

SolverLLM: Leveraging Test-Time Scaling for Optimization Problem via LLM-Guided Search

Dong Li¹, Xujiang Zhao², Linlin Yu³, Yanchi Li², Wei Cheng², Zhengzhang Chen²,
Zhong Chen⁴, Feng Chen⁵, Chen Zhao¹, Haifeng Chen²

¹Baylor University, ²NEC Labs America, ³Augusta University,

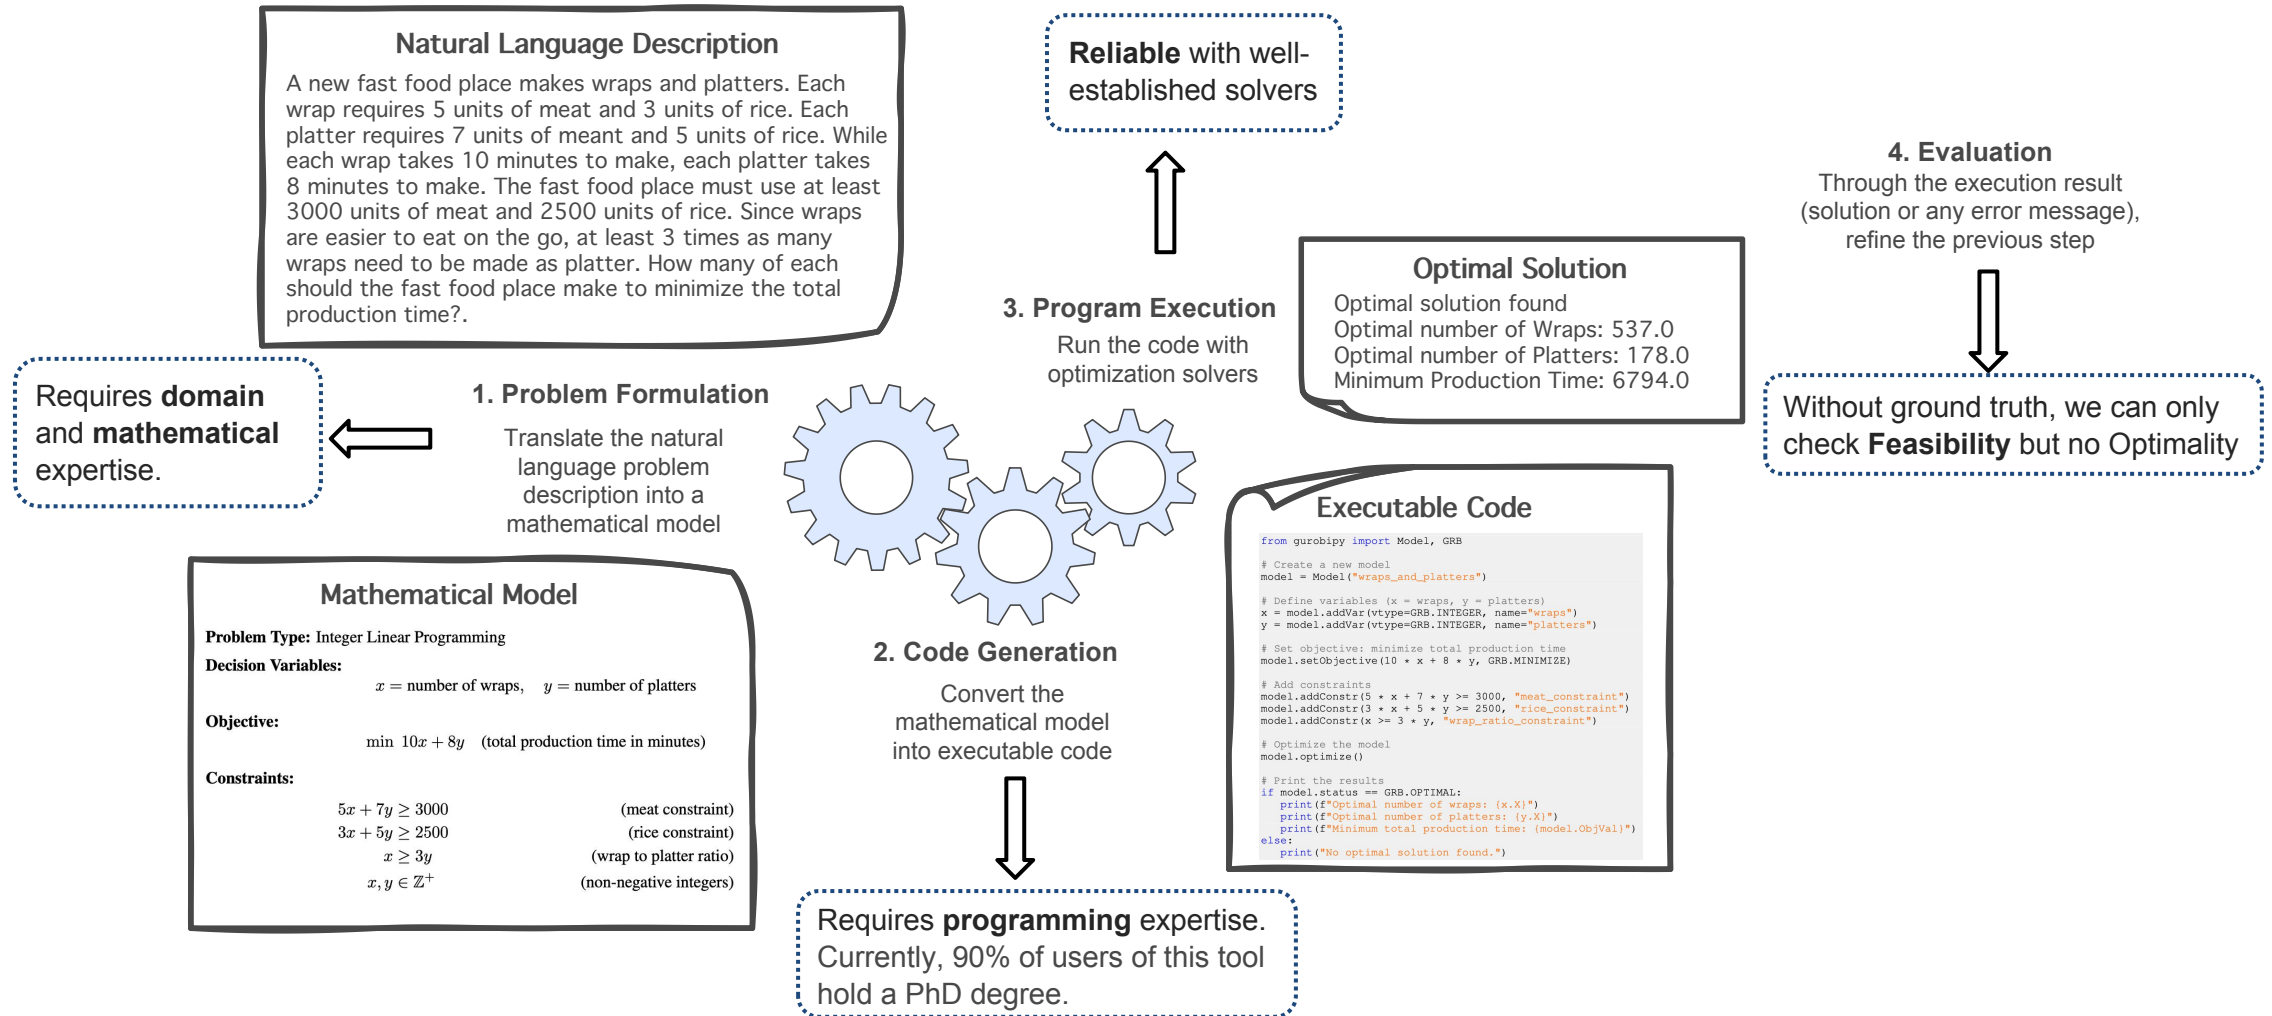
⁴Southern Illinois University, ⁵University of Texas at Dallas



