ICLR 2026

MetaMuse

Algorithm Generation via Creative Ideation

Ruiying Ma[1] Chieh-Jan Mike Liang[2] Yanjie Gao[2] Francis Y. Yan[3]

[1]UC Berkeley [2]Microsoft Research [3]UIUC

35.76%
Cache Miss Reduction

30.93%
Bin Usage Reduction

1.8×
More Diverse Solutions

1.3×
Higher Entropy

ICLR 2026

The Challenge: LLM Availability Bias

- LLMs have broad knowledge of **existing algorithms** — but struggle to invent genuinely new ones
- They exhibit **availability bias** : solutions cluster around well-known human heuristics
- The bias is consistent across three LLM backbones — an **architectural tendency** , not a model quirk

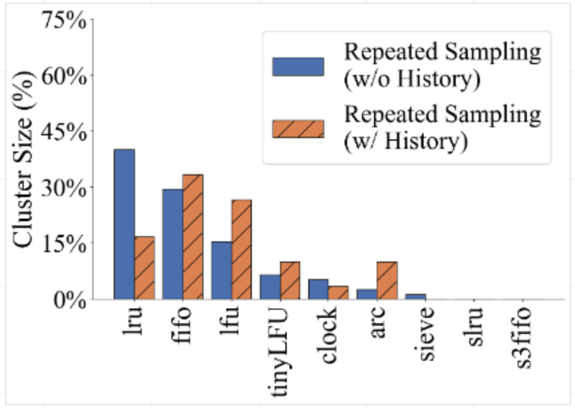

GPT-4o

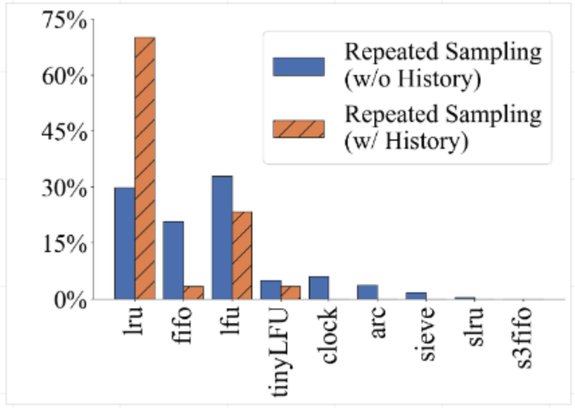

Llama 3.3-70B

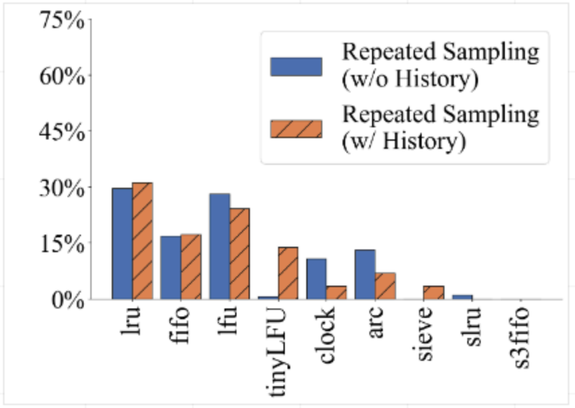

DeepSeek-V3

MetaMuse: Three Self-Reflection Principles

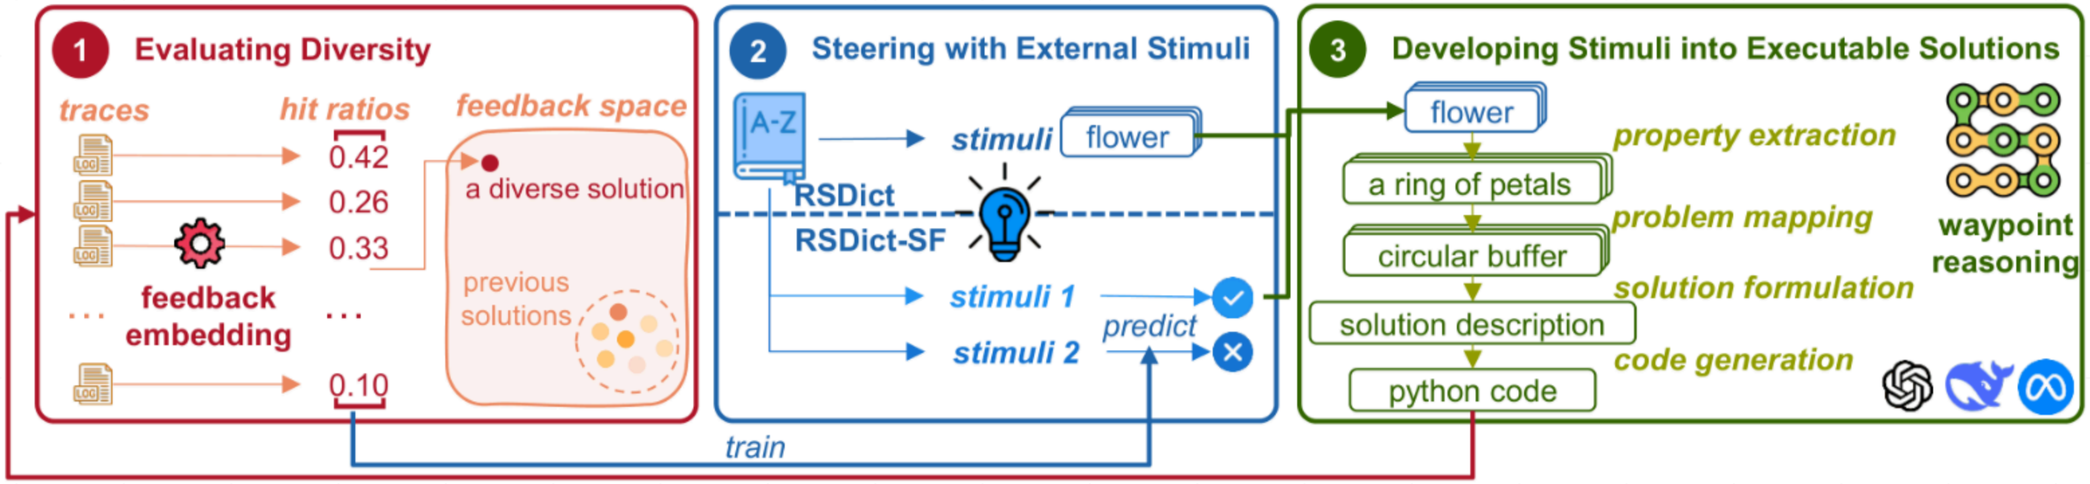

① Performance-Based Diversity Evaluation

Measures diversity in measurable performance space, not abstract idea space

② External Stimuli Steering

Guides LLM ideation via external stimuli rather than internal randomness

③ Waypoint Reasoning

Constructs executable algorithms through structured waypoints, not free-form CoT

