev, Ning Dong, Ning Zhang, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratan-chandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Rohan Maheswari, Russ Howes, Ruty Rinott, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Kohler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vitor Albiero, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaofang Wang, Xiao-jian Wu, Xiaolan Wang, Xide Xia, Xilun Wu, Xinbo Gao, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yuchen Hao, Yundi Qian, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, and Zhiwei Zhao. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.

Jiazhan Feng, Ruochen Xu, Junheng Hao, Hiteshi Sharma, Yelong Shen, Dongyan Zhao, and Weizhu Chen. 2023. Language models can be logical solvers. *arXiv preprint arXiv:2311.06158*.

- Simon Frieder, Luca Pinchetti, Ryan-Rhys Griffiths, Tommaso Salvatori, Thomas Lukasiewicz, Philipp Petersen, and Julius Berner. 2024. Mathematical capabilities of chatgpt. *Advances in Neural Information Processing Systems*, 36.
- Dongdong Ge, Qi Huangfu, Zizhuo Wang, Jian Wu, and Yinyu Ye. 2023. Cardinal Optimizer (COPT) user guide. <https://guide.coap.online/copt/en-doc>.
- Frank R. Giordano, Steven B. Horton, and William P. Fox. 2014. *A first course in mathematical modeling*, 5 edition. Brooks/Cole.
- Frank R. Giordano, Maurice D. Weir, and William P. Fox. 2013. *A First Course in Mathematical Modeling*, 5th edition. Brooks/Cole, Cengage Learning.
- Gurobi Optimization, LLC. 2024. [Gurobi Optimizer Reference Manual](#).
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, et al. 2024. Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. *arXiv preprint arXiv:2402.14008*.
- Joy He-Yueya, Gabriel Poesia, Rose E Wang, and Noah D Goodman. 2023. Solving math word problems by combining language models with symbolic solvers. *arXiv preprint arXiv:2304.09102*.
- Glenn Hurlbert. 2010. *Linear Optimization: The simplex workbook*. Springer.
- The MathWorks Inc. 2022. [Matlab version: 9.13.0 \(r2022b\)](#).
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L  lio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Th  ophile Gervet, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2024. [Mixture of experts](#). *Preprint*, arXiv:2401.04088.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with pagedattention](#). *Preprint*, arXiv:2309.06180.
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. 2022. Solving quantitative reasoning problems with language models. *Advances in Neural Information Processing Systems*, 35:3843–3857.
- Margaret L. Lial, Raymond N. Greenwell, and Nathan P. Ritchey. 2017. *Calculus with applications*, 11 edition. Pearson.
- Hongwei Liu, Zilong Zheng, Yuxuan Qiao, Haodong Duan, Zhiwei Fei, Fengzhe Zhou, Wenwei Zhang, Songyang Zhang, Dahua Lin, and Kai Chen. 2024. [Mathbench: Evaluating the theory and application proficiency of llms with a hierarchical mathematics benchmark](#). *Preprint*, arXiv:2405.12209.
- Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. 2023. [Faithful chain-of-thought reasoning](#). *Preprint*, arXiv:2301.13379.
- Iman Mirzadeh, Keivan Alizadeh, Hooman Shahrokhi, Oncel Tuzel, Samy Bengio, and Mehrdad Farajtabar. 2024. [Gsm-symbolic: Understanding the limitations of mathematical reasoning in large language models](#). *Preprint*, arXiv:2410.05229.
- Erik Nijkamp, Bo Pang, Hiroaki Hayashi, Lifu Tu, Huan Wang, Yingbo Zhou, Silvio Savarese, and Caiming Xiong. 2022. [Codegen: An open large language model for code with multi-turn program synthesis](#). *arXiv preprint arXiv:2203.13474*.
- R OpenAI. 2023. [Gpt-4 technical report](#). arxiv 2303.08774. *View in Article*, 2(5).
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Liangming Pan, Alon Albalak, Xinyi Wang, and William Yang Wang. 2023. [Logic-lm: Empowering large language models with symbolic solvers for faithful logical reasoning](#). *arXiv preprint arXiv:2305.12295*.
- Rindranirina Ramamonjison, Timothy T. Yu, Raymond Li, Haley Li, Giuseppe Carenini, Bissan Ghaddar, Shiqi He, Mahdi Mostajabdaveh, Amin Banitalebi-Dehkordi, Zirui Zhou, and Yong Zhang. 2023. [NL4Opt Competition: Formulating Optimization Problems Based on Their Natural Language Descriptions](#). *arXiv preprint*. ArXiv:2303.08233 [cs].
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y.K. Li, Y. Wu, and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#).
- James Stewart, Dan Clegg, and Saleem Watson. 2014. *Calculus: Early transcendentals*, 8 edition. Cengage.
- Zhengyang Tang, Xingxing Zhang, Benyou Wan, and Furu Wei. 2024. [Mathscale: Scaling instruction tuning for mathematical reasoning](#). *arXiv preprint arXiv:2403.02884*.