BAYLOR
UNIVERSITY



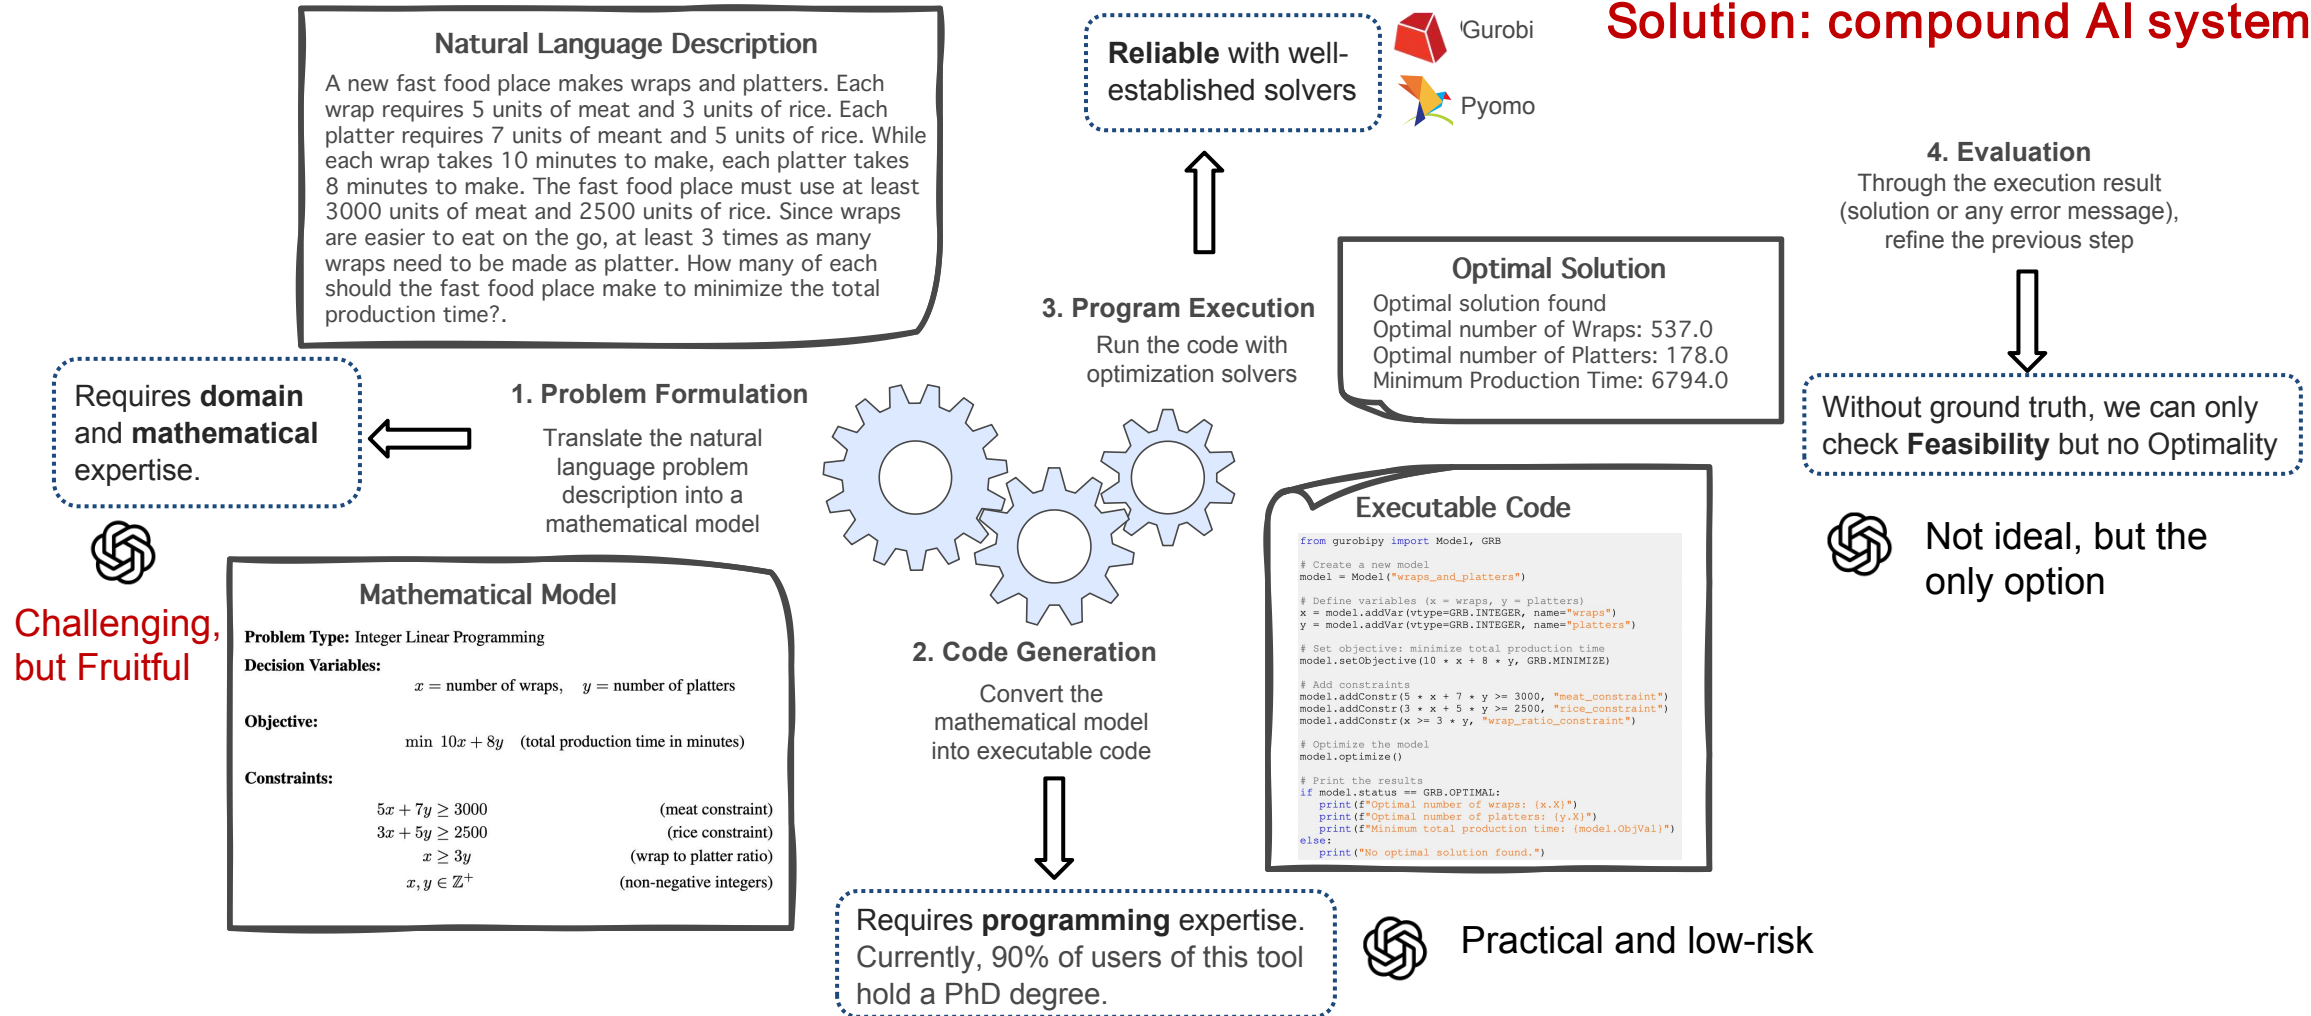
Solving the Optimization Problem



Demands expertise in application domain and mathematical programming

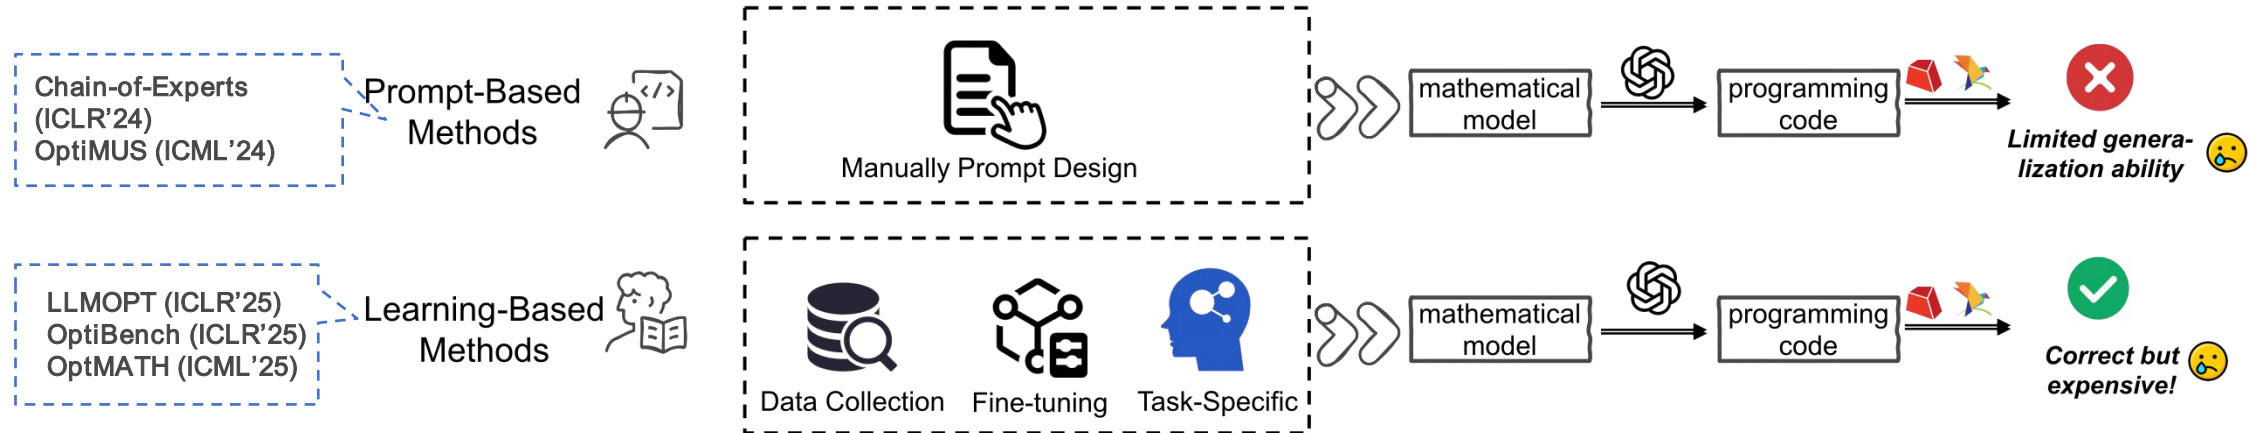
Automating Optimization Problem Solving

Solution: compound AI system



Compound AI: Language Understanding + Reliable Logic Solution

Compound AI System Design (Literature)



Prompt-Based Methods

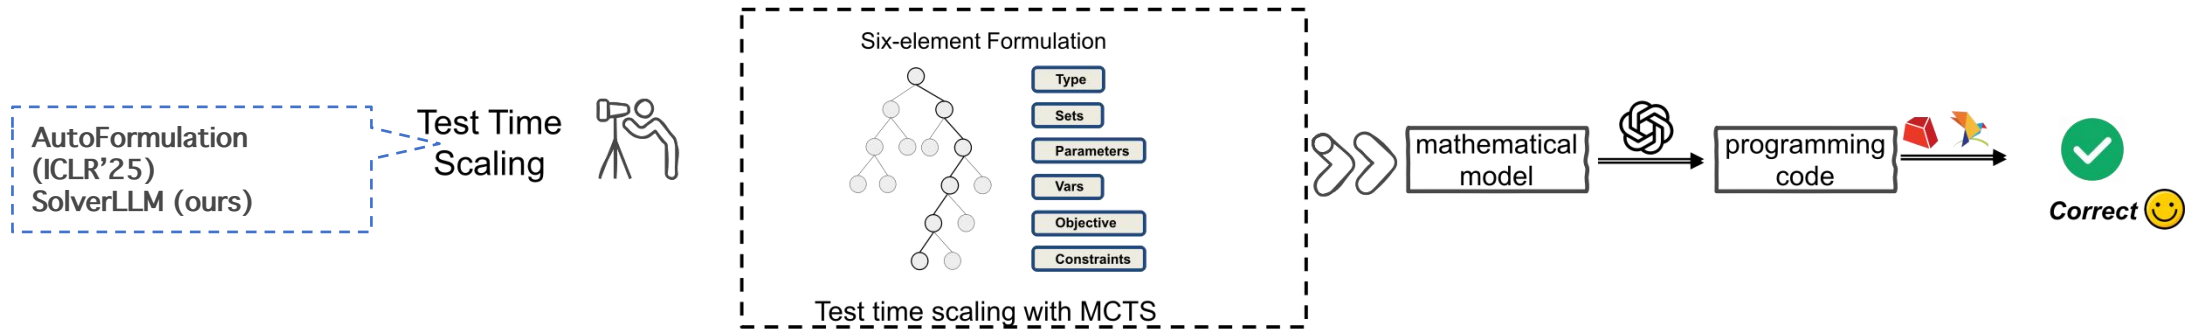
- ❖ Coordinate specialized agents under a carefully designed workflow.
- ❖ But? Sensitive to prompt choice, making them fragile on unfamiliar optimization tasks.

Learning-Based Methods

- ❖ Fine-tune a general LLM on curated problem–solution pairs.
- ❖ But? Effectiveness depends on significant dataset labeling and model fine-tuning cost

Automated optimization problem solving is an active research area


Compound AI System Design: Test Time Scaling (TTS)



No need for training



Optimization problems follow established paradigms, and there is abundant unstructured information on the Internet, which is likely part of LLM training data.



Token consumption and solution exploration are expensive and time-consuming.