Optimization: Learning to Steer Better

Builds on previous solutions to select the most effective stimuli for the next iteration. Uses **Gaussian Process Regression** trained on (stimuli → diversity + performance) pairs from prior runs.

Train

Fit GP models on the history of stimuli sets and the diversity + performance of the solutions they produced

Predict

Generate two candidate stimuli sets; GP predicts expected diversity and usefulness for each

Select

Choose the set that fits the current goal — *exploration* (diversity) or *exploitation* (performance)

Application 1: Cache Replacement

- Setup: 96 real-wrold workloads from 4 scenarios; 14 baselines (5 LLM-driven methods + 9 human heuristics)
- MetaMuse reduces cache miss ratios by **up to 35.76%**

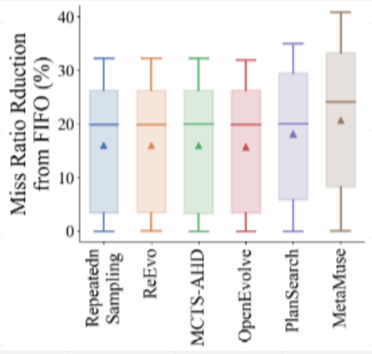

GPT-4o

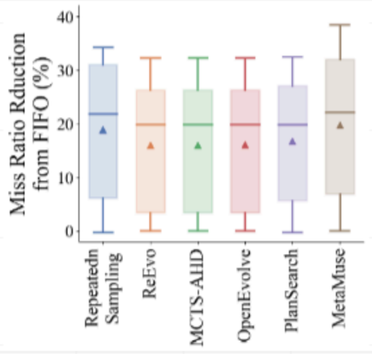

Llama 3.3-70B

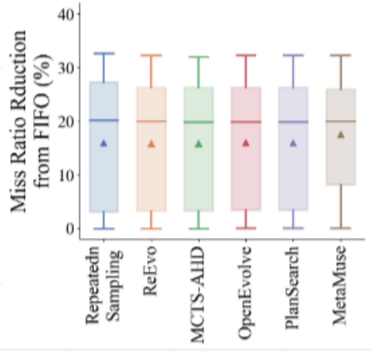

DeepSeek-V3

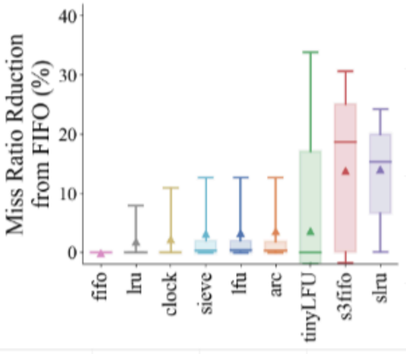

Human Heuristics

Application 2: Online Bin Packing

- Setup: 288 real-wrold workloads from 3 scenarios; 11 baselines (5 LLM-driven methods + 6 human heuristics)
- MetaMuse reduces bin usage by up to 30.93%

