

# To Code or not to Code? Adaptive Tool Integration for Math Language Models via Expectation-Maximization

Haozhe Wang<sup>1\*†</sup>, Long Li<sup>2\*</sup>, Chao Qu<sup>2</sup>, Fengming Zhu<sup>1</sup>,  
Weidi Xu<sup>2</sup>, Wei Chu<sup>2</sup>, Fangzhen Lin<sup>1†</sup>

Hong Kong University of Science and Technology<sup>1</sup>, INFLY Tech<sup>2</sup>

## Abstract

Recent advances in mathematical problem-solving with language models (LMs) integrate chain-of-thought (CoT) reasoning and code execution to harness their complementary strengths. However, existing hybrid frameworks exhibit a critical limitation: they depend on externally dictated instructions or rigid code-integration templates, lacking metacognitive awareness – the capacity to dynamically evaluate intrinsic capabilities and autonomously determine when and how to integrate tools. This rigidity motivates our study of autonomous code integration, enabling models to adapt tool-usage strategies as their reasoning abilities evolve during training.

While reinforcement learning (RL) shows promise for boosting LLM reasoning at scale (e.g., DeepSeek-R1), we demonstrate its inefficiency in learning autonomous code integration due to inadequate exploration of the vast combinatorial space of CoT-code interleaving patterns. To address this challenge, we propose a novel Expectation-Maximization (EM) framework that synergizes structured exploration (E-step) with off-policy RL optimization (M-step), creating a self-reinforcing cycle between metacognitive tool-use decisions and evolving capabilities. Experiments reveal our method achieves superior results through improved exploration. Notably, our 7B model improves over 11% on MATH500 and 9.4% on AIME without o1-like CoT. Code, models and data are released on <https://github.com/HaozheH3/AutoCode>.

## 1 Introduction

Large Language Models (LLMs) have demonstrated remarkable performance across various domains (Kaddour et al., 2023; Achiam et al., 2023; Dubey et al., 2024; Team et al., 2023; Yang et al.,

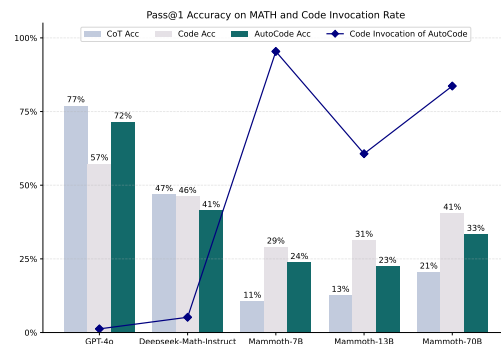


Figure 1: **Analysis of existing tool-integrated math LLMs.** While tool integration can in principle bring complementary benefits to CoT reasoning, existing models show critical rigidity when incorporating code, lacking the metacognitive capacity to earn the synergistic benefits from code integration.

2024a). Yet, solving complex mathematical problems still remains challenging, as the task requires hybrid skills in abstract reasoning, symbolic manipulation, and precise numerical computation (Gao et al., 2023; Yue et al., 2023; Gou et al., 2023; Li, 2024). Current approaches adopt two complementary paradigms: (1) *chain-of-thought (CoT) reasoning*, which decomposes problems into intermediate reasoning steps (Wei et al., 2022; Yu et al., 2023), and (2) *external tool integration*, where models generate code snippets to offload computations to interpreters or symbolic solvers (Toshniwal et al., 2024; Yue et al., 2023). While CoT reasoning excels at semantic parsing and stepwise logic, its reliance on token-level autoregressive generation often propagates numerical errors. Conversely, tool-based approaches ensure computational precision but suffer from a semantic-to-symbolic translation gap, where even minor syntactic errors or contextual misinterpretations disrupt execution (Li, 2024).

Recent hybrid frameworks like Mammoth (Yue et al., 2023), Deepseek-Math (Gou et al., 2023; Shao et al., 2024), and Qwen-2.5-Math (Yang et al., 2024b) attempt to combine these paradigms through interleaved CoT-code reasoning. How-

\*Equal contribution.

†Corresponding author.

ever, as our empirical analysis reveals (Fig. 1), current methods exhibit a critical rigidity: they either default to CoT reasoning unless explicitly prompted for code generation or adhere to static templates for tool invocation. We trace this limitation to prevailing supervised fine-tuning (SFT) paradigms that condition models to (1) passively follow user instructions (e.g., "Let's write a Python program" (Yue et al., 2023)), (2) replicate fixed code-integration patterns from curated datasets (Yang et al., 2024b), or (3) imitate teacher-forced tool-use trajectories (Gou et al., 2023; Shao et al., 2024). Consequently, LLMs lack *metacognitive awareness* – the capacity to dynamically evaluate their intrinsic capabilities against problem contexts and autonomously determine when and how to integrate tools. This deficiency motivates our central research question:

*How can mathematical LLMs learn autonomous code integration (AutoCode) that optimally complements their inherent reasoning capabilities?*

Reinforcement learning (RL) offers a promising pathway by optimizing policies through self-generated trajectories, as evidenced by recent successes like DeepSeek R1 (Guo et al., 2025). However, we empirically observe that standard RL methods is inefficient in learning autonomous code integration (*AutoCode*) strategies (see Sec. 4.2). This stems from RL's tendency to exploit local policy neighborhoods, thereby insufficiently exploring the vast combinatorial space of potential CoT-code interleaving patterns. Such myopic exploration constrains the discovery of high-reward reasoning paths that judiciously blend both modalities, particularly as the model's reasoning capabilities evolve during training.

To address this challenge, we propose a novel Expectation-Maximization (EM) framework that synergizes guided exploration with policy optimization. Our key innovation lies in formulating code-integration decisions as latent variables within an EM paradigm, creating a self-reinforcing cycle: the E-step identifies high-potential code-integration decisions through guided exploration, while the M-step optimizes policy parameters for joint metacognitive tool-usage and reasoning.

This dual mechanism enables models to adapt tool-use strategies as their capabilities evolve during training. Practically, we achieve efficiency through two design choices: (1) an offline data

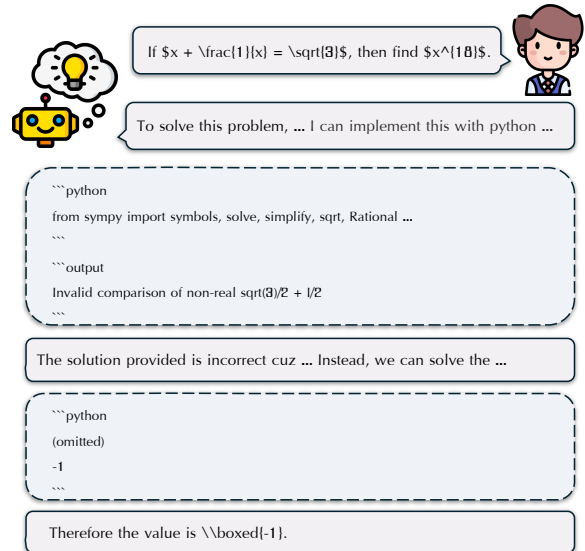


Figure 2: **Example of Autonomous Code Integration.** We aim to enable LLMs to determine tool-usage strategies based on their own capability boundaries. In the example, the model write code to solve the problem that demand special tricks, strategically bypassing its inherent limitations.

curation step (E-step) that prioritizes high-return code invocation decisions through guided exploration, and (2) an off-policy RL optimization step (M-step) that jointly improves tool-usage and reasoning. This approach offers enhanced control and efficiency compared to standard RL, which is particularly beneficial for resource-constrained companies or researchers.