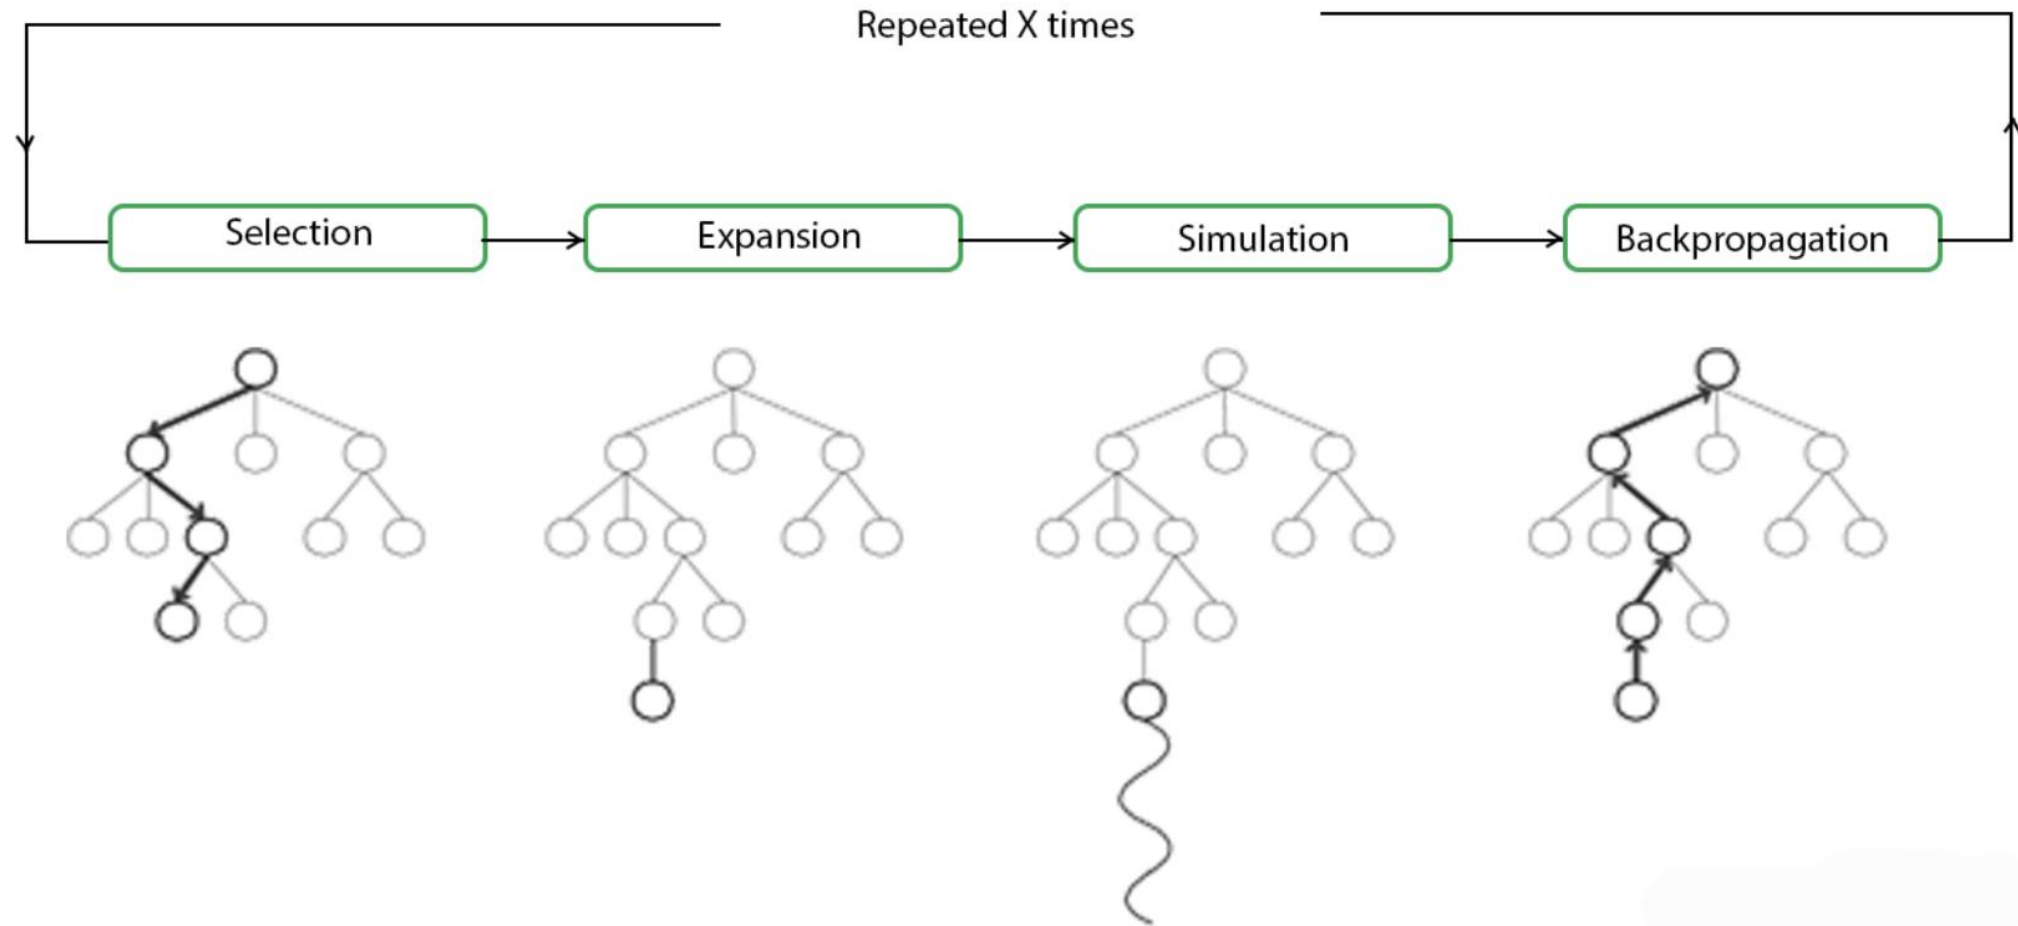
Risk of randomness and hallucination during generation.