- Qwen Team. 2024. [Qwen2.5: A party of foundation models](#).
- George B. Thomas, Maurice D. Weir, Joel Hass, and Frank R. Giordano. 2016. *Thomas' calculus*, 13 edition. Pearson.
- Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. 2020. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature methods*, 17(3):261–272.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V Le, Denny Zhou, and Xinyun Chen. 2023. Large language models as optimizers. *arXiv preprint arXiv:2309.03409*.
- Zheng Yuan, Hongyi Yuan, Chuanqi Tan, Wei Wang, and Songfang Huang. 2023. How well do large language models perform in arithmetic tasks? *arXiv preprint arXiv:2304.02015*.
- Beichen Zhang, Kun Zhou, Xilin Wei, Xin Zhao, Jing Sha, Shijin Wang, and Ji-Rong Wen. 2024. Evaluating and improving tool-augmented computation-intensive math reasoning. *Advances in Neural Information Processing Systems*, 36.
- Dennis G. Zill. 2013. *A first course in differential equations with modeling applications*, 10 edition. Cengage.

A GPT-4-0613 as Code Modifier

Selecting GPT-4 as a code modifier generally results in better performance than using the LLM itself for self-improvement. For instance, DeepSeek-v2’s accuracy in ODE tasks improves from 36.41% with self-modification to 48.80% when GPT-4 is used as the code modifier. This pattern indicates that GPT-4 is a more effective code modifier, enhancing the accuracy and performance of various models more reliably than self-improvement.

B Cross Review

In ODE part, we have the following review process:

1. We employed GPT-4 to confirm that the newly generated questions are amenable to being solved by solvers using the final-state approach. There were only 10 out of 383 questions are invalid. All were filtered, 373 questions were left.
2. Then we conducted a thorough evaluation of the questions. Our dataset initially comprised 373 questions. We identified 29 questions that lacked sufficient information, 21 that had incorrect answers, and 18 that featured unclear statements. Following our assessment, 27 questions were deleted, and 41 were corrected to ensure clarity and correctness.

Finally, we have 346 problems in ODE part.

We have also conducted the cross-review process. See the details in Appendix B.1 and the results on Appendix B.1.1.

In the optimization section, particularly with easy LP part, out of 688 entries, 12 were found to be completely invalid (not meeting the criteria), 26 had incorrect answers, and six had misleading descriptions. All invalid entries were filtered out, and the remaining issues were corrected manually. Among them, eight of them were corrected and others (36 questions) were filtered out. In complex LP part, for the original 211 questions, 40 were found formulated by wrong mathematical models. All of them were corrected.

B.1 Cross-Review Process

Furthermore, adhering to the scaling law evident in the dataset’s distribution, we selected 50 questions from the ODE parts in our dataset for a deeper analysis. Four independent reviewers (see their

qualification in Section C) were tasked with assessing 50 questions each ⁵. The responses could either be numeric or labeled as “error” if a question was deemed problematic. This section outlines the metrics used to evaluate the reviewers’ responses against pre-defined correct answers, which also include the “error” label for any flawed questions. We use the same metric as described in Subsection 4.1.2. Here are our results:

B.1.1 Results

The effectiveness of the review process was quantified using the metrics defined in Subsection 4.1.2, focusing on the accuracy and inter-rater reliability among reviewers. The statistical outcomes are summarized as follows:

- **Average Cohen’s Kappa:** The average Cohen’s Kappa across all reviewer pairs was 0.60, indicating a moderate to substantial agreement. This suggests that reviewers generally agreed on the classification of answers, though there were variations in some cases.
- **Minimum and Maximum Cohen’s Kappa:** The minimum Kappa value recorded was 0.50, and the maximum was 0.71. The spread in these values highlights areas where alignment and training could potentially enhance consistency.
- **Accuracy Rates:** The individual accuracy rates were as follows:
 - Reviewer 1: 90.0%
 - Reviewer 2: 84.0%
 - Reviewer 3: 88.0%
 - Reviewer 4: 74.0%

These rates reflect the precision with which each reviewer matched the standard answers, including their recognition of problematic questions correctly labeled as “error.”

These results provide insights into the overall effectiveness and areas of improvement for the cross-review process, particularly in terms of aligning understanding and interpretation of the evaluation criteria among reviewers.

⁵The hourly wage is approximately 20 US dollars (equivalent in local currency), with a total participant compensation of around 473 US dollars (equivalent in local currency).

