第七章 大模型智能决策

修贤超

https://xianchaoxiu.github.io

致谢：本教案由李建豪协助准备

目录

- 7.1 智能决策
- 7.2 模型构建
- 7.3 算法设计
- 7.4 方案验证
- 7.5 场景应用
- 7.6 小结

智能决策

- 大模型是指**具有大规模参数和复杂计算结构的机器学习模型**，广泛应用于文本生成、图片生成和视频生成等场景

- Attention is All You Need, 2017

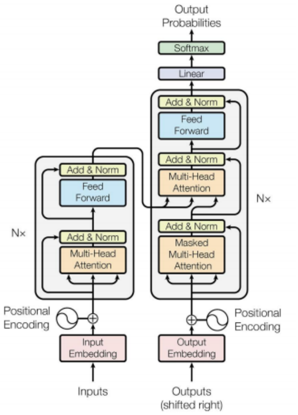

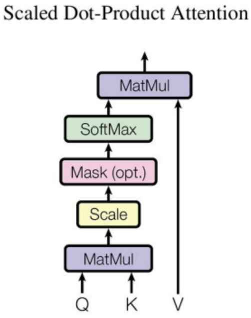

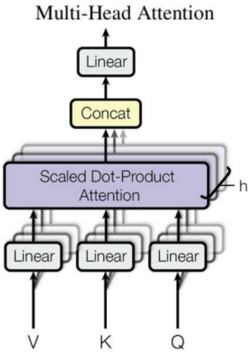

- Large Language Models for Mathematical Reasoning: Progresses and Challenges, 2024

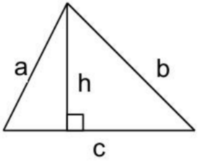

$\mathcal{Q}$: a=7 inches; b=24 inches; c=25 inches; h=6.72 inches; What is its area? (Unit: square inches)
$\mathcal{A}$: 84

$\mathcal{Q}$: Beth bakes 4, or 2 dozen batches of cookies in a week. If these cookies are shared amongst 16 people equally, how many cookies does each person consume?
$\mathcal{R}$: Beth bakes 4 2 dozen batches of cookies for a total of $4*2 =<< 4*2 = 8 >>$ 8 dozen cookies. There are 12 cookies in a dozen and she makes 8 dozen cookies for a total of $12*8 =<< 12*8 = 96 >>$ 96 cookies. She splits the 96 cookies equally amongst 16 people so they each eat $96/16 =<< 96/16 = 6 >>$ 6 cookies.
$\mathcal{A}$: 6