Model evaluation and backward reasoning are non-trivial without correctness labels.

TTS is promising, but efficiency and reasoning is critical

Monte Carlo Tree Search (MCTS)

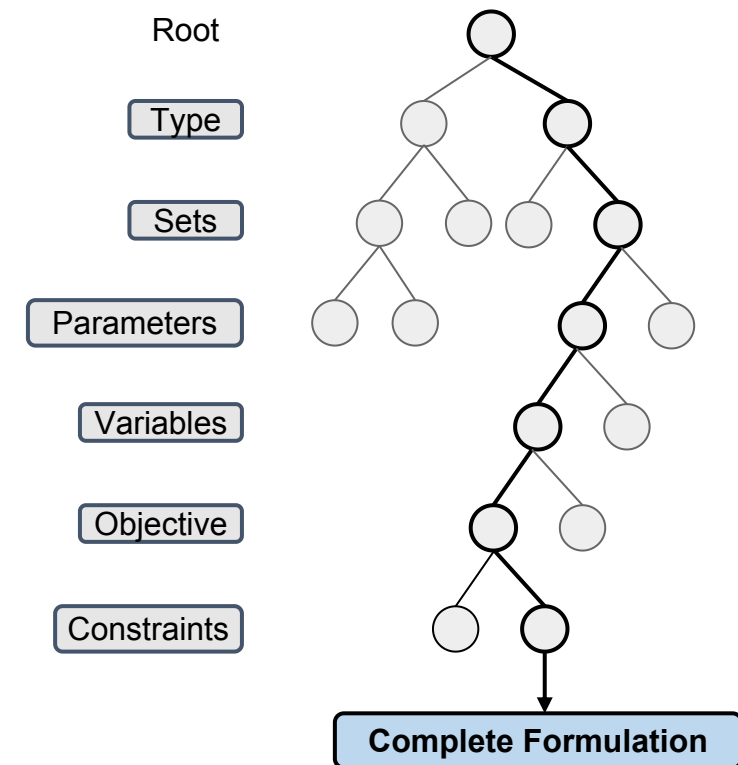


Six-element Schema Formulation for Optimization Problem




Search Tree

- Starting from the root, each level corresponds to an element and is associated with a **local expert-guided knowledge base**, yielding a tree of depth seven.
- Each node maintains its **visit count**, **accumulated reward**, and a **trigger** that indicates whether it is **active**.