Models	ODE			Optimization		Overall (%)
	First order (%)	Second order (%)	System (%)	Easy LP (%)	Complex LP (%)	
Proprietary Models						
o1-preview ††	71.43	50.00	45.00	80.21	37.91	67.49
GPT-4o ††	68.37	37.27	40.00	87.42	29.38	68.07
GPT-4-turbo ††	65.82	33.64	40.00	88.19	27.96	67.58
GPT-4 ††	60.20	29.09	42.50	86.96	24.17	64.93
GPT-3.5-turbo ††	44.90	20.00	20.00	85.43	9.48	57.48
Claude-3-Opus ††	59.18	36.36	47.50	84.20	14.22	62.37
Claude-3-Sonnet ††	56.63	26.36	32.50	85.12	27.96	63.44
Claude-3-Haiku ††	51.53	25.45	22.50	86.81	18.96	61.54
Gemini-1-pro ††	40.31	18.18	10.00	80.21	15.64	54.51
Gemini-1.5-pro ††	65.82	30.91	50.00	85.28	34.12	67.08
Open-source Models						
DeepSeek-v2.5 ††	67.86	36.36	32.50	86.20	13.33	64.20
DeepSeek-math-7b-base ††	13.78	9.09	5.00	59.05	9.48	36.72
DeepSeek-math-7b-rl ††	13.78	7.27	7.50	47.70	1.90	29.20
Llama-3.1-8b-instruct ††	33.67	8.18	15.00	82.06	13.27	53.27
Llama-3.1-70b-instruct ††	62.76	39.09	32.50	86.81	23.70	65.76
Llama-3.1-405b-instruct ††	67.86	40.00	30.00	86.66	34.60	68.40
Qwen-2.5-32b-instruct ††	61.73	38.18	35.00	84.05	13.27	62.28
Qwen-2.5-72b-instruct ††	60.71	37.27	32.50	89.57	11.43	64.61
Mixtral-8x7b-instruct ††	32.14	11.82	20.00	75.46	9.95	49.38
Mixtral-8x22b-instruct ††	55.10	17.27	22.50	85.43	22.27	61.21

Table 6: Evaluation result on Mamo Benchmark. The remark †† indicates it additionally uses the GPT-4-0613 to rectify syntax error of the code and .lp file generated by the corresponding LLM.

C Collector Qualifications

The benchmarking process benefits from the expertise of reviewers with robust backgrounds in mathematics, underscored by their academic accomplishments. The team is composed of individuals whose education encompasses critical mathematical disciplines, including calculus (I and II), linear algebra (both introductory and advanced levels), optimization, probability, and ordinary differential equations. The average Grade Point Average (GPA) across these foundational courses is approximately 3.89 out of a 4.0 scale. This high level of academic achievement demonstrates the reviewers’ strong grasp of essential mathematical concepts and their application. Furthermore, the team is enriched by members who hold Ph.D. degrees in mathematics, further solidifying the depth of expertise and analytical skills available for the benchmarking process.

C.1 Qualification of Reviewers

We select the reviewers mentioned in Appendix B.1 based on the following criteria:

1. Undergraduate students who have taken fundamental mathematical courses (Calculus and Linear Algebra) and have an average GPA of at least 3.5/4.0.
2. Have completed or are currently taking courses in Ordinary Differential Equations (ODE).
3. Meet one of the following English proficiency requirements: TOEFL > 80, IELTS > 6.5, a minimum grade of B+ in all English classes, or a Gaokao English score > 120.

These criteria ensure their ability to read questions in English and solve ODE problems.

D Understand Code Modifier

Figure 5 shows the evaluation process involving the code modifier. The introduction of code modifiers is to maintain the fairness of the evaluation: we try to reduce conditions when the LLM outputs the code with a correct mathematical model, but

the code ends with a compile error due to the limit of coding ability. We consider two factors when applying code modifier:

- (a) How often does a format error occur?
- (b) How many passes on average are needed?

The factor (a) is exactly the Table 5. As for (b), the number of passes required to fix format errors depends on the code modifier chosen. To study the factor (b), To study factor (b), we applied GPT-4-0613 as a code modifier on the chosen format error files and iteratively called it until the code was executed successfully. Table 7 displays the distribution of the number of pass the code modifier was called until the errors were fixed. The chosen files were generated by GPT-4o and DeepSeek-v2.5, representing a close-source and an open-source model, respectively.

	1	2	3	4	≥5	all
LP	74	35	10	6	11	136
ODE	88	3	1	1	9	112

Table 7: Number of pass to apply code modifier

On average, 1.68 passes were needed to correct format errors when using GPT-4-0613. We found that a single pass was sufficient to correct the majority of format errors. To balance efficiency and accuracy, we only applied a single pass for the code modifier throughout our experiments.

E Metric

The comparison process is as the following:

$$C_{A,\hat{A}} = \begin{cases} 1, & \text{if } \left| \frac{\hat{A} - A}{A} \right| \leq 1 \times 10^{-4} \text{ (1)} \\ & \text{or} \\ & |\hat{A} - A| < 1 \times \min\{10^{-n}, 10^{-2}\} \text{ (2),} \\ 0, & \text{otherwise.} \end{cases}$$

Our consideration: (1) The difference is much less than the Answers, so can be omitted.