智能决策

- A Survey on Large Language Models for Code Generation, 2024

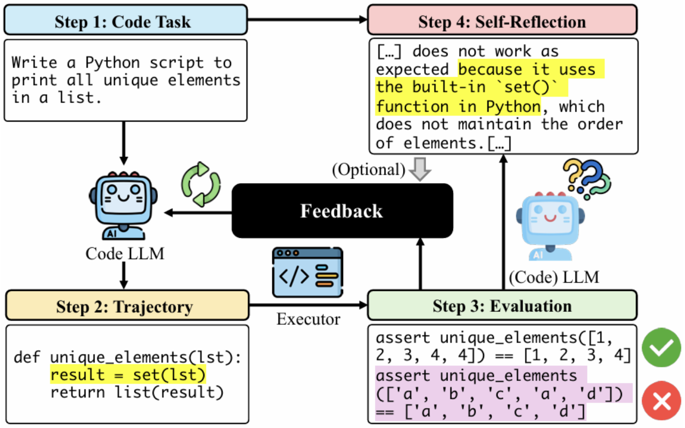

智能决策

- http://www.llm4ad.com/

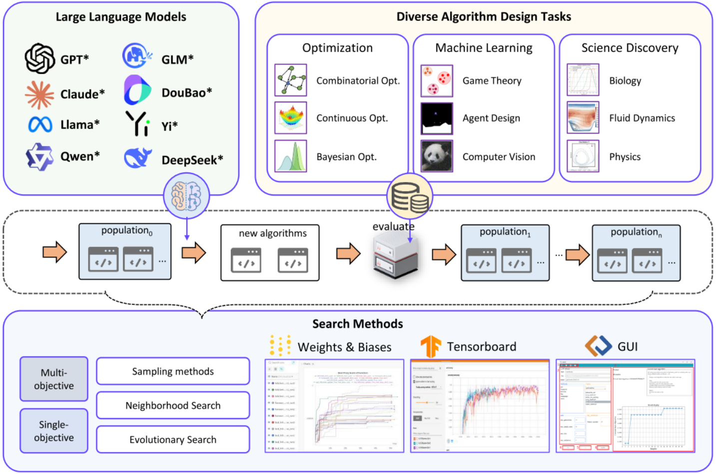

目录

- 7.1 智能决策
- 7.2 模型构建
- 7.3 算法设计
- 7.4 验证机制
- 7.5 应用场景
- 7.6 小结

基于提示的方法

- Chain-of-Experts: When LLMs Meet Complex Operations Research Problems, 2023

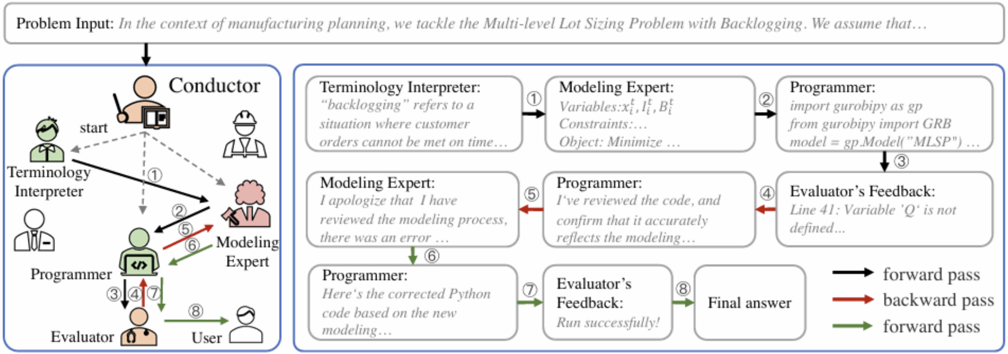

■ 实验结果

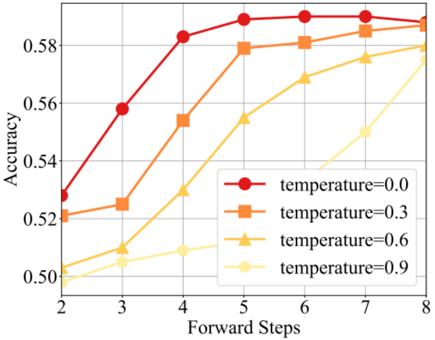

(a) CoE performance on different parameter settings.

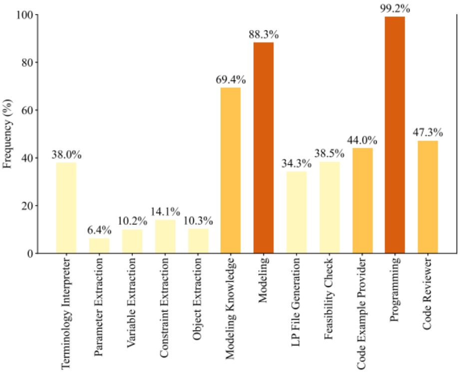

(b) Selection frequency of individual expert.

基于提示的方法

- Scalable Optimization Modeling with (MI)LP Solvers and Large Language Models, 2024

Task: *Debug the runtime error for the material capacity limit constraint*

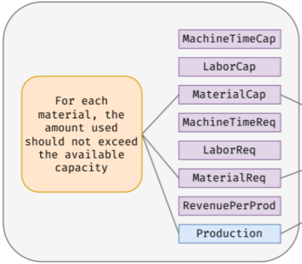

Prompt

The execution of the following code results in a runtime error:

```
import numpy as np
import gurobipy as gp

R = data["R"] # scalar parameter
P = data["P"] # scalar parameter
MaterialCap = np.array(data["MaterialCap"]) # ['R']
MaterialReq = np.array(data["MaterialReq"]) # ['R', 'P']


Prod = model.addVars(P, vtype=gp.GRB.CONTINUOUS, name="production")

# Add constraints for the quantity of raw material usage not exceeding available amounts
for j in range(N):
    model.addConstr(gp.quicksum(MaterialReq[j, i] * Prod[i] for i in range(P)) \\
<= Available[j], name=f"material_usage_limit_{j}")
```

Here is the error message:

IndexError: index 4 is out of bounds for axis 0 with size 4

Identify the error and fix it.

- 实验结果

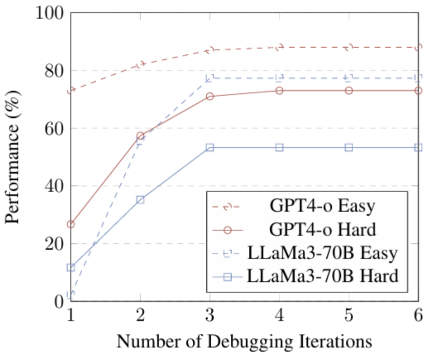
100
80
60
40
20
0
1
2
3
4
5
6
Performance (%)
Number of Debugging Iterations
GPT4-o Easy
GPT4-o Hard
LLaMa3-70B Easy
LLaMa3-70B Hard

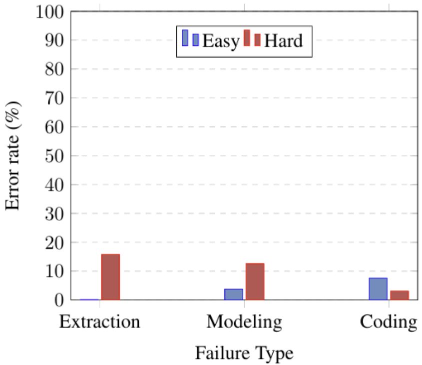
100
90
80
70
60
50
40
30
20
10
0
Error rate (%)
Easy
Hard
Extraction
Modeling
Coding
Failure Type

- ORLM: A Customizable Framework in Training Large Models for Automated Optimization Modeling, 2025

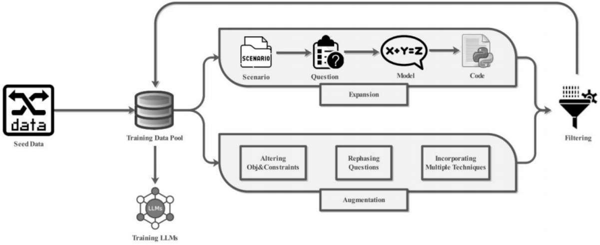

Table 7 Ablation study on OR-Instruct augmentations.