Extensive experiments demonstrate that our method (a) preserves higher training efficiency while achieving better performance, and (b) learns intelligent code integration strategies that achieves higher accuracy than either CoT or code prompted in isolation. Notably, our show consistent improvements across different benchmarks, raising MATH500 from 60.4% to 71.4%.

Our contribution is summarized as follows: (1) We diagnose a critical gap in mathematical LLM – the inability to autonomously integrate tools based on metacognitive awareness – and demonstrate standard RL's inefficiency in addressing it. (2) We propose a novel EM-based framework that jointly adapts the tool-usage strategies with evolving reasoning abilities, with a simple yet efficient implementation. (3) We demonstrate superior results in both training efficiency and accuracy on challenging benchmarks.

## 2 Background

**Problem Statement.** Modern tool-augmented language models address mathematical problems  $x_q \in \mathcal{X}_Q$  by generating step-by-step solutions that interleave natural language reasoning with executable Python code (Fig. 2). Formally, given a problem  $x_q$ , a model  $\mathcal{M}_\theta$  iteratively constructs a solution  $y_a = \{y_1, \dots, y_T\}$  by sampling components  $y_t \sim p(y_t | y_{<t}, x_q)$ , where  $y_{<t}$  encompasses both prior reasoning steps, code snippets and execution results  $\mathbf{e}_t$  from a Python interpreter. The process terminates upon generating an end token, and the solution is evaluated via a binary reward  $r(y_a, x_q) = \mathbb{I}(y_a \equiv y^*)$  indicating equivalence to the ground truth  $y^*$ . The learning objective is formulated as:

$$\max_{\theta} \mathbb{E}_{x_q \sim \mathcal{X}_Q} [r(y_a, x_q)]$$

**Challenge and Motivation.** Developing autonomous code integration (AutoCode) strategies poses unique challenges, as optimal tool-usage behaviors must dynamically adapt to a model’s intrinsic capabilities and problem-solving contexts. While traditional supervised fine-tuning (SFT) relies on imitation learning from expert demonstrations, this paradigm fundamentally limits the emergence of self-directed tool-usage strategies. Unfortunately, current math LLMs predominantly employ SFT to orchestrate tool integration (Yue et al., 2023; Gou et al., 2023; Shao et al., 2024; Li, 2024), their rigid adherence to predefined reasoning templates therefore struggles with the dynamic interplay between a model’s evolving problem-solving competencies and the adaptive tool-usage strategies required for diverse mathematical contexts.

Reinforcement learning (RL) offers a promising alternative by enabling trial-and-error discovery of autonomous behaviors. Recent work like DeepSeek-R1 (Guo et al., 2025) demonstrates RL’s potential to enhance reasoning without expert demonstrations. However, we observe that standard RL methods (e.g., PPO (Schulman et al., 2017)) suffer from a critical inefficiency (see Sec. 4.2): Their tendency to exploit local policy neighborhoods leads to insufficient exploration of the vast combinatorial space of code-integrated reasoning paths, especially when only given a terminal reward in mathematical problem-solving.

To bridge this gap, we draw inspiration from human metacognition – the iterative process where

learners refine tool-use strategies through deliberate exploration, outcome analysis, and belief updates. A novice might initially attempt manual root-finding via algebraic methods, observe computational bottlenecks or inaccuracies, and therefore prompting the usage of calculators. Through systematic reflection on these experiences, they internalize the contextual efficacy of external tools, gradually forming stable heuristics that balance reasoning with judicious tool invocation.

To this end, *our focus diverges from standard agentic tool-use frameworks (Yuan et al., 2025)*, which merely prioritize successful tool execution. Instead, *we aim to instill human-like metacognition in LLMs, enabling them to (1) determine tool-usage based on their own capability boundaries (see the analysis in Sec. 4.2), and (2) dynamically adapt tool-usage strategies as their reasoning abilities evolve (via our EM framework).*

## 3 Methodology

Inspired by human metacognitive processes, we introduce an Expectation-Maximization (EM) framework that trains LLMs for autonomous code integration (AutoCode) through alternations (Fig. 3):

1. *Guided Exploration (E-step)*: Identifies high-potential code-integrated solutions by systematically probing the model’s inherent capabilities.
2. *Self-Refinement (M-step)*: Optimizes the model’s tool-usage strategy and chain-of-thought reasoning using curated trajectories from the E-step.

### 3.1 The EM Framework for AutoCode

A central challenge in AutoCode lies in the code triggering decisions, represented by the binary decision  $c \in \{0, 1\}$ . While supervised fine-tuning (SFT) suffers from missing ground truth for these decisions, standard reinforcement learning (RL) struggles with the combinatorial explosion of code-integrated reasoning paths. Our innovation bridges these approaches through systematic exploration of both code-enabled ( $c = 1$ ) and non-code ( $c = 0$ ) solution paths, constructing reference decisions for policy optimization.

We formalize this idea within a maximum likelihood estimation (MLE) framework. Let  $P(r = 1|x_q; \theta)$  denote the probability of generating a correct response to query  $x_q$  under model  $\mathcal{M}_\theta$ . Our objective becomes:

$$\mathcal{J}_{\text{MLE}}(\theta) \doteq \log P(r = 1|x_q; \theta) \quad (1)$$

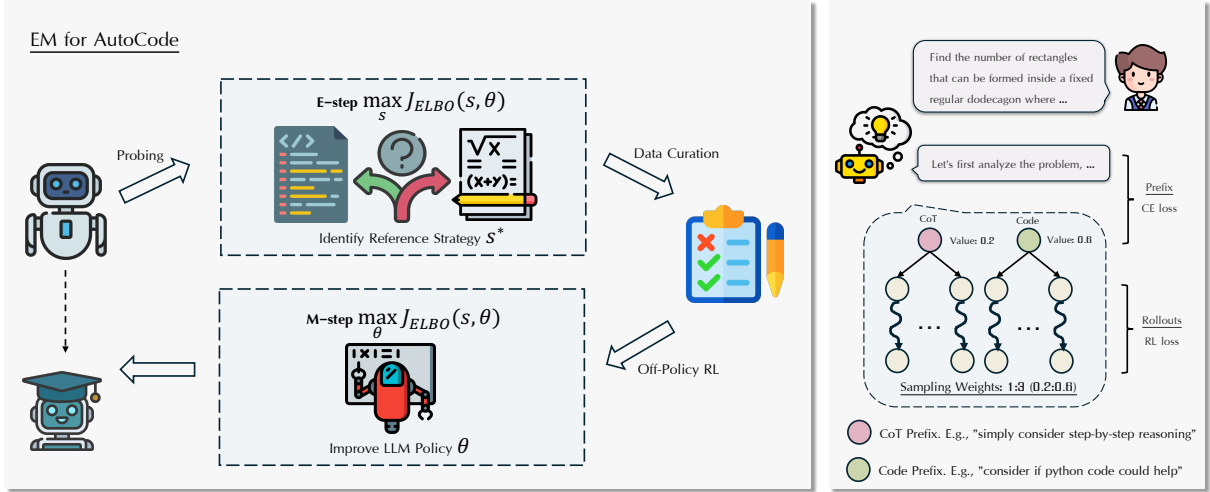


Figure 3: **Method Overview.** (Left) shows an overview for the EM framework, which alternates between finding a reference strategy for guided exploration (E-step) and off-policy RL (M-step). (Right) shows the data curation for guided exploration. We generate  $K$  rollouts, estimate values of code-triggering decisions and subsample the initial data with sampling weights per Eq. 5.