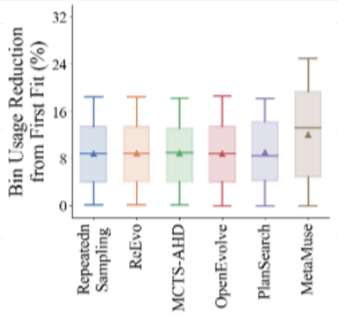

GPT-4o

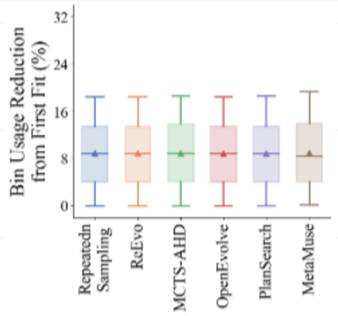

Llama 3.3-70B

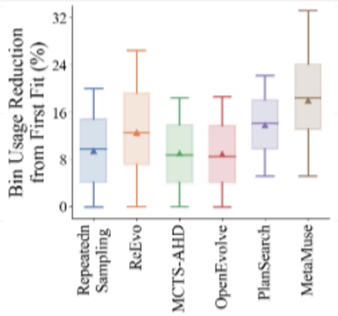

DeepSeek-V3

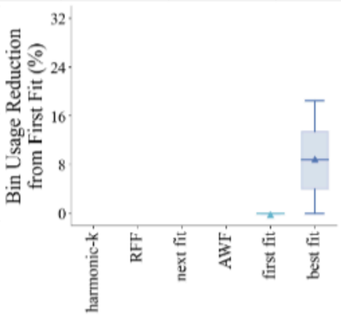

Human Heuristics

BIN PACKING

Less Bias, More Diversity

1.8×
More diverse solutions vs. best baseline

1.3×
Higher solution distribution entropy

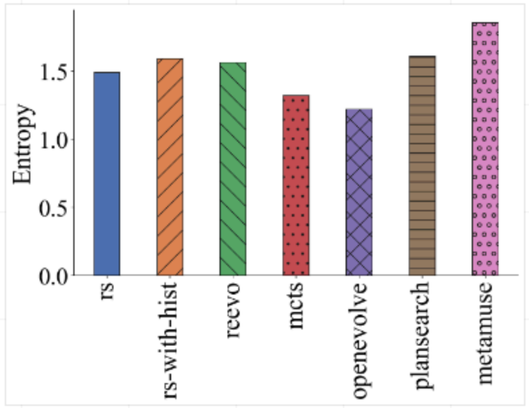

GPT-4o

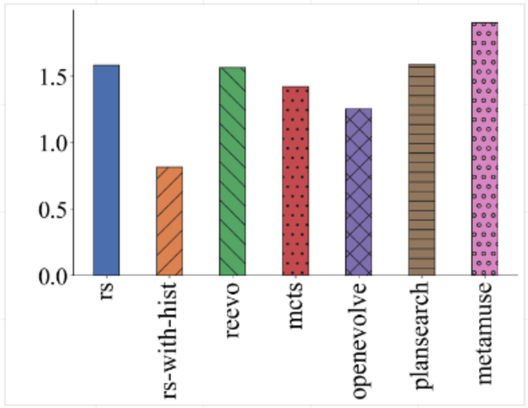

Llama 3.3-70B

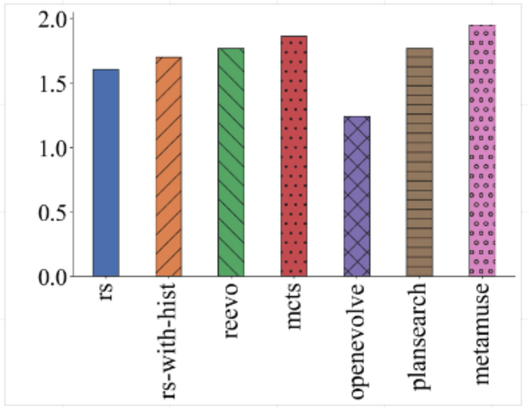

DeepSeek-V3

DIVERSITY

Summary: What MetaMuse Achieves

Bias Mitigation

Three self-reflection principles break LLMs' availability bias and push exploration beyond familiar human heuristics

Superior Diversity

1.8× more diverse solutions with 1.3× higher entropy — broadest solution coverage among all LLM-driven methods

Real-World Performance

Up to 35.76% cache miss reduction and 30.93% bin usage reduction on production cloud workloads

Adaptive Optimization

Using Gaussian Process Regression to balance exploration and exploitation through learned stimuli selection

github.com/illinois-nsai/MetaMuse