Method	NL4OPT	MAMO EasyLP	MAMO ComplexLP	IndustryOR	Micro Avg	Macro Avg
Full Augmentations	78.3%	80.6%	43.1%	21.0%	68.6%	55.7%
w/o Altering Obj&Const	77.5%	79.2%	36.4%	20.0%	66.4%	53.2%
w/o Rephrasing Questions	74.2%	77.3%	41.1%	15.0%	65.1%	51.9%
w/o Multiple Modeling	78.3%	78.0%	38.8%	18.0%	66.2%	53.2%

- LLMOPT: Learning to Define and Solve General Optimization Problems from Scratch, 2025

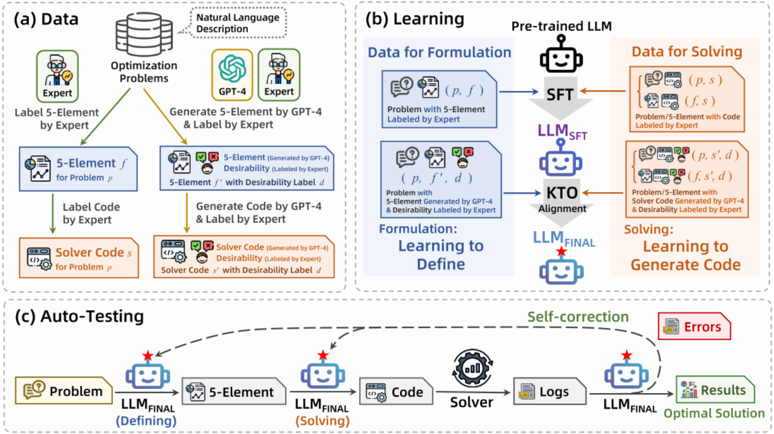

实验结果

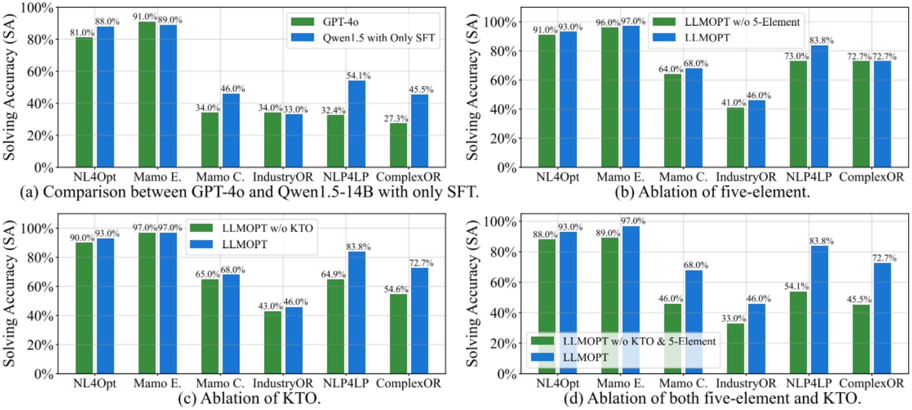

(a) Comparison between GPT-4o and Qwen1.5-14B with only SFT.

(b) Ablation of five-element.

(c) Ablation of KTO.

(d) Ablation of both five-element and KTO.

基于学习的方法

- OptMATH: A Scalable Bidirectional Data Synthesis Framework for Optimization Modeling, 2025

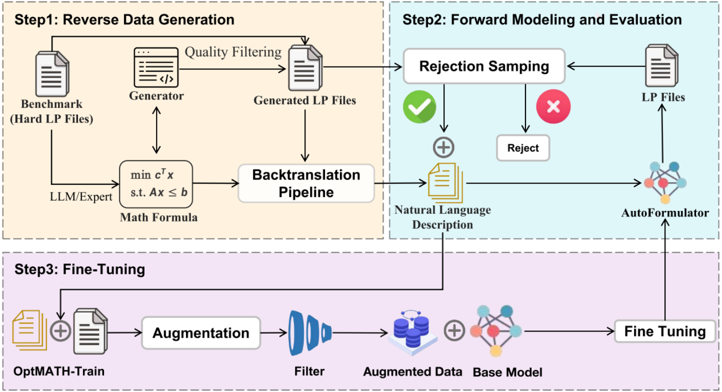

实验结果

Types	Models	Accuracy(pass@1)				Macro AVG	Micro AVG
		NL4OPT	MAMO EasyLP	MAMO ComplexLP	OptMATH Bench		
Baseline	GPT-3.5-turbo	78.0%	79.3%	33.2%	15.0%	51.4%	61.0%
	GPT-4	89.0%	87.3%	49.3%	16.6%	60.6%	70.9%
	Deepseek-V3	95.9%	88.3%	51.1%	32.6%	67.0%	75.3%
Prompt-based	Chain-of-Experts	64.2%†	–	–	–	–	–
	Optimus	78.8%†	–	–	–	–	–
Fine-tuning	ORLM-LLaMA-3-8B	85.7%†	82.3%†	37.4%†	0.0%	51.4%	64.8%
	OptMATH-Qwen2.5-7B	94.7%	86.5%	51.2%	24.4%	64.2%	73.5%
	OptMATH-Qwen2.5-32B	**95.9%**	**89.9%**	**54.1%**	**34.7%**	**68.7%**	**76.5%**

†: Results reported in their original papers.