This likelihood depends on two latent factors: (1) the code triggering decision  $\pi_\theta(c|x_q)$  and (2) the solution generation process  $\pi_\theta(y_a|x_q, c)$ . Here, for notation-wise clarity, we consider code-triggering decision at a solution’s beginning ( $c$  following  $x_q$  immediately). We show generalization to mid-reasoning code integration in Sec. 3.2.

The EM framework provides a principled way to optimize this MLE objective in the presence of latent variables (Bishop and Nasrabadi, 2006). We derive the evidence lower bound (ELBO):  $\mathcal{J}_{ELBO}(s, \theta) \doteq$

$$\mathbb{E}_{s(c|x_q)} \left[ \log \frac{\pi_\theta(c|x_q) \cdot P(r=1|c, x_q; \theta)}{s(c|x_q)} \right] \quad (2)$$

where  $s(c|x_q)$  serves as a surrogate distribution approximating optimal code triggering strategies. It is also considered as the reference decisions for code integration.

**E-step: Guided Exploration** computes the reference strategy  $s(c|x_q)$  by maximizing the ELBO, equivalent to minimizing the KL-divergence:  $\max_s \mathcal{J}_{ELBO}(s, \theta) =$

$$-D_{KL}(s(c|x_q) || P(r=1, c|x_q; \theta)) \quad (3)$$

The reference strategy  $s(c|x_q)$  thus approximates the posterior distribution over code-triggering decisions  $c$  that maximize correctness, i.e.,  $P(r=1, c|x_q; \theta)$ . Intuitively, it guides exploration by prioritizing decisions with high potential: if decision  $c$  is more likely to lead to correct solutions, the reference strategy assigns higher

probability mass to it, providing guidance for the subsequent RL procedure.

**M-step: Self-Refinement** updates the model parameters  $\theta$  through a composite objective:

$$\max_\theta \mathcal{J}_{ELBO}(s, \theta) = \mathbb{E}_{\substack{c \sim s(c|x_q) \\ y_a \sim \pi_\theta(y_a|x_q, c)}} \left[ r(x_q, y_a) - \mathcal{CE}(s(c|x_q) || \pi_\theta(c|x_q)) \right] \quad (4)$$

The first term implements reward-maximizing policy gradient updates for solution generation, while while the second aligns native code triggering with reference strategies through cross-entropy minimization (see Fig. 3 for an illustration of the optimization). This dual optimization jointly enhances both tool-usage policies and reasoning capabilities.

### 3.2 Practical Implementation

In the above EM framework, we alternate between finding a reference strategy  $s$  for code-triggering decisions in the E-step, and perform reinforcement learning under the guidance from  $s$  in the M-step. We implement this framework through an iterative process of offline data curation and off-policy RL. **Offline Data Curation.** We implement the E-step through Monte Carlo rollouts and subsampling. For each problem  $x_q$ , we estimate the reference strategy as an energy distribution:

$$s^*(c|x_q) = \frac{\exp(\alpha \cdot \pi_\theta(c|x_q) Q(x_q, c; \theta))}{Z(x_q)} \quad (5)$$

where  $Q(x_q, c; \theta)$  estimates the expected value through  $K$  rollouts per decision,  $\pi_\theta(c|x_q)$  represents the model’s current prior and the  $Z(x_q)$  is

the partition function to ensure normalization. Intuitively, the strategy will assign higher probability mass to the decision  $c$  that has higher expected value  $Q(x_q, c; \theta)$  meanwhile balancing its intrinsic preference  $\pi_\theta(c|x_q)$ .

Our curation pipeline proceeds through:

- Generate  $K$  rollouts for  $c = 0$  (pure reasoning) and  $c = 1$  (code integration), creating candidate dataset  $\mathcal{D}$ .
- Compute  $Q(x_q, c)$  as the expected success rate across rollouts for each pair  $(x_q, c)$ .
- Subsample  $\mathcal{D}_{\text{train}}$  from  $\mathcal{D}$  using importance weights according to Eq. 5.

To explicitly probe code-integrated solutions, we employ prefix-guided generation – e.g., prepending prompts like “Let’s first analyze the problem, then consider if python code could help” – to bias generations toward free-form code-reasoning patterns.

This pipeline enables guided exploration by focusing on high-potential code-integrated trajectories identified by the reference strategy, contrasting with standard RL’s reliance on local policy neighborhoods. As demonstrated in Sec. 4.2, this strategic data curation significantly improves training efficiency by shaping the exploration space.

**Off-Policy RL.** To mitigate distributional shifts caused by mismatches between offline data and the policy, we optimize a clipped off-policy RL objective. The refined M-step (Eq. 4) becomes:

$$\mathbb{E}_{(x_q, y_a)} \left[ \text{clip} \left( \frac{\pi_\theta(y_a|x_q)}{\pi_{\text{ref}}(y_a|x_q)}, 1 - \epsilon, 1 + \epsilon \right) \cdot A \right] - \mathbb{E}_{(x_q, c)} \left[ \log \pi_\theta(c|x_q) \right] \quad (6)$$

where  $(x_q, c, y_a)$  is sampled from the dataset  $\mathcal{D}_{\text{train}}$ . The importance weight  $\frac{\pi_\theta(y_a|x_q)}{\pi_{\text{ref}}(y_a|x_q)}$  accounts for off-policy correction with PPO-like clipping. The advantage function  $A(x_q, y_a)$  is computed via query-wise reward normalization (Schulman et al., 2017).

**Generalizing to Mid-Reasoning Code Integration.** Our method extends to mid-reasoning code integration by initiating Monte Carlo rollouts from partial solutions  $(x_q, y_{<t})$ . Notably, we observe emergence of mid-reasoning code triggers after initial warm-up with prefix-probed solutions. Thus, our implementation requires only two initial probing strategies: explicit prefix prompting for code integration and vanilla generation for pure reasoning, which jointly seed diverse mid-reasoning code usage in later iterations.

## 4 Experiments

Our experiments investigate three key research questions:

*Q1: Method Effectiveness.* How does our approach enhance performance across both in-domain and out-of-domain mathematical benchmarks compared to existing math LLMs?

*Q2: Baseline Comparisons.* How does our method compare to standard RL and SFT baselines in terms of training efficiency and exploration patterns?

*Q3: AutoCode Analysis.* What strategies does the model learn for code integration, and how do these strategies contribute to performance gains?