SolverLLM: Optimization Formulation via LLM-Guided MCTS

Selection
Select until reaching a leaf
(upper) / **active (lower)** node



Expansion on leaf node

Expansion on non-leaf node



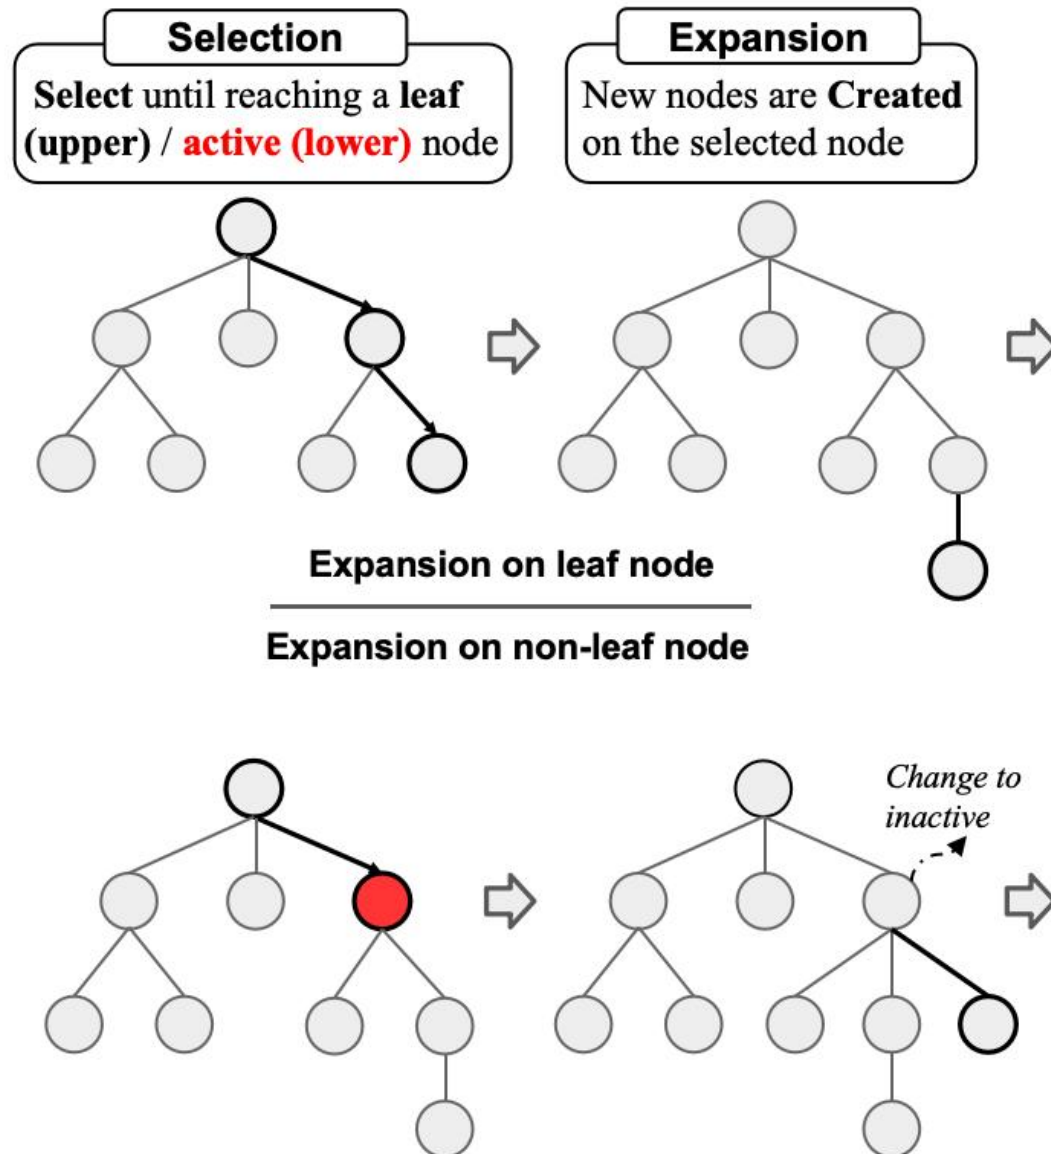
Selection

- Starting from the root node, nodes are selected sequentially downward until a leaf or an **active non-leaf** node is reached.