(2) The only difference lies in the last few digits of the required significant figure of the answer. By selecting different thresholds $\xi \times \min\{10^{-n}, 10^{-2}\}$, we can adjust the tolerance for the discrepancy, ensuring that $|\hat{A} - A| < \xi \times \min\{10^{-n}, 10^{-2}\}$. We have explored the relationship between various thresholds and the accuracy of GPT-4o.

Figure 6 plots the results. In our metric, we choose strict $\xi = 1$.

F Error Analysis

In this section, we will analyze some typical errors.

F.1 Deviation from Instruction

It is common for some models that show a low capability of few-shot learning, especially for small models. For example, in a case Llama-3-8B output ‘I wish you good luck’ when prompting with some few-shot demonstrations in mathematical modeling.

F.2 Syntax Error

In ODE parts, the syntax errors is obviously the syntax errors in the reposed python code. In Optimization part, the .lp file require more strict syntax than the universal optimization formulation. The Figure 7 and 8 show the typical syntax errors in .lp format.

F.3 Modeling Error

One of the most common errors in modeling is the lack of recognition of decision variable types. The types of decision variables in optimization problems are divided into discrete (sometimes integer, sometimes binary) and continuous. Figure 9 is an example.

G Scope of Mamo

The scope of Mamo is chosen under the consideration in Section 3.1. As Table 9 shows, Mamo focus on the ODE and Optimization areas.

H Examples of Different \hat{M}

Different models will typically lead to different final answers after solving by the solvers. For example, in ODE part:

$$\hat{M}_1 : y' + y = 0$$

and

$$\hat{M}_2 : y' + 2y = 0$$

Suppose we ask for $y(1)$, then we have $\hat{A}_1 = e^{-1}$ and $\hat{A}_2 = e^{-2}$, they have different values.

Then for the optimization,

Table 8: Accuracy at various thresholds for different categories(%), where $\xi_i = i$

ξ	0	1	2	3	4	5	6	7	8	9
ODE	48.55	50.87	54.62	56.36	57.51	58.96	59.83	60.12	60.69	61.85
Easy LP	87.27	87.27	87.27	87.27	87.27	87.27	87.27	87.27	87.27	87.27
Complex LP	22.75	22.75	22.75	22.75	22.75	22.75	22.75	22.75	22.75	22.75




Figure 6: Accuracy of different threshold in different category

\widehat{M}_1 :

$$\begin{aligned} & \max_{x,y} && x + y \\ & \text{subject to} && x + y \leq 20, \\ & && x - y \geq 0, \\ & && 0 \leq x \leq 10, \\ & && 0 \leq y \leq 5, \\ & && x, y \in \mathbb{Z}. \end{aligned}$$

\widehat{M}_2 :

$$\begin{aligned} & \max_{x,y} && x - y \\ & \text{subject to} && x + y \leq 20, \\ & && x - y \geq 0, \\ & && 0 \leq x \leq 10, \\ & && 0 \leq y \leq 5, \\ & && x, y \in \mathbb{Z}. \end{aligned}$$

the optimal value of \widehat{M}_1 is $\widehat{A}_1 = 10$, while the optimal value of \widehat{M}_2 is $\widehat{A}_2 = 0$. The minor difference in the mathematical models will result in different final values.

I Examples of Syntax Errors

Some outputs of LLM contain the correct model but not follow the syntax of solvers. For examples, the standard .lp file for

$$\begin{aligned} & \max_{x,y} && x + y \\ & \text{subject to} && x + y \leq 20, \\ & && x - y \geq 0, \\ & && 0 \leq x \leq 10, \\ & && 0 \leq y \leq 5, \\ & && x, y \in \mathbb{Z}. \end{aligned}$$

is

```
Minimize
obj: x + y
Subject To
    c1: x + y <= 20
    c2: x - y >= 0
Bounds
    0 <= x <= 10
    0 <= y <= 5
Generals
```

```

Minimize
obj: 0

Subject To
S2: x2 >= 4 - x5
S3: x3 >= 5 - x5
S4: x4 >= 5 - x2
S5: x5 >= 4 - x3
S6: x6 >= 7 - x3
S7: x7 >= 2 - x4
ST: xT >= 1 - x5
T7: x7 + xT >= 1 - xT
T6: x6 + xT >= 3 - x7
T3: x3 + xT >= 5 - x5
T2: x2 + xT >= 8 - x4

Bounds
0 <= x5 <= 9
0 <= x2 <= 4
0 <= x3 <= 5
0 <= x4 <= 5
0 <= x5 <= 4
0 <= x6 <= 7
0 <= x7 <= 2
0 <= xT <= 1

Generals
x5
x2
x3
x4
x5
x6
x7
xT

End

```

Figure 7: The syntax error is caused by the constraint, in .lp format, the variables should be placed on the left hand side of the inequality

```

Minimize
10 x + 20 y

Subject To
c1: 50 <= x + y
c2: 2 x + y <= 120

Bounds
x >= 0
y >= 0

Generals
x
y