**Datasets and Benchmarks.** Our method only requires a query set for training. We collect public available queries from MATH (Hendrycks et al., 2021) and Numina (Li et al., 2024), and sample  $7K$  queries based on difficulties. We upload the collected data to the anonymous repo. For evaluation, we employ: GSM8k (Cobbe et al., 2021), MATH500 (Hendrycks et al., 2021), Gaokao-Math2023 (Liao et al., 2024), OlympiadBench (He et al., 2024), the American Invitational Mathematics Examination (AIME24), and the American Mathematics Competitions (AMC23). This benchmark suite spans elementary to Olympiad-level mathematics. We adopt Pass@1 accuracy (Chen et al., 2021; Guo et al., 2025) as our primary metric, using evaluation scripts from DeepseekMath (Shao et al., 2024) and Qwen2Math (Yang et al., 2024a). For competition-level benchmarks (AIME/AMC), we use 64 samples with temperature 0.6 following Deepseek R1 protocols.

**Baselines and Implementation.** We compare against three model categories:

- Proprietary models: o1 (Jaech et al., 2024), GPT-4 (Achiam et al., 2023) and Claude (Anthropic, 2024)
- Recent math-specialized LMs: NuminaMath (Li et al., 2024), Mathstral (AI, 2024), Mammoth (Yue et al., 2023), ToRA (Gou et al., 2023), DartMath (Tong et al., 2024). We do not compare with models that rely on test-time scaling, such as MCTS or long CoT.
- Foundation models enhanced with our method: Qwen2Math (Yang et al., 2024a), DeepseekMath (Shao et al., 2024) and Qwen-2.5 (Yang et al., 2024b).

Our implementation uses  $K = 8$  rollouts per query (temperature=1.0, top-p=0.9). Training completes in about 10 hours on  $8 \times$  A100 (80GB) GPUs

Table 1: **Public SFT Data Used in this Work.** We collect public query set for AutoCode Training. After Deduplication, the total amount of query used is 119K. If the base model is not trained to write code for math problems, we use the SFT annotations associated with the above queries.

Dataset	Size	Unique Queries	CoT Solutions	Code Solutions
Openmath (Toshniwal et al., 2024)	129917	70002	25116	104801
Math-Instruct (Yue et al., 2023)	237781	219607	188644	49137
Metamath (Yue et al., 2023)	285000	161337	285000	0
MMOS (Chen et al., 2024)	134610	69007	0	134610

across three epochs of 7K queries. We list the collected public SFT data in Tab. 1.

#### 4.1 Main Results

Notably, we observe a minimum performance gain of 11% on the MATH500 benchmark, escalating to an impressive 9.4% absolute improvement on the highly challenging AIME benchmark. Across in-domain benchmarks, our method yields an average improvement of 8.9%, and for out-of-domain benchmarks, we achieve a substantial average gain of 6.98%. These results validate the effectiveness of our approach across model families and problem difficulty levels.

#### 4.2 Ablation Study

We conduct three primary analyses: (a) comparison with standard RL and SFT baselines to validate our method’s effectiveness in facilitating exploration, (b) visualization of exploration patterns to reveal limitations in the standard RL paradigm, and (c) behavioral analysis of code integration strategies. These analyses collectively demonstrate our method’s benefits in facilitating guided exploration and explains how it improves performance.

**Training Efficiency.** We evaluated the learning dynamics of our approach in direct comparison to three established training paradigms:

- *Base+RL*: On-policy Reinforcement Learning (RL) initialized from a base model without Supervised Fine-Tuning (SFT). This follows the methodology of DeepSeek R1, designed to isolate and assess the pure effects of RL training.
- *SFT*: Supervised Fine-Tuning, the prevailing training paradigm widely adopted in current tool-integrated math Language Models (LMs).
- *SFT+RL*: Standard RL applied after SFT, serving as a conventional baseline for evaluating our EM-based RL method.

From the figure, we make the following key observations:

- While Reinforcement Learning directly from the base model (*Base+RL*) exhibits consistent performance improvement, its training efficiency is lower than training paradigms incorporating SFT. In addition, the model rarely explores code-integrated solutions, with the code invocation rate below 5%. This strongly suggest that *reinforcement learning tool-usage behavior from scratch is inherently inefficient*.
- SFT effectively provides a strong initialization point, but *SFT alone exhibits limited asymptotic performance*. This suggests that SFT lacks the capacity to adapt and optimize beyond the scope of the expert demonstrations, thereby limiting further improvement.
- Standard RL applied after SFT shows initial further improvement but subsequently plateaus, *even after an extended training stage*. This suggests *the exploration-exploitation dilemma when applying RL for LLM post-training*: standard RL with vanilla rollout exploration tends to exploit local optima and insufficiently explores the combinatorial code-integrated trajectories.

To further substantiate the exploration limitations inherent in the conventional *SFT+RL* paradigm, we present a visualization of the exploration patterns. We partitioned the model-generated responses during self-exploration into three distinct training phases and analyzed the statistical distribution of code invocation rates across queries as the model’s policy evolved throughout training. As depicted in Figure 5, the distribution of code invocation progressively concentrates towards the extremes – either minimal or maximal code use – indicating the model’s growing tendency to exploit its local policy neighborhood. This exploitation manifests as a focus on refining established code-triggering decisions, rather than engaging in broader exploration of alternative approaches.

These empirical observations lend strong support to our assertion that standard RL methods are

Table 2: **Effectiveness of AutoCode4Math.** The column "Code?" indicates whether code integration is involved, with ★ representing autonomous determination of code integration by the model. The improvement over code-driven inference is highlighted in the colored row, denoted as  $\Delta$ .

Model	Code?	In-domain		Out-of-domain			
		GSM8K	MATH500	GaoKao	Olympiad	AIME24	AMC23
Proprietary Model							
OpenAI-o1-preview (Jaech et al., 2024)	✗		85.5	62.1	52.1	44.6	81.8
GPT-4o (Achiam et al., 2023)	✗	92.9	76.4	67.5	43.3	9.3	45.8
Claude-3.5-Sonnet-1022 (Anthropic, 2024)	✗	95	78.3			16.0	
Open-Source Models							
Mammoth-70B (Yue et al., 2023)	✓	76.9	41.8	25.2			
ToRA-70B (Gou et al., 2023)	✓	84.3	49.7	31.7			
NuminaMath-72B (Li et al., 2024)	✓	91.4	59.2	49.4	36.7	6.5	40.6
Mathstral-7B (AI, 2024)	✗	84.9	56.6	46	21.5		
Mammoth-Mistral-7B (Yue et al., 2023)	✓	74.22	37.8	22.08	9.63	6.67	20.0
NuminaMath-7B-CoT (Li et al., 2024)	✗	81.27	53.0	48.83	22.22	3.33	25.0
Dart-Math-DeepSeek-7B (Tong et al., 2024)	✗	87.64	50.0	45.45	18.52	3.33	35.0
Dart-Math-Llama3-8B (Tong et al., 2024)	✗	82.71	45.0	34.80	23	0.0	17.5
AutoCode Training							
Qwen2Math-Base-7B (Yang et al., 2024a)	✗	80.74	48.80	43.37	21.62	6.5	19.8
<b>AutoCode4Math-Qwen2</b>	★	88.1	61.86	50.13	26.37	13.2	30.0
$\Delta$		7.36 $\uparrow$	13.06 $\uparrow$	6.76 $\uparrow$	4.75 $\uparrow$	6.7 $\uparrow$	10.2 $\uparrow$
DeepseekMath-Instruct-7B (Shao et al., 2024)	✓	84.46	51.00	44.68	20.44	1.6	17.4
<b>AutoCode4Math-DeepSeek</b>	★	89.26	63.32	50.53	26.95	9.5	28.8
$\Delta$		4.8 $\uparrow$	12.32 $\uparrow$	5.85 $\uparrow$	6.51 $\uparrow$	7.9 $\uparrow$	11.4 $\uparrow$
Qwen-2.5-Base-7B (Yang et al., 2024b)	✗	84.88	60.4	45.45	30.37	13.2	39.38
<b>AutoCode4Math-Qwen2.5</b>	★	89.12	71.4	51.69	32.6	22.6	45.18
$\Delta$		4.24 $\uparrow$	11.0 $\uparrow$	6.24 $\uparrow$	2.23 $\uparrow$	9.4 $\uparrow$	5.8 $\uparrow$

susceptible to premature exploitation of the local policy space when learning AutoCode strategies. In sharp contrast, our proposed EM method facilitates a more guided exploration by sub-sampling trajectories according to the reference strategy (Sec. 3.2). This enables continuous performance (evidenced in Sec. 4.1) and mitigating the risk of converging to suboptimal local optima (Fig. 4).