- The selection is guided by the Upper Confidence Bound for Trees (UCT). Given a parent node s , the next child $s_{child} \in \text{Child}(s)$ is selected according to:

$$s_{child} = \arg \max_{s' \in \text{Child}(s)} \left[Q_{s'} + c \sqrt{\frac{2 \log N_s}{N_{s'}}} \right]$$

average reward ← $Q_{s'}$ $2 \log N_s$ ← *visit count* c ← *exploration factor*

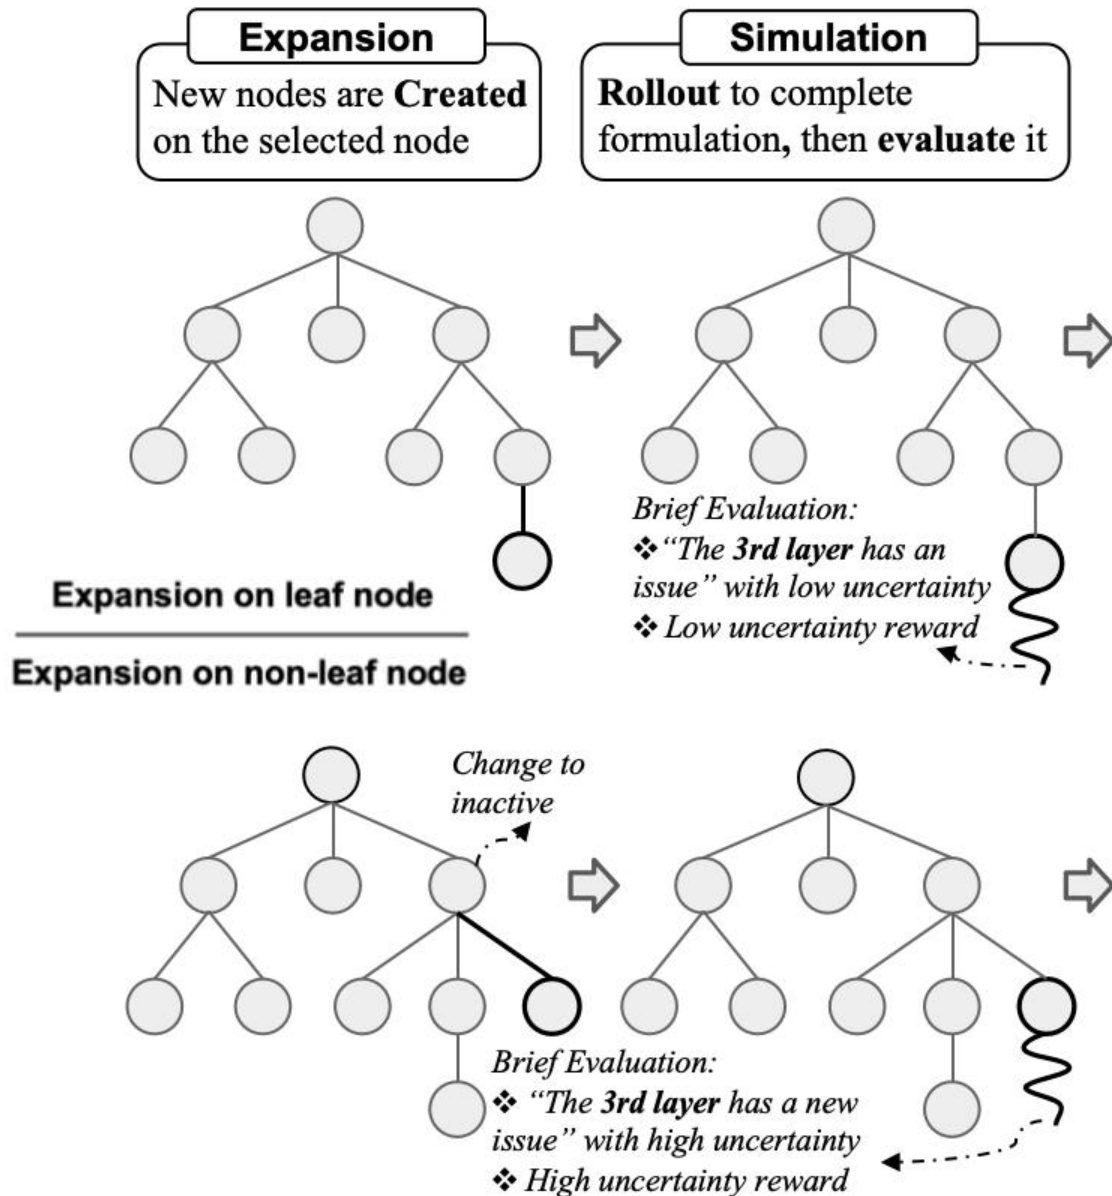
SolverLLM: Optimization Formulation via LLM-Guided MCTS