```

Figure 8: Same type of syntax error as Figure 7

```

x
y
End

```

However, if we replace the constraint $c1$ with $x \leq 20 - y$, the code will be reported as an error, even if the two constraints are equivalent. To specifically test the modeling ability of the LLMs, we apply code modifier to fix the syntax errors.

```

Minimize
obj: 29 x12 + 97 x13 + 72 x14 + 29 x21 + 69 x23 + 52 x24 + 97 x31 + 69 x32 + 91 x34 + 72 x41 + 52 x42 + 91 x43

Subject To
c1: x12 + x13 + x14 = 1
c2: x21 + x23 + x24 = 1
c3: x31 + x32 + x34 = 1
c4: x41 + x42 + x43 = 1
c5: x12 - x21 = 0
c6: x13 - x31 = 0
c7: x14 - x41 = 0
c8: x21 - x12 - x23 - x24 = 0
c9: x31 - x13 - x32 - x34 = 0
c10: x41 - x14 - x42 - x43 = 0

Bounds
0 <= x12 <= 1
0 <= x13 <= 1
0 <= x14 <= 1
0 <= x21 <= 1
0 <= x23 <= 1
0 <= x24 <= 1
0 <= x31 <= 1
0 <= x32 <= 1
0 <= x34 <= 1
0 <= x41 <= 1
0 <= x42 <= 1
0 <= x43 <= 1

Generals
x12
x13
x14
x21
x23
x24
x31
x32
x34
x41
x42
x43

End

```

Figure 9: In this TSP problem, the decision variable should be binary, instead of "General", which means "integer"

Category	In Mamo
Changes and Differences	✗
Proportionality and Geometric Similarity	✗
Optimization Problem Modeling	✓
Differential Equations	✓
Probabilistic Modeling	✗

Table 9: The taxonomy of mathematical modeling. In Mamo benchmark, we only include the categories where sophisticated Solvers is available.

J Example of Mathematical Model Correction

In Figure 10, we display an example of correction of the mathematical model in Complex_LP. In the optimization model of the max flow problem, variables were initially set in the wrong signs in two constraints(C_1, C_3). The corrected constraints(C'_1, C'_3) are also displayed.

Maximum Flow Problem:

The network consists of n nodes connected by m edges. The goal is to maximize the total flow from the source node (node 0) to the sink node (node $n - 1$), subject to the following constraints:

Maximize: c

Subject to:

$$C_1 : \sum_{j=1}^{n-1} x_{j0} - \sum_{j=1}^{n-1} x_{0j} - \mathbf{c} = \mathbf{0} \quad (\text{Inflow Conservation at Source})$$

$$C_2 : \sum_{\substack{j=0 \\ j \neq i}}^{n-1} x_{ji} - \sum_{\substack{j=0 \\ j \neq i}}^{n-1} x_{ij} = 0 \quad \forall i \in \{1, 2, \dots, n-2\} \quad (\text{Intermediate Nodes})$$

$$C_3 : \sum_{j=0}^{n-2} x_{jn-1} - \sum_{j=0}^{n-2} x_{n-1j} + \mathbf{c} = \mathbf{0} \quad (\text{Outflow Conservation at Sink})$$

$$C_4 : 0 \leq x_{ij} \leq \text{capacity}(i, j) \quad \forall (i, j) \in \text{Edges} \quad (\text{Capacity Constraints})$$

In the **Inflow Conservation** and **Outflow Conservation** part, the correct constraints should be:

$$C'_1 : \sum_{j=1}^{n-1} x_{j0} - \sum_{j=1}^{n-1} x_{0j} + \mathbf{c} = \mathbf{0}$$

$$C'_3 : \sum_{j=0}^{n-2} x_{jn-1} - \sum_{j=0}^{n-2} x_{n-1j} - \mathbf{c} = \mathbf{0}$$

Figure 10: The example of model correction

K Few-shot Experiment

To test the sensitivity of different models to the number of prompts, we conducted a sensitivity analysis using the scaling law to examine the effect of the number of shots on model performance. In our analysis, we tested the performances of 11 models on ODE problems using prompts with 0 shot, 1 shot, 3 shots, 5 shots, and 10 shots. The results, shown in Figure 11, illustrate the ODE problem-solving accuracy of different models under varying numbers of shots. Similarly, we performed sensitivity tests on optimization problems, assessing the accuracy of 10 models with prompts of 0 shot, 1 shot, 3 shots, and 5 shots. The results are depicted in Figure 12.

Additionally, we specifically tested the 0-shot performance of GPT-4-0613 on all questions in Easy LP and Complex LP categories. The accuracy of GPT-4-0613 was 66.56% in Easy LP and 14.69% in Complex LP, resulting in an overall LP accuracy of 53.65%, consistent with the results obtained from sampling.