**Analysis on Code Integration Behaviors.** We investigated the properties of the learned code integration strategies to gain deeper insights into the mechanisms behind our method’s performance gains. Our central hypothesis posits that optimal code integration unlocks synergistic performance benefits by effectively combining the strengths of CoT and code executions. This synergy presents a "free lunch" scenario: a well-learned metacognitive tool-usage strategy can elevate overall performance, provided the model demonstrates competence in solving *distinct* subsets of queries using either CoT or code execution.

To empirically validate this "free lunch" principle and demonstrate the superiority of our approach in realizing it, we benchmarked our model against

baselines that inherently support both code execution and Chain-of-Thought (CoT) reasoning: GPT-4, Mammoth-70B, and DeepseekMath-Instruct-7B. Our analysis evaluated the model’s autonomous decision to invoke code when not explicitly instructed on which strategy to employ. We compared this "AutoCode" performance against scenarios where models were explicitly prompted to utilize either code or CoT reasoning. We also considered the theoretical "free lunch" upper bound – the accuracy achieved by combining the successful predictions from either strategy (i.e., taking the union of queries solved by CoT or code).

As visually presented in Figure 6, existing baseline models exhibit inferior performance in AutoCode mode compared to scenarios where code invocation is explicitly prompted, e.g., DeepseekMath-Instruct-7B shows a degradation of 11.54% in AutoCode mode. This suggests that their AutoCode strategies are often suboptimal, performing closer to random selection between CoT and code (selection accuracy near 50%), resulting in AutoCode falling between the performance of explicitly triggered CoT and code. In contrast, our

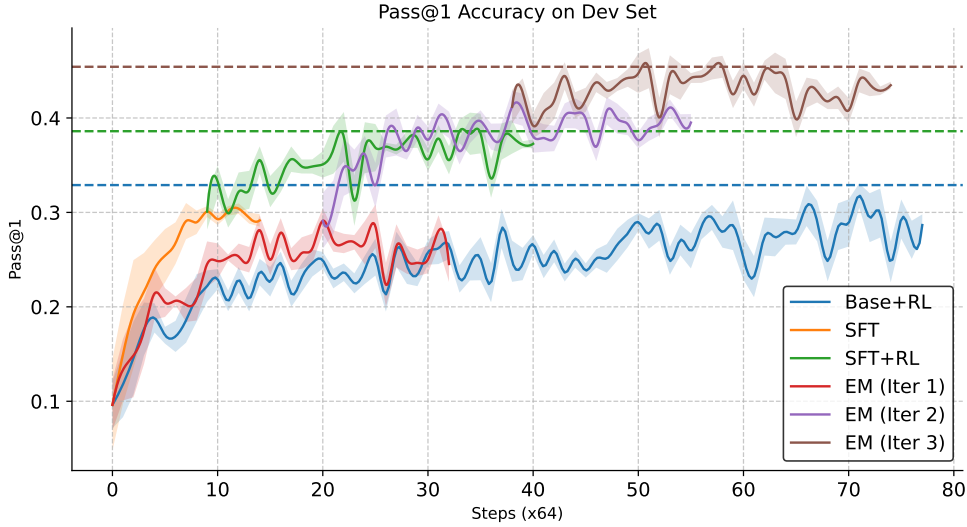


Figure 4: **Training Efficiency and Convergence.** We benchmark the learning dynamics of our approach against three training paradigms: supervised fine-tuning and reinforcement learning (RL). The Pass@1 accuracy is evaluated on an held-out dev-set. We use Qwen-2.5-Base as the base model. SFT is conducted using collected public data (Toshniwal et al., 2024; Yue et al., 2023). The dashed lines indicate asymptotic performance.

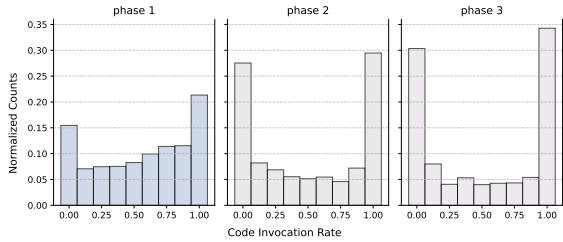


Figure 5: **Visualization of Exploration in the SFT+RL paradigm.** The distribution of code invocation rates *across queries* to visualize policy’s exploration of code-integrated trajectories. Without external guidance, LLM tends to exploit its local policy neighborhood, concentrating code usage toward extremes as training phase evolves.

models learn more effective code integration strategies. AutoCode4Math-Qwen2.5, for example, improves upon explicitly code-triggered performance by 7%, indicating a true synergistic integration of reasoning and code execution.

To quantify the effectiveness of these learned "AutoCode" strategies, we calculated the CoT/code selection accuracy. We used the outcome of explicit instruction (i.e., performance when explicitly prompted for CoT or code) as a proxy for the ground-truth optimal method selection. Our model achieves a selection accuracy of 89.53%, showcasing the high efficacy of the learned code integration strategy.

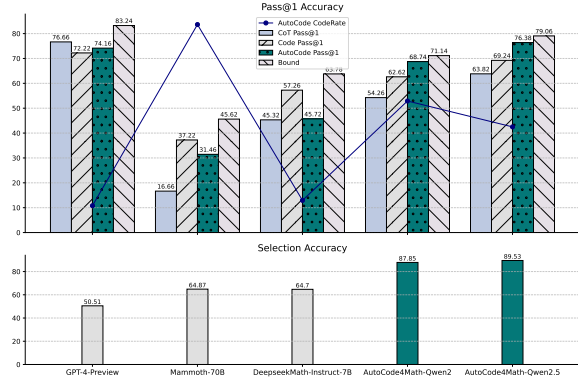


Figure 6: **Analysis of AutoCode Strategies.** We compare AutoCode performance against scenarios where models explicitly prompted to utilize code or CoT, and consider the union of solved queries as the bound for AutoCode performance. Existing models show inferior AutoCode performance than explicit instructed, with their AutoCode strategies close to random (50%). Our approach consistently improves AutoCode performance, with AutoCode selection accuracy near 90%.