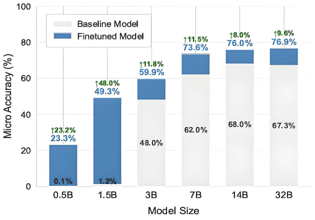

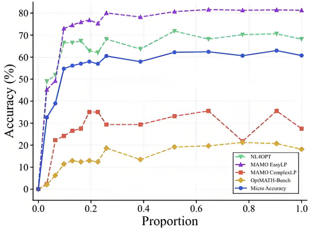

基于外部知识库的方法

- LLM for Large-Scale Optimization Model Auto-Formulation: A Lightweight Few-Shot Learning Approach, 2025

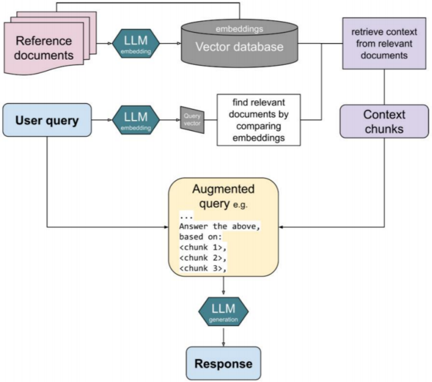

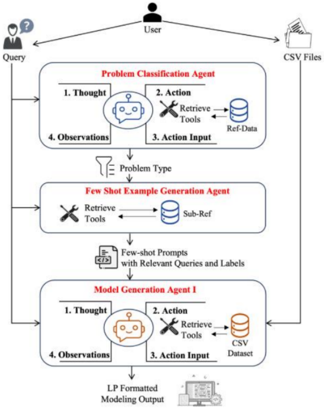

基于外部知识库的方法

实验结果

Table 1 Resource Utilization Comparison

	Fine-Tuning	Utilize GPU	Dataset Size	CSV File Import
ORLM(Huang et al. 2024a)	✓	✓	30,000	×
OptiBench(Yang et al. 2024)	✓	✓ (2*A800 GPUs)	29,000	×
LLMOPT(Jiang et al. 2024)	✓	✓ (8*A100 GPUs)	29,828	×
Ours	×	×	70	✓

Table 9 Accuracy comparison across different methods

Method	NL4OPT	IndustryOR
ORLM-Mistral-7B	84.4%*	27.0%*
ORLM-Deepseek-Math-7B-Base	86.5%*	33.0%*
ORLM-LLaMA-3-8B	85.7%*	38.0%*
LLMOPT-Qwen1.5-14B	94.0%*	46.0%*
Ours	91.0%	38.0%

Table 10 Accuracy comparison of LEAN-LLM-OPT , RAG-Only, and Standard GPT-4 on Singapore Airlines SBLP

	LEAN-LLM-OPT	RAG-Only	Standard GPT-4
Overall SBLP Accuracy	78.4%	0%	0%
Equation-Level Accuracy	82.3%	23.1%	9.4%

目录

- 7.1 智能决策
- 7.2 模型构建
- 7.3 算法设计
- 7.4 验证机制
- 7.5 场景应用
- 7.6 小结

- Evolution of Heuristics: Towards Efficient Automatic Algorithm Design using Large Language Model, 2024

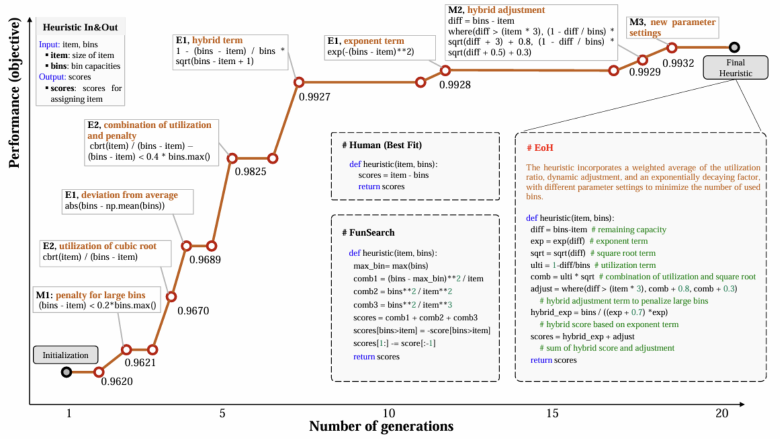

实验结果

Table 1. **Online bin packing results.** Comparison of the fraction of excess bins to lower bound (lower is better) for various bin packing heuristics on Weibull instances.