L Testing Process

We use 3-shot learning in testing; see examples in Appendix R. The examples of testing process in ODE and optimization is shown in Figure 13 and Figure 14.

M Probability of Getting Correct \hat{A} from Wrong \hat{M}

There are chances for the **wrong** mathematical model \hat{M} to lead to correct answer \hat{A} . However, the probability is low due to the wide spread of the continuous answer space and lack of prior knowledge. We examined the distribution of all the answers in Mamo dataset. The probability of guessing an ODE answer within a 0.01 interval is about 2.9%, and for LP, it's around 7.3%. These low probabilities highlight the difficulty of correctly guessing the answer without proper modeling.




Figure 11: Results of models accuracy and shots on ODE problems




Figure 12: Results of models accuracy and shots on optimization problems




Figure 13: Example of testing process in optimization

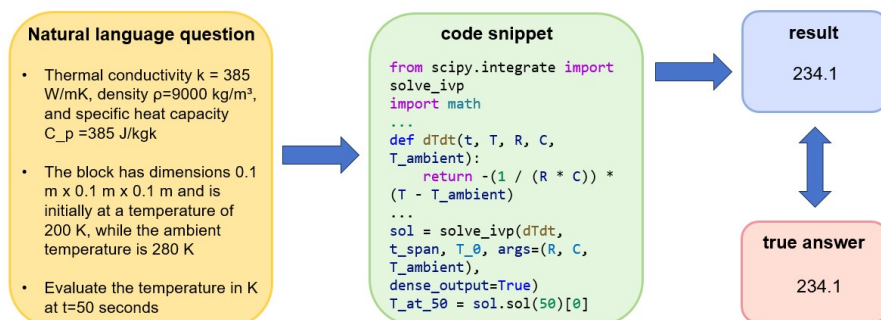


Figure 14: Example of testing process in ODE

N Question Formulation

Figure 15 shows the example of synthesis process in ODE. The process of synthesizing data starts with the creation of typical mathematical constructs, into which random parameters are introduced to give them specificity. This is followed by the computation of answers to establish a ground truth. Subsequently, GPT-4 is employed to craft real-life scenarios based on these predefined models, effectively translating abstract mathematical concepts into tangible, context-rich problems.

O Data Description and Word Cloud

We present a comparative table showcasing the number of scenarios across various categories, from manufacturing to natural science, in Table 10. We also make the word cloud of the combined data in Figure 16

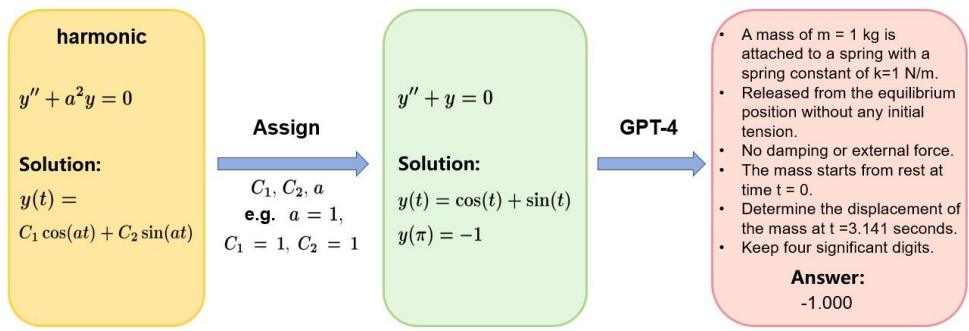


Figure 15: Example of synthesis process in ODE

Category	Manufacturing	Environment	Business	Daily Life	Natural Science	Others
LP	75	114	265	258	44	77
ODE	85	74	54	22	109	2

Table 10: Comparison of ODE and LP dataset of different categories



Figure 16: The word cloud of the data after filtering the requirement words such as 'significant figures', 'rounded'

P Evaluation script

The evaluation script is as following:

```
def comp(output, standard_answer):
    dif = abs(float(output) - float(standard_answer))
    if float(standard_answer) == 0:
        rate = dif * 1e-4
    else:
        rate = dif / float(standard_answer)
    if abs(rate) < 1e-4:
        return 1
    else:
        return 0

def compare_output_with_standard(output, standard_answer):
    try:
        float_output = float(output)
    except ValueError:
        return False

    if '.' in standard_answer:
        digit = len(standard_answer.split('.')[1])
        if digit <= 2:
            digit = 2
        s_ans = float(standard_answer) * 10 ** digit
        ans = float_output * 10 ** digit
        return (abs(ans - s_ans) < 1 or comp(output, standard_answer))
    else:
        digit = 2
        s_ans = float(standard_answer) * 10 ** digit
        ans = float_output * 10 ** digit
        return (abs(ans - s_ans) < 1 or comp(output, standard_answer))
```

Q Settings

We have the following settings:

- We tested the following models by calling API: All the closed source models except o1-preview, DeepSeek-V2.5, Llama-3.1 series, Qwen-2.5 series with temperature=0.8. o1-preview was called under default temperature since the change of temperature is not allowed.
- Other models was deployed and inference using VLLM (Kwon et al., 2023) on machine in slurm cluster with 8 * A100, with temperature 0.8 and max_length 4096 tokens.