## 5 Related Work and Discussion

**Tool-Integrated Math LLMs.** Math language models adopted two major paradigms: Chain-of-Thought (CoT) reasoning and the use of external tools, such as Python programs (Yu et al., 2023; Yue et al., 2023; Toshniwal et al., 2024). Each paradigm offers unique benefits, and recent hybrid frameworks (Yue et al., 2023; Gou et al., 2023; Li, 2024; Shao et al., 2024; Yang et al., 2024b) increasingly seek to combine them for synergy. However, current models exhibit critical rigidity, motivating



our work to realize the true metacognitive capacity that enjoys synergistic benefits of CoT and code.

**EM for RL.** Expectation-Maximization (EM) has proven effective for maximum likelihood problems involving hidden variables, such as Expert Iteration (Anthony et al., 2017), Iterative Maximum Likelihood (Wu, 2016; Agarwal et al., 2019), Meta-Reinforcement Learning (Zintgraf et al., 2019; Wang et al., 2020), and Adversarial Games (Wang et al., 2023). In the context of math LLMs, the most relevant works are (Singh et al., 2023) and (Ni et al., 2022), which apply EM-style iterative self-training to math problem-solving. Unlike these approaches, we leverage the EM framework for guided exploration during reinforcement learning of language models.

## 6 Conclusion

Existing tool-integrated math language models lack the metacognitive capacity to effectively determine code integration, hindering their ability to fully realize the synergistic benefits of tool integration and CoT. To address this critical gap, we propose a novel EM-based framework that combines guided exploration with policy optimization. Our experiments demonstrate the limitations of standard SFT and RL in efficiently exploring the combinatorial space of code-integrated trajectories and highlight the superior training efficiency and performance of our approach.

## 7 Limitations

The scope of our work is primarily focused on mathematical problem-solving. While we observe promising results on challenging benchmarks like MATH500, the generalizability of our approach to other domains requiring the metacognitive capacity of tool integration and CoT, such as scientific reasoning or code generation for general-purpose tasks, remains to be explored. Future work should investigate the effectiveness of our framework across a wider range of tasks and domains.

## References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Rishabh Agarwal, Chen Liang, Dale Schuurmans, and Mohammad Norouzi. 2019. Learning to generalize

from sparse and underspecified rewards. In *International conference on machine learning*, pages 130–140. PMLR.

Mistral AI. 2024. [Mathstral](#).

Thomas Anthony, Zheng Tian, and David Barber. 2017. Thinking fast and slow with deep learning and tree search. *Advances in neural information processing systems*, 30.

Anthropic. 2024. The claude 3 model family: Opus, sonnet, haiku.

Christopher M Bishop and Nasser M Nasrabadi. 2006. *Pattern recognition and machine learning*, volume 4. Springer.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.

Zui Chen, Yezeng Chen, Jiaqi Han, Zhijie Huang, Ji Qi, and Yi Zhou. 2024. [An empirical study of data ability boundary in llms' math reasoning](#). *Preprint*, arXiv:2403.00799.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. Pal: Program-aided language models. In *International Conference on Machine Learning*, pages 10764–10799. PMLR.

Zhibin Gou, Zhihong Shao, Yeyun Gong, Yujiu Yang, Minlie Huang, Nan Duan, Weizhu Chen, et al. 2023. Tora: A tool-integrated reasoning agent for mathematical problem solving. *arXiv preprint arXiv:2309.17452*.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, et al. 2024. Olympiad-bench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. *arXiv preprint arXiv:2402.14008*.

- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.
- Jean Kaddour, Joshua Harris, Maximilian Mozes, Herbie Bradley, Roberta Raileanu, and Robert McHardy. 2023. Challenges and applications of large language models. *arXiv preprint arXiv:2307.10169*.
- Jia Li, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Huang, Kashif Rasul, Longhui Yu, Albert Q Jiang, Ziju Shen, et al. 2024. Numinamath: The largest public dataset in ai4maths with 860k pairs of competition math problems and solutions.
- Long Li. 2024. How do humans write code? large models do it the same way too. *arXiv preprint arXiv:2402.15729*.
- Minpeng Liao, Wei Luo, Chengxi Li, Jing Wu, and Kai Fan. 2024. Mario: Math reasoning with code interpreter output—a reproducible pipeline. *arXiv preprint arXiv:2401.08190*.
- Ansong Ni, Jeevana Priya Inala, Chenglong Wang, Oleksandr Polozov, Christopher Meek, Dragomir Radev, and Jianfeng Gao. 2022. Learning math reasoning from self-sampled correct and partially-correct solutions. *arXiv preprint arXiv:2205.14318*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, YK Li, Yu Wu, and Daya Guo. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Avi Singh, John D Co-Reyes, Rishabh Agarwal, Ankesh Anand, Piyush Patil, Peter J Liu, James Harrison, Jaehoon Lee, Kelvin Xu, Aaron Parisi, et al. 2023. Beyond human data: Scaling self-training for problem-solving with language models. *arXiv preprint arXiv:2312.06585*.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Yuxuan Tong, Xiwen Zhang, Rui Wang, Ruidong Wu, and Junxian He. 2024. DART-math: Difficulty-aware rejection tuning for mathematical problem-solving. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Shubham Toshniwal, Ivan Moshkov, Sean Narenthiran, Daria Gitman, Fei Jia, and Igor Gitman. 2024. Openmathinstruct-1: A 1.8 million math instruction tuning dataset. *arXiv preprint arXiv:2402.10176*.
- Haozhe Wang, Chao Du, Panyan Fang, Li He, Liang Wang, and Bo Zheng. 2023. Adversarial constrained bidding via minimax regret optimization with causality-aware reinforcement learning. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 2314–2325.
- Haozhe Wang, Jiale Zhou, and Xuming He. 2020. Learning context-aware task reasoning for efficient meta reinforcement learning. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1440–1448.
- 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.
- Yonghui Wu. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. 2024a. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024b. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.
- Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. 2023. Metamath: Bootstrap your own mathematical questions for large language models. *arXiv preprint arXiv:2309.12284*.
- Siyu Yuan, Zehui Chen, Zhiheng Xi, Junjie Ye, Zhengyin Du, and Jiecao Chen. 2025. Agent-r: Training language model agents to reflect via iterative self-training. *arXiv preprint arXiv:2501.11425*.
- Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhui Chen. 2023. Mammoth: Building math generalist models through hybrid instruction tuning. *arXiv preprint arXiv:2309.05653*.
- Luisa Zintgraf, Kyriacos Shiarlis, Maximilian Igl, Sebastian Schulze, Yarin Gal, Katja Hofmann, and Shimon Whiteson. 2019. Varibad: A very good method

for bayes-adaptive deep rl via meta-learning. *arXiv preprint arXiv:1910.08348*.

## A Appendix

### A.1 Derivation of the EM

We first model RL as a maximum likelihood problem. Specifically, we consider ‘maximizing expected return’ as ‘maximizing the likelihood of observing a correct response’, this is equivalent to maximizing the following log-likelihood,

$$\begin{aligned} & \max_{\theta} \log P(r = 1|x_q; \theta) \\ &= \max_{\theta} \log \sum_c p_{\theta}(c|x_q) \sum_{y_a} p_{\theta}(y_a|x_q, c) \\ & \cdot P(r = 1|y_a, c, x_q) \\ & \doteq \mathcal{J}(\theta), \end{aligned}$$

where we factorize the language model as  $p_{\theta}(y_a|x_q) = \sum_c p_{\theta}(c|x_q)p_{\theta}(y_a|x_q, c)$ .