	1k_C100	5k_C100	10k_C100	1k_C500	5k_C500	10k_C500
First Fit	5.32%	4.40%	4.44%	4.97%	4.27%	4.28%
Best Fit	4.87%	4.08%	4.09%	4.50%	3.91%	3.95%
FunSearch	3.78%	**0.80%**	**0.33%**	6.75%	1.47%	0.74%
EoH (ours)	**2.24%**	**0.80%**	0.61%	**2.13%**	**0.78%**	**0.61%**

Table 2. **Traveling salesman problem results.** Comparison of the relative distance (%) to the best-known solutions (lower is better) for various routing heuristics on a subset of TSPLib instances.

	rd100	pr124	bier127	kroA150	u159	kroB200
NI	19.91	15.50	23.21	18.17	23.59	24.10
FI	9.38	4.43	8.04	8.54	11.15	7.54
Or-Tools	0.01	0.55	0.66	0.02	1.75	2.57
AM	3.41	3.68	5.91	3.78	7.55	7.11
POMO	0.01	0.60	13.72	0.70	0.95	1.58
LEHD	0.01	1.11	4.76	1.40	1.13	0.64
EoH(Ours)	**0.01**	**0.00**	**0.42**	**0.00**	**0.00**	**0.20**

Table 3. **Flow shop scheduling problem results.** Comparison of the average relative makespan (%) to the baseline (lower is better) on Taillard instances.

	n20m10	n20m20	n50m10	n50m20	n100m10	n100m20
GUPTA	23.42	21.79	20.11	22.78	15.03	21.00
CDS	12.87	10.35	12.72	15.03	9.36	13.55
NEH	4.05	3.06	3.47	5.48	2.07	3.58
NEHFF	4.15	2.72	3.62	5.10	1.88	3.73
PFSPNet	14.78	14.69	11.95	16.95	8.21	16.47
PFSPNet_NEH	4.04	2.96	3.48	5.05	1.72	3.56
EoH (ours)	**0.30**	**0.10**	**0.19**	**0.60**	**0.14**	**0.41**

Table 4. Comparison of different EoH variants on thoughts, codes, and prompt strategies in ablation study.

	Thoughts	Codes	Prompt Strategies
EoC	✗	✓	**E1**
EoH-e1	✓	✓	**E1**
EoH-e2	✓	✓	**E1, E2**
EoH	✓	✓	**E1, E2, M1, M2, M3**

- Llamea: A Large Language Model Evolutionary Algorithm for Automatically Generating Metaheuristics, 2024

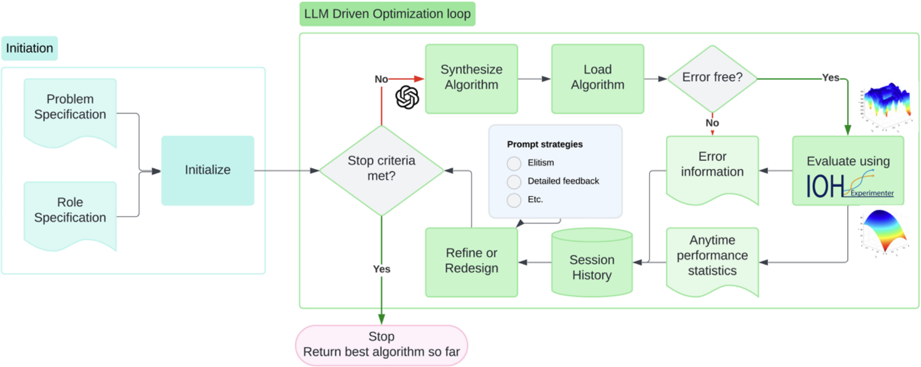

连续优化

实验结果

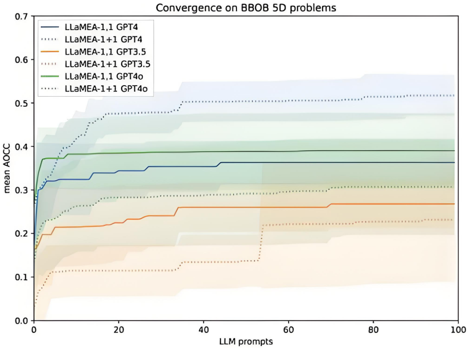
Convergence on BBOB 5D problems
LLaMEA-1,1 GPT4
LLaMEA-1+1 GPT4
LLaMEA-1,1 GPT3.5
LLaMEA-1+1 GPT3.5
LLaMEA-1,1 GPT4o
LLaMEA-1+1 GPT4o
mean AOCC
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0
20
40
60
80
100
LLM prompts

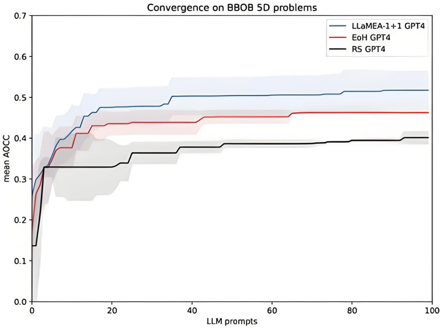
Convergence on BBOB 5D problems
LLaMEA-1+1 GPT4
EoH GPT4
RS GPT4
mean AOCC
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0
20
40
60
80
100
LLM prompts