R Prompts

The following figures show the prompts we use in the benchmarking. Figure 17,18 display the prompt we use in the evaluation process, while Figure 19,20 list the prompts we use for calling code modifiers. As shown in Figure 19,20, the code modifier is prompted to adjust the code based solely on the code and error information during execution, without any information about the problem itself. This ensures that the code modifier does not alter the underlying mathematical model based on problem details.

[Instruction]

Assume you are a virtual assistant with expertise in optimization, specifically in creating .lp files for linear programming problems. Your task is to translate given natural language problems into optimization models, formatted as .lp files. When you receive a question, it might include mathematical expressions in LaTeX format. Your job is to interpret these expressions accurately and model the problem in .lp format. Your response must adhere to the following guidelines:

- The optimization model must be written in .lp format, adhering to the conventions and syntax appropriate for linear programming problems.
- The model should be designed so that solving it yields the optimal value, which directly answers the question posed.
- Your response should be an entire .lp file, ready to be processed by a linear programming solver. Ensure that the file contains no comments or extraneous content beyond the model itself.
- Handle LaTeX expressions with care to ensure that the mathematical aspects of the problem are accurately represented in the .lp model.
- If the solution needs to be rounded to an integer, make use of the 'General' integer constraint in the .lp file to specify integer variables, please do not use 'General' if there are no requirements for the integer variables.

Here comes the examples:

[Example_1:]

(the input)

A manufacturing company produces two types of products: X and Y . The production cost for each unit of product X is 50, while the cost for each unit of product Y is

10. There are constraints in the production process, such that twice the number of units produced from product X , plus the number of units from product Y , cannot exceed 200 due to resource limitations. In addition, to meet market demand, four times the number of units produced from product X , plus the number of units from product Y , must be at least 50.

Considering these constraints and given that both products can only be produced in whole numbers due to their physical nature, what is the minimum total cost needed for producing these items while satisfying all conditions? Provide your answer rounded to the nearest dollar.

Your response:

Minimize

obj: $50x + 10y$

Subject To

c1: $2x + y \leq 200$

c2: $4x + y \geq 50$

Bounds

$x \geq 0$

$y \geq 0$

Generals

x

y

End

[...]

Please craft the .lp file according to these instructions, focusing on delivering a model that is directly solvable to obtain the answer. And Please follow the syntax like examples to write the .lp file.

Here comes the question:

[question]

Generate the contents of an .lp file for this problem, starting with the objective function and followed by the constraints, without any additional sentences. The constraints should be formatted as 'variable + variable \geq number' for inequalities, all the variables should be on the left hand side of the inequality. Ensure there is a space between variables and their coefficients.

Your response:

Figure 17: The prompt for testing in optimization part. [question] refers to a question in the benchmark. [...] refers to the examples. We use 3-shot in testing

S Licensing Information

This content(including benchmark) is licensed under the Creative Commons Attribution-ShareAlike 4.0 International (CC BY-SA 4.0) license.

[Instruction]

Assume you are a virtual assistant with expertise in ordinary differential equations (ODEs), particularly in formulating ODE models from natural language descriptions and solving them using Python's `'solve_ivp'`, `'odeint'` from SciPy, and `'dsolve'` from SymPy. Your primary task is to convert the given natural language problems into ODE models and then apply ODE solvers to find the solutions.

When you receive a question, it will often contain mathematical expressions in LaTeX format, which you need to interpret accurately. Your response must be in the form of Python code that directly outputs the final solution to the problem upon execution. This Python code should adhere to the following criteria:

1. The output should only display the final answer of the problem.
2. Utilize a 'final-state-approach' where the code immediately prints the solution after solving the ODE, without showing any intermediate steps.
3. Ensure the answer is rounded to the specified number of significant figures or decimal places. Use Python's 'round' function for decimal rounding. For significant figures, include and use the following function in your code:

```
``` python
def round_to_significant_figures(num, sig_figs):
if num != 0:
return round(num, -int(math.floor(math.log10(abs(num)))) + (1 - sig_figs))
else:
return 0 # Handles the case of num being 0
```
```

4. Your response should be entirely in Python code, formatted to run directly without modifications.
5. Handle LaTeX expressions in the problem statement carefully to ensure accurate modeling and solution.

Here comes the examples:

[Examples]

[...]