Since the variable  $c$  is unobservable for lack of reliable supervision, we resort to the EM framework. Specifically, we treat  $c$  as a hidden variable, and introduce a proposal distribution  $s(c|x_q)$ , which represents the belief of  $c$  given the fixed LLM  $\mathcal{M}_{\theta}$  and the data evidence. We have the following derivations,  $\mathcal{J}(\theta) =$

$$\begin{aligned} & \log \sum_c s(c|x_q) \cdot \frac{p_{\theta}(\cdot|x_q)}{s(\cdot|x_q)} \\ & \cdot \sum_{y_a} p_{\theta}(y_a|x_q, c) \cdot P(r = 1|y_a, c, x_q) \\ & \geq \sum_c s(c|x_q) \\ & \left[ \log \frac{p_{\theta}(c|x_q) \sum_{y_a} p_{\theta}(y_a|x_q, c) P(r = 1|y_a, c, x_q)}{s(c|x_q)} \right] \\ & = \mathbb{E}_{s(c|x_q)} \left[ \log \frac{p_{\theta}(c|x_q) \cdot P(r = 1|c, x_q; \theta)}{s(c|x_q)} \right] \\ & \doteq \mathcal{J}_{\text{ELBO}}(s, \theta), \end{aligned}$$

where  $P(r = 1|c, x_q; \theta)$  denotes the likelihood of generating correct responses given  $(x_q, c)$  following the solution-generation policy  $p_{\theta}(y_a|x_q, c)$ .

In the E-step, we are essentially minimizing a KL divergence,

$$\max_s \mathcal{J}_{\text{ELBO}}(s, \theta) = \min_s D_{\text{KL}}(s(c|x_q) | s^*(c|x_q)), \quad (7)$$

where the minimizer is  $s^*(c|x_q) \propto p_{\theta}(c|x_q) \cdot P(r = 1|c, x_q; \theta)$ . Intuitively, the belief over the methodology  $c$  following a posterior update: it is based on the prior of the current model  $p_{\theta}(c|x_q)$  and the likelihood of data  $P(r = 1|c, x_q; \theta)$ . The

optimal methodology-selection strategy assigns higher probability to method  $c$  if following the current LLM it has higher likelihood to generate correct responses or higher prior over it.

To compute the optimal strategy, we define  $P(r = 1|c, x_q; \theta) = \frac{\exp(\alpha \cdot p_{\theta}(c|x_q) Q^{\theta}(x_q, c))}{Z(x_q)}$  as an energy-based distribution, where the negative energy function evaluates the expected return induced by the current solution-generation policy:  $Q(x_q, c; \theta) = \mathbb{E}_{p_{\theta}(y_a|x_q, c)}[R(x_q, y_a)]$ ,  $\alpha > 0$  is a temperature controlling the sharpness of the distribution. Then the minimizer can be computed by enumerating over  $c$ .

In the M-step, we optimize  $\max_{\theta} \mathcal{J}_{\text{ELBO}}(s, \theta) =$

$$\begin{aligned} & \max_{\theta} \mathbb{E}_{s(c|x_q)} [\log P(r = 1|c, x_q; \theta)] \\ & - D_{\text{KL}}(s(c|x_q) | p_{\theta}(c|x_q)) \\ & = \max_{\theta} \mathbb{E}_{s(c|x_q)} [Q(x_q, c; \theta)] \\ & - D_{\text{KL}}(s(c|x_q) | p_{\theta}(c|x_q)) \end{aligned}$$

which maximizes the expected return while imitating the optimal strategy.

### A.2 Distribution of Multi-round

We show the distribution of multi-round responses of the model AutoCode4Math-Deepseek in Tab. 3.

## B Examples

### B.1 Success Cases

Listing 1: Success case for multi-round.

Round 1:

Question: Henry and 3 of his friends order 7 pizzas for lunch. Each pizza is cut into 8 slices. If Henry and his friends want to share the pizzas equally, how many slices can each of them have?

Solution:

There are 7 pizzas in total, and each pizza is cut into 8 slices. So, there are  $7 * 8 = 56$  slices in total.

Henry and his 3 friends want to share the pizzas equally, so they will each have  $56 / 4 = 14$  slices.

Each of them can have 14 slices.

The answer is  $\boxed{14}$ .

Table 3: Distribution of Multi-Round Responses.

Model	Dataset	Queries	Round1	Round2	Round3
Deepseek-math-AutoCode	MATH	5000	1252	3719	29
	GSM8k	1319	1013	304	2

-----

Round2:

Question: Henry and 3 of his friends order 7 pizzas for lunch. Each pizza is cut into 8 slices. If Henry and his friends want to share the pizzas equally, how many slices can each of them have?

Solution:

```
```python
def pizza_slices_per_person():
    """Henry and 3 of his friends order 7
    pizzas for lunch. Each pizza is cut
    into 8 slices. If Henry and his
    friends want to share the pizzas
    equally, how many slices can each of
    them have?"""
    total_pizzas = 7
    slices_per_pizza = 8
    total_slices = total_pizzas *
        slices_per_pizza
    people = 4 # Henry and his 3 friends
    slices_per_person = total_slices //
        people # Use floor division to get
        an integer result
    return slices_per_person

slices_per_person =
    pizza_slices_per_person()
print(slices_per_person)
```
```

```output

```
14
```
Each of them can have  $\boxed{14}$ 
slices.
```

-----

Round3:

Question: Peter has 4 boxes with the same number of chocolate bars in each, while Martha has 7 boxes with the same number of chocolate bars in

each. If Peter and Martha have totals of 64 and 56 chocolate bars respectively, how many more chocolate bars does Peter have in each box than Martha?

Solution:

If Peter has 4 boxes and Martha has 7 boxes, then Martha has  $7/4 = 1.75$  times as many boxes as Peter. If Martha has 56 chocolate bars and Peter has 64 chocolate bars, then Peter has  $64/56 = 1.142857$  times as many chocolate bars as Martha. If Martha has 1.75 times as many boxes as Peter, and Peter has 1.142857 times as many chocolate bars as Martha, then Martha has  $1.75/1.142857 = 1.5$  times as many chocolate bars in each box as Peter. Therefore, Peter has  $1.5 - 1 = 0.5$  more chocolate bars in each box than Martha.

The answer is  $\boxed{0.5}$ .

Sorry, the above solution may be incorrect. It's better to write a program.

```
```python
def chocolate_bars_difference():
    """Peter has 4 boxes with the same
    number of chocolate bars in each,
    while Martha has 7 boxes with the
    same number of chocolate bars in
    each. If Peter and Martha have
    totals of 64 and 56 chocolate bars
    respectively, how many more
    chocolate bars does Peter have in
    each box than Martha?"""
    peter_total = 64
    martha_total = 56
    peter_boxes = 4
    martha_boxes = 7
    peter_bars_per_box = peter_total /
        peter_boxes
    martha_bars_per_box = martha_total /
        martha_boxes
```

```

difference = peter_bars_per_box -
    martha_bars_per_box
return difference

difference = chocolate_bars_difference()
print(difference)
'''
'''output
8.0
'''

Peter has  $\boxed{8}$  more chocolate
bars in each box than Martha.

```

## B.2 Failure Cases

Listing 2: Failure case for multi-round.