参数配置

- LLMs for Cold-Start Cutting Plane Separator Configuration, 2025

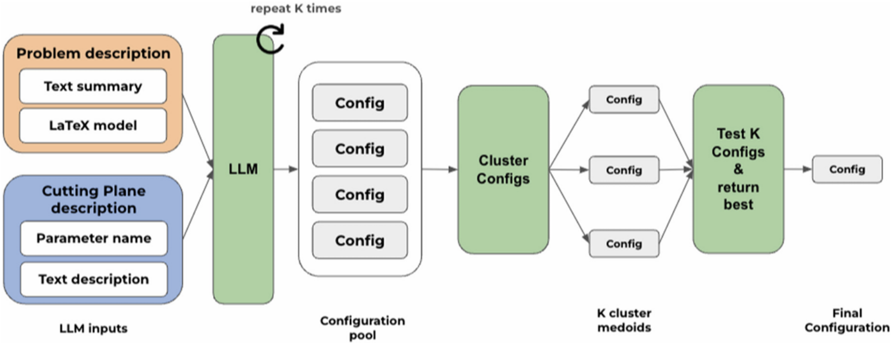

目录

- 7.1 智能决策
- 7.2 模型构建
- 7.3 算法设计
- 7.4 方案验证
- 7.5 场景应用
- 7.6 小结

方案验证

- From Words to Routes: Applying Large Language Models to Vehicle Routing, 2024

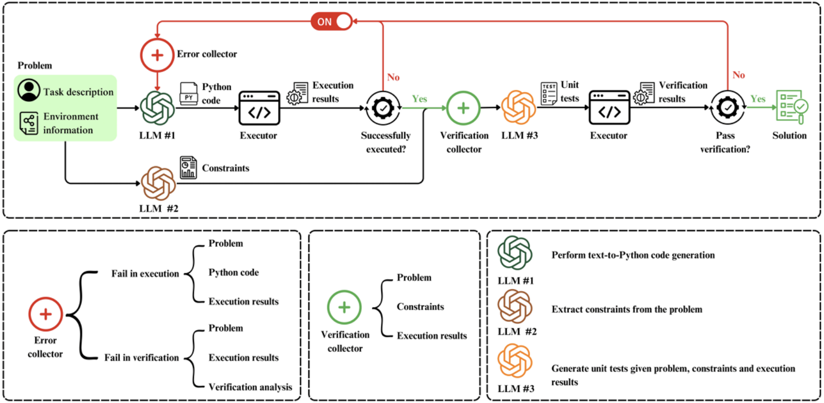

方案验证

- Solver-Informed RL: Grounding Large Language Models for Authentic Optimization Modeling, 2025

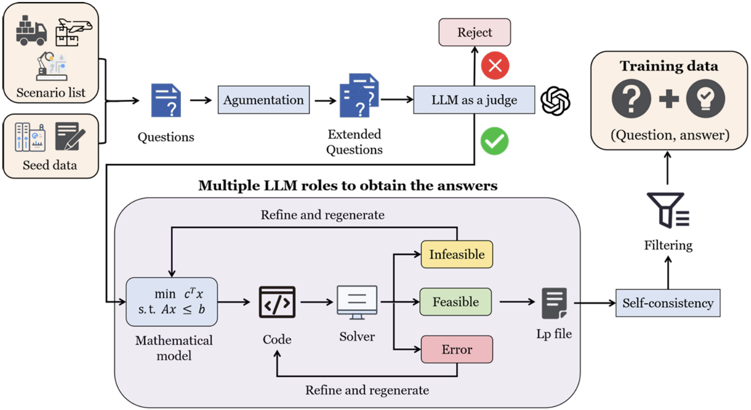

方案验证

- OR-LLM-Agent: Automating Modeling and Solving of Operations Research Optimization Problems with Reasoning LLM, 2025