Please process the problem according to these instructions, focusing solely on delivering the Python code that meets these requirements. And Please directly generate the code without any explanation(except the comments in the code).

the following lines are forbidden:

“

Here's the Python code that solves the given problem and meets the specified requirements:

```
``` python
```

```
```
```

This code uses the `'solve_ivp'` function from SciPy to solve the initial value problem for the given differential equation. The `'round_to_significant_figures'` function is included to round the final answer to the specified number of significant figures. Upon execution, the code will directly output the amount of pollutant left in the tank after 5 minutes, rounded to five significant figures.

”

Please avoid the above sentences.

Take a deep breathe before answering the question. This is a piece of cake to you.

Here comes the question:

[question]

Your response:

Figure 18: The prompt for testing in ODE part. [question] refers to a question in the benchmark. [...] refers to the examples.

[Instruction]

You are experienced python engineer, the following codes may have some errors (in syntax), with the error information **[error_info]** , please fix the errors to make sure the code can run successfully.
 If there is no error, please return the original code.
 And please do not change the original logic of the code. Your response should be entirely in Python code, formatted to run directly without modifications.
 Take a deep breathe before answering the question. This is a piece of cake to you.
 PLEASE do not response anything except the code, No other comments outside the code are allowed.

The followings are forbidden:

“

The code is almost correct, but there is a syntax error in the function 'round_to_significant_figures'. The round function is missing a closing parenthesis. Here is the corrected version:

```

  \ \ \ python
  \ \ \
  ”

```

Please avoid resposing above sentences.
 Please output only the corrected code, with no additional text or explanations.
 Here comes the code:
[code]
 The correct code is:

Figure 19: The prompt for fixing syntax error in ODE. **[error]** refers to the message when executing the Python code. While **[code]** refer to the code which contains syntax error

[Instruction]

You are experienced engineer in optimization, please fix the errors to make sure the code can run successfully.
 If there is no error, please return the original code.
 And please do not change the original logic of the lp file. Your response should be entirely in .lp format, formatted to run directly without modifications.
 Take a deep breathe before answering the question. This is a piece of cake to you.
 PLEASE do not response anything except the code, No comments are allowed.
 The followings are forbidden:

“

This is the correct .lp file:

```

  \ \ \ lp
  \ \ \
  ”

```

Please avoid resposing above sentences.
 Please output only the corrected code, with no additional text or explanations.
 Here comes the lp:
[lp_code]
 Generate the correct .lp format, starting with the objective function and followed by the constraints, without any additional sentences. The constraints should be formatted as 'variable + variable >= number' for inequalities, all the variables shoule on the left hand side of the inequality. For example, make $a + b + c \leq d$ into $a + b + c - d \leq 0$. Ensure there is a space between variables and their coefficients (coefficients should be numerical).
 The correct lp code is:

Figure 20: The prompt for fixing syntax error in LP. **[lp_code]** refer to the lp code which contains syntax error

T Explanation of Criteria

In this section, we explain the reasoning behind the validation of each proposition shown in Figure 3, based on the properties outlined in our benchmark dataset.

T.1 Property 1: Final-State Approach Solvability

Property 1 states that problems should be solvable directly, using a 'final-state approach'. This means that once the model is formulated, the solver should be able to compute the final answer without requiring additional transformations. We assess two propositions based on this property:

- **Proposition 1:** $A = y(t_0)$ is **valid** as it asks for the value of the solution function $y(t)$ at a specific time t_0 , which is a final-state solution. A solver can compute this directly.
- **Proposition 1:** $A = a \times y(t_0) + b$ is **invalid** because it requires an additional transformation (i.e., multiplying by a and adding b), which does not adhere to the final-state solvability approach.

T.2 Property 2: Unified and Numerical Answers

Property 2 requires the answer to be a unified and numerical value, particularly for ODE problems where the function's value at a specific time is sought. For optimization problems, the optimal value should be clear.

- **Proposition 2:** $A = \max(x + y)$ for $x, y \leq 0 = 0.0$ is **valid** as it yields a clear numerical result, specifically the maximum of $x + y$ within the defined constraints.
- **Proposition 2:** $A = \sqrt{3}$ is **invalid** because while $\sqrt{3}$ represents a number, it is not expressed as a numerical answer. In practice, this could lead to ambiguities if different solvers provide decimal approximations (e.g., 1.732) instead of the exact expression.

T.3 Property 3: Significant Figures and Precision

Property 3 mandates that each question specifies the required precision or number of significant figures for the answer.

- **Proposition 3:** $A = 6.66$ is **valid** because it presents the number to three significant figures, which is explicitly required.

- **Proposition 3:** $A = 6.7$ is **invalid** as it is given only to two significant figures, which does not meet the required precision.

T.4 Property 4: Real-World Problem Context

Finally, Property 4 introduces the need for real-world context in problem framing, without explicitly stating the mathematical model. This requires the solver to abstract the appropriate mathematical model from the problem description.

- **Proposition 4:** "In a company..." aligns with this property, as the problem is framed as a real-world scenario. The mathematical model is not explicitly given, challenging the solver to formulate it based on the provided context.