```

Round 1:
Question: The solution to the inequality
 $\frac{x + c}{x^2 + ax + b} \leq 0$  is
 $x \in (-\infty, -1) \cup [1, 2)$ .
Find  $a + b + c$ .
Solution:
Let's reason step by step.
The inequality is equivalent to
 $(x + c)(x^2 + ax + b) \leq 0$ . Since
the roots of  $x^2 + ax + b$  are  $-1$ 
and  $1$ , we can write
 $x^2 + ax + b = (x + 1)(x - 1) = x^2 -
1$ . Then
 $(x + c)(x^2 + ax + b) = (x + c)(x^2 -
1) = (x + c)(x + 1)(x - 1)$ . The
inequality is satisfied when  $x \in
(-\infty, -1) \cup [1, 2)$ , so
 $(x + c)(x + 1)(x - 1) \leq 0$  when  $x \in
(-\infty, -1) \cup [1, 2)$ . This
means that one of  $x + c$ ,  $x + 1$ ,
or  $x - 1$  is nonpositive, and the
other two are nonnegative.

If  $x + c \leq 0$ , then  $x \leq -c$ . Since
 $x \in [1, 2)$ , we must have  $-c \geq
2$ , or  $c \leq -2$ .

If  $x - 1 \leq 0$ , then  $x \leq 1$ . Since
 $x \in (-\infty, -1)$ , we must have  $-1 \leq 1$ , which is always true.

If  $x + 1 \leq 0$ , then  $x \leq -1$ . Since
 $x \in [1, 2)$ , we must have  $-1 \leq
1$ , which is always true.

```

Therefore, the solution is  $x \in (-\infty, -c] \cup [1, 2)$ . Since  $x \in [1, 2)$ , we must have  $-c \geq 2$ , or  $c \leq -2$ .

Hence,  $a + b + c = 0 - 1 + (-2) = \boxed{-3}$ .  
The answer is  $\boxed{-3}$ .

Round2:

Question: The girls are trying to raise money for a carnival. Kim raises \$320 more than Alexandra, who raises \$430, and Maryam raises \$400 more than Sarah, who raises \$300. How much money, in dollars, did they all raise in total?

Solution:

```

'''python
def total_money_raised():
    """The girls are trying to raise
    money for a carnival. Kim raises
    $320 more than Alexandra, who
    raises $430, and Maryam raises
    $400 more than Sarah, who raises
    $300. How much money, in
    dollars, did they all raise in
    total?"""
    alexandra_money = 430
    kim_money = alexandra_money + 320
    sarah_money = 300
    maryam_money = sarah_money + 400
    total_money = alexandra_money +
        kim_money + sarah_money +
        maryam_money
    return total_money

```

```

total_money = total_money_raised()
print(total_money)
'''
'''output
2180
'''

```

The girls raised a total of  $\boxed{2180}$  dollars.

Round3:

Question: The digits from 1 to 6 are

arranged to form a six-digit multiple of 5. What is the probability that the number is greater than 500,000? Express your answer as a common fraction.

Solution:

Let's reason step by step.

To form a six-digit multiple of 5, the last digit must be either 0 or 5.

Since the number must be greater than 500,000, the last digit must be 5.

The remaining five digits can be arranged in  $5!$  ways.

Therefore, the probability is  $\frac{5!}{6!} = \boxed{\frac{1}{6}}$ .

The answer is  $\boxed{\frac{1}{6}}$ .

Based on the above reasoning, we can verify by writing a program.

```
```python
from sympy import Rational

def
    probability_greater_than_five_million
    ():
    """The digits from 1 to 6 are
        arranged to form a six-digit
        multiple of 5. What is the
        probability that the number is
        greater than 500,000? Express
        your answer as a common fraction
        ."""
    # Total number of arrangements
    total_arrangements = 6*5*4*3*2*1

    # Number of arrangements that start
    with 5 or 6
    valid_arrangements = 2*5*4*3*2*1

    # Number of arrangements that start
    with 1
    invalid_arrangements_1 = 1*5*4*3*2*1

    # Number of arrangements that start
    with 2
    invalid_arrangements_2 = 1*4*3*2*1

    # Number of arrangements that start
    with 3
    invalid_arrangements_3 = 1*3*2*1

    # Total number of invalid
    arrangements
```

```
invalid_arrangements =
    invalid_arrangements_1 +
    invalid_arrangements_2 +
    invalid_arrangements_3
```

```
# Calculate the probability
probability = Rational(
    valid_arrangements,
    total_arrangements)
```

```
return probability
```

```
result =
    probability_greater_than_five_million
    ()
```

```
print(result)
```

```
```
```

```
```output
```

```
1/3
```

```
```
```

The probability that the number is greater than 500,000 is  $\boxed{\frac{1}{3}}$ .

---

### B.3 AutoCode 4shots

Listing 3: Autocode 4shots example.

User: Simplify  $\sqrt{25000}$ .

Assistant:

Let's write a program.

```
```python
from sympy import sqrt
```

```
def simplify_sqrt():
    """Simplify  $\sqrt{25000}$ ."""
    simplified_sqrt = sqrt(25000)

    return simplified_sqrt
```

```
result = simplify_sqrt()
```

```
print(result)
```

```
```
```

User: The function  $f(x)$  satisfies  $[f(x) + f(x)] = 4f(x)$  for all  $x$ , and  $f(1) = 4$ . Find  $f(21)$ .

Assistant:

Let's reason step by step.

We can start by plugging in  $x = 1$  into the given equation. This gives us

$[f(1) + f(1)] = 4f(1)$ .

Since we know that  $f(1) = 4$ , we can substitute this into the equation to get

$$f(1 + 4) = 4 \cdot 4.$$

Simplifying, we have

$$f(5) = 16.$$

Next, we can plug in  $x = 5$  into the given equation. This gives us

$$f(5 + f(5)) = 4f(5).$$

Since we know that  $f(5) = 16$ , we can substitute this into the equation to get

$$f(5 + 16) = 4 \cdot 16.$$

Simplifying, we have

$$f(21) = 64.$$

Therefore, the value of  $f(21)$  is  $\boxed{64}$ .

The answer is  $\boxed{64}$ .

User: Find the matrix that corresponds to reflecting over the  $x$ -axis.

Assistant:

Let's reason step by step.

The reflection of a point  $(x, y)$  over the  $x$ -axis is  $(x, -y)$ .

Thus, the matrix that corresponds to reflecting over the  $x$ -axis is  $\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ .

The answer is  $\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ .

User: If  $a$  is three times  $b$ ,  $b$  is twice  $c$ , and  $c$  is four times  $d$ , what is the value of  $\frac{a \cdot c}{b \cdot d}$ ?

Assistant:

Let's write a program.

```

```python
from sympy import symbols, simplify

def ratio_of_products():
    a, b, c, d = symbols('a b c d')

    # Given relations
    a = 3 * b
    b = 2 * c
    c = 4 * d

    # Substitute the relations into the
    # expression
    expr = a * c / (b * d)
    expr = expr.subs({a: 3 * b, b: 2 * c
                      , c: 4 * d})

```

```

# Simplify the expression
simplified_expr = simplify(expr)

return simplified_expr

result = ratio_of_products()
print(result)
```

```