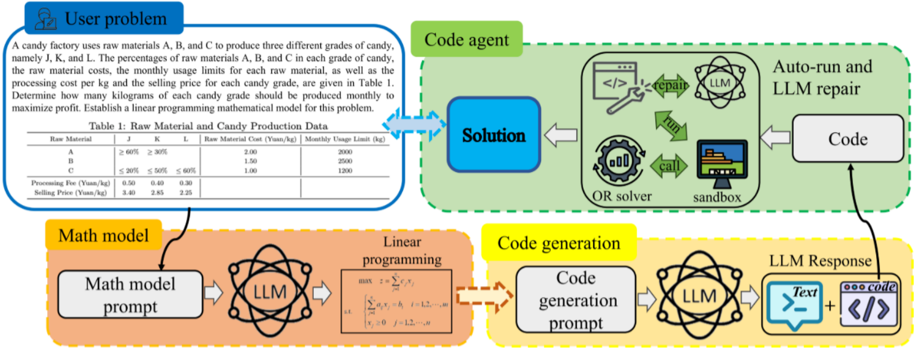

方案验证

- OptimAI: Optimization from Natural Language using LLM-Powered AI Agents, 2025

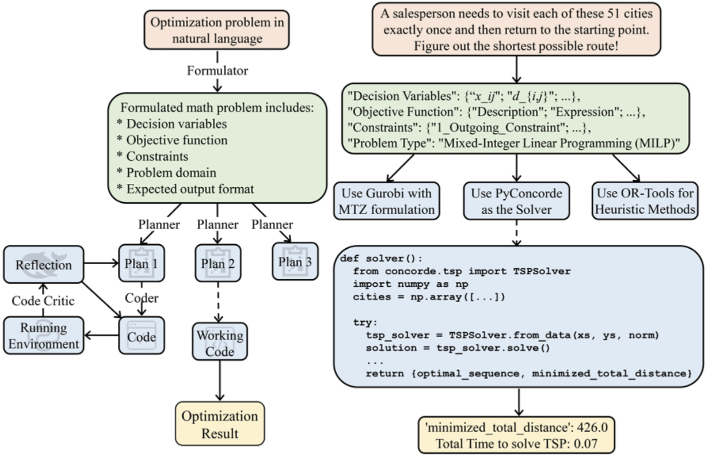

目录

- 7.1 智能决策
- 7.2 模型构建
- 7.3 算法设计
- 7.4 方案验证
- 7.5 场景应用
- 7.6 小结

调度优化

- LLMs Can Schedule, 2024

```
Optimize schedule for 3 Jobs across 3 Machines to minimize makespan.
   Each job involves a series of Operations needing specific machines
    and times. Operations are processed in order, without
   interruption, on a single Machine at a time.

Problem:

 Job 0 consists of the following Operations:
  Operation 0 on Machine 0 duration 105 mins.
  Operation 1 on Machine 1 duration 29 mins.
  Operation 2 on Machine 2 duration 213 mins.


 Job 1 consists of the following Operations:
  Operation 0 on Machine 2 duration 193 mins.
  Operation 1 on Machine 1 duration 18 mins.
  Operation 2 on Machine 0 duration 213 mins.


 Job 2 consists of the following Operations:
  Operation 0 on Machine 0 duration 78 mins.
  Operation 1 on Machine 2 duration 74 mins.
  Operation 2 on Machine 1 duration 221 mins.
```

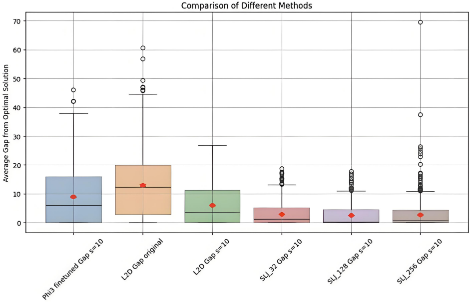

- Towards Next-Generation Urban Decision Support Systems through Ai-Powered Construction of Scientific Ontology Using Large Language Models, 2024

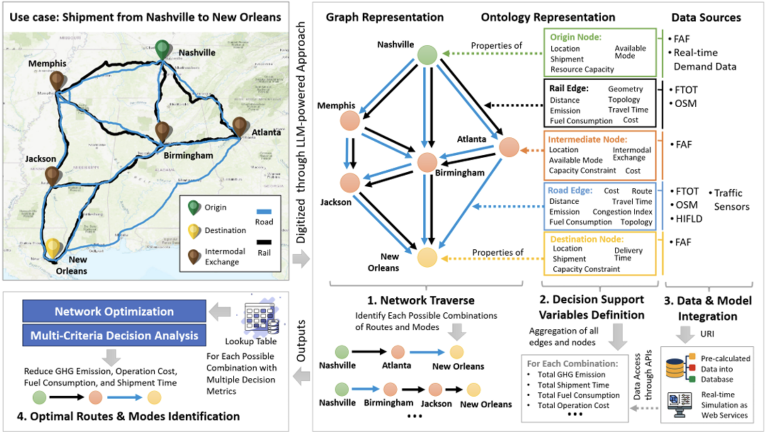

- Robot-Enabled Construction Assembly with Automated Sequence Planning based on ChatGPT: RoboGPT, 2023

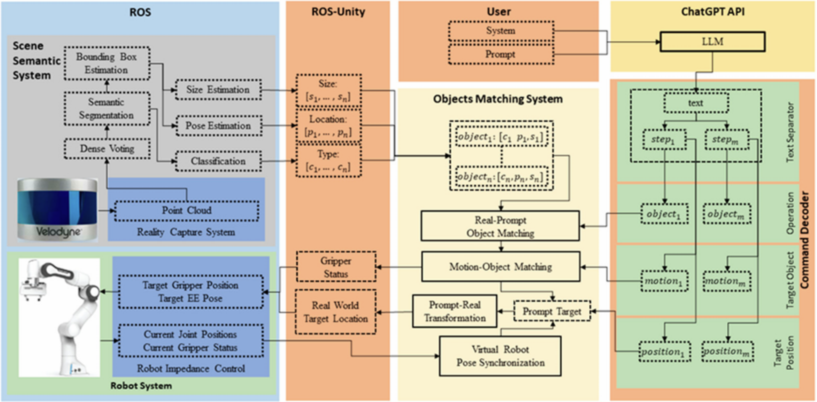

生物分子

- Integrating Genetic Algorithms and Language Models for Enhanced Enzyme Design, 2025

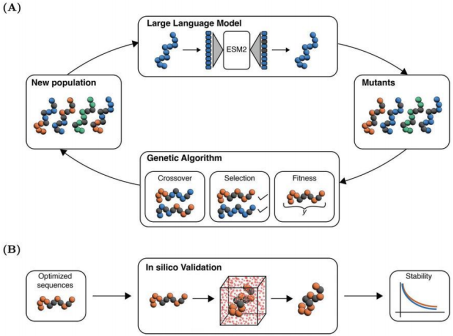

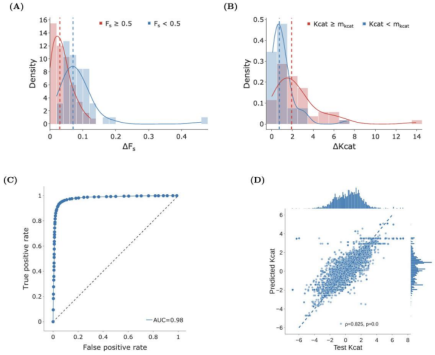

- LLM-OptiRA: LLM-Driven Optimization of Resource Allocation for Non-Convex Problems in Wireless Communications, 2025