Dynamic Expansion

- Expand a new node to the selected leaf or **active non-leaf** node.
- LLM-guided expansion with **local expert-guided knowledge base**.

SolverLLM: Optimization Formulation via LLM-Guided MCTS




Simulation

- Expand continues on expanded node until a complete formulation is obtained, after which alignment and evaluation are performed to yield the following signals:
- A **reward** capturing the factors of feasibility, optimality, and error penalty.
- Layer-wise reasoning feedback, which integrates an **activation trigger**, **reasoning guidance**, and **selection rationale** to inform subsequent search and refinement.

SolverLLM: Optimization Formulation via LLM-Guided MCTS

Backpropagation

- Extend standard MCTS with prompt backpropagation and uncertainty propagation to update the search tree and inform future decisions.
- Prompt Backpropagation: The layer-wise reasoning feedback signals are propagated back to each layer, where the **local uncertainty** is computed via predictive entropy derived from the **selection rationale**.
- Uncertainty Backpropagation: The semantic uncertainty at each evaluated node is estimated as the **global uncertainty**, which is then used to update each node's **reward** through an uncertainty-weighted averaging scheme.



Comparison of Solving Accuracy (SA) Among Different Types of Methods

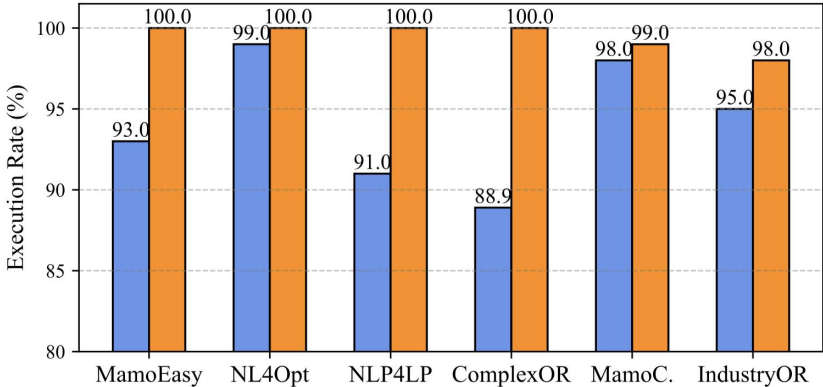
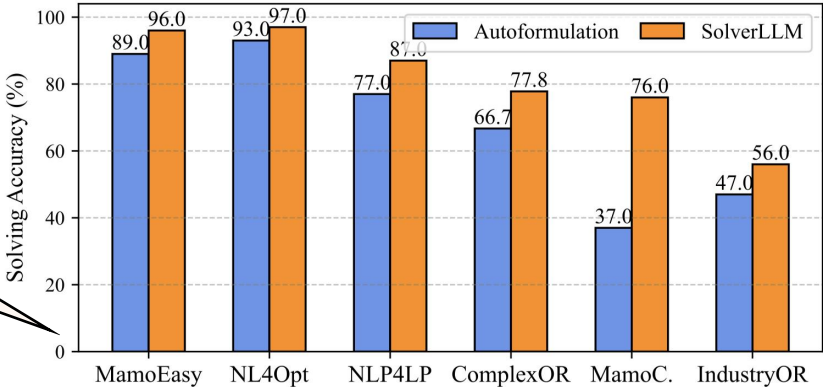
Prompt-Based methods

| | NL4Opt | NLP4LP | ComplexOR |
|-------------------------|--------------|--------------|--------------|
| GPT-4 Directly | 47.3% | 35.8% | 9.5% |
| GPT-4o Directly | <u>81.0%</u> | 32.4% | 27.3% |
| Reflexion | 53.0% | 46.3% | 19.1% |
| Chain-of-Experts | 64.2% | 53.1% | 38.1% |
| OptiMUS | 78.8% | <u>72.0%</u> | <u>66.7%</u> |
| SolverLLM (Ours) | 97.0% | 87.0% | 77.8% |

Learning-Based methods

| | MamoEasy | NL4Opt | MamoComplex | IndustryOR |
|-------------------------|--------------|--------------|--------------|--------------|
| GPT-4 Directly | 66.5% | 47.3% | 14.6% | 28.0% |
| GPT-4o Directly | 91.0% | 81.0% | 34.0% | 34.0% |
| ORLM-Mistral | 81.4% | 84.4% | 32.0% | 27.0% |
| ORLM-Deepseek | 82.2% | 86.5% | 37.9% | 33.0% |
| ORLM-LLaMa3 | 82.3% | 85.7% | 37.4% | 38.0% |
| LLMOPT | 97.0% | <u>93.0%</u> | <u>68.0%</u> | <u>46.0%</u> |
| SolverLLM (Ours) | <u>96.0%</u> | 97.0% | 76.0% | 56.0% |

Test-Time Scaling method



Thanks for Your Wathcing!