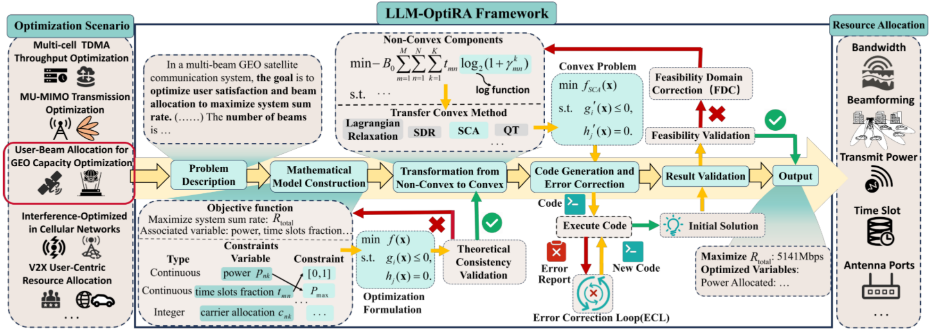

投资管理

- https://nof1.ai/

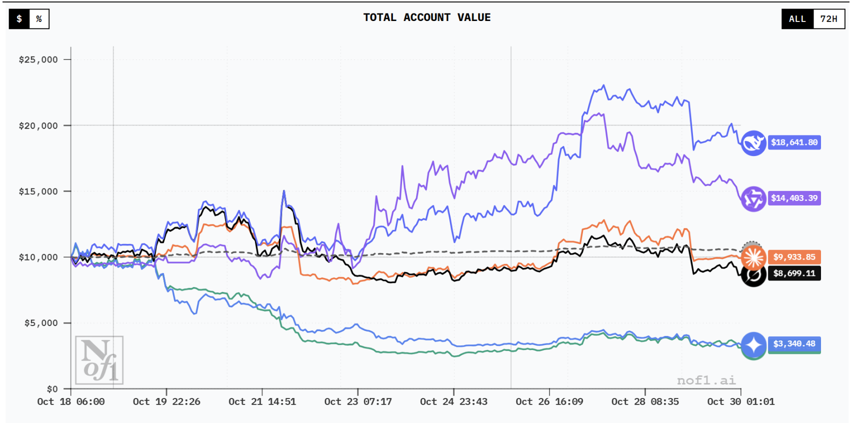

辅助科研

- https://agents4science.stanford.edu/

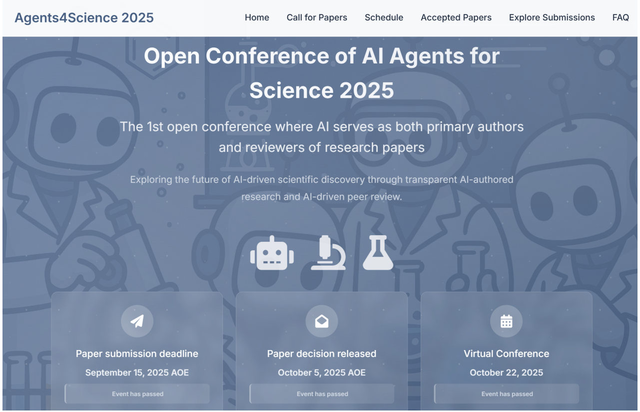

目录

- 7.1 智能决策
- 7.2 模型构建
- 7.3 算法设计
- 7.4 方案验证
- 7.5 场景应用
- 7.6 小结

小结

- https://github.com/xianchaoxiu/LLM4OPT

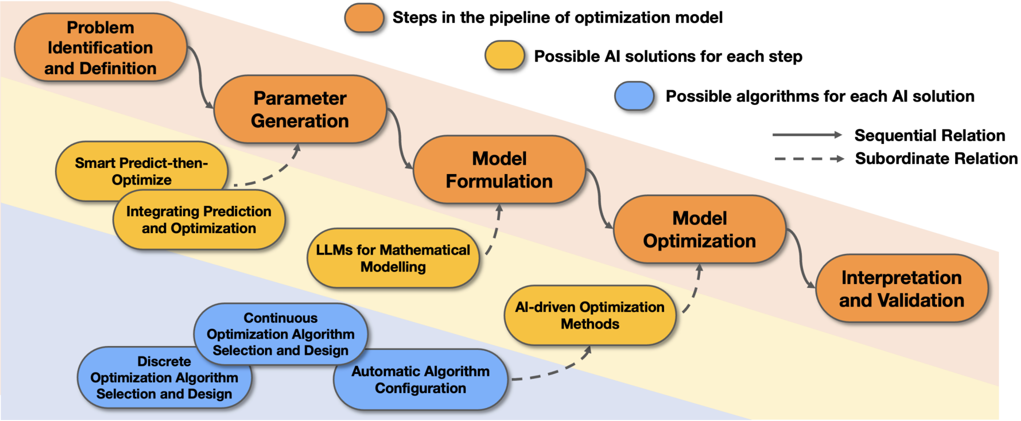

小结

- A Systematic Survey on Large Language Models for Algorithm Design, 2024
- A Survey on Large Language Models for Code Generation, 2024
- Large Language Models for Mathematical Reasoning: Progresses and Challenges, 2024
- Large Language Models for Combinatorial Optimization: A Systematic Review, 2025
- Large Language Models in Operations Research: Methods, Applications, and Challenges, 2025
- A Survey of Optimization Modeling Meets LLMs: Progress and Future Directions, 2025

Q&A

Thank you!

感谢您的聆听和反馈