

Beyond Accuracy: Unveiling Inefficiency Patterns in Tool-Integrated Reasoning

Qisheng Su^{1,2}, Shiting Huang¹, Zhen Fang¹, Ziyang Chen¹, Zehui Chen¹, Feng Zhao^{1*}

¹MoE Key Lab of BIPC, University of Science and Technology of China

²Shanghai Innovation Institute

nicksu@mail.ustc.edu.cn fzhao956@ustc.edu.cn

Abstract

In real-world Tool-Integrated Reasoning (TIR) scenarios, where LLMs interleave reasoning with external tool calls, a major source of inefficiency is that the toolcalls create pauses between LLM requests and cause KV-Cache eviction, forcing recomputation. Also, the long, unfiltered response returned by external tools inflates the KV-Cache, so each decode step spends more time loading the growing cache and thus becomes steadily slower as context length increases. However, existing efficiency metrics like token counts and toolcall counts fail to capture the real model inference latency. To address this, we introduce PTE (Prefill Token Equivalents), a hardware-aware TIR-efficiency metric that unifies internal reasoning and external tool-use costs while explicitly accounting for non-reusable KV-Cache and long-tool-response scenarios. Validation in a high-concurrency industrial setting indicates that PTE aligns significantly better with wall-clock latency than standard token counts, while maintaining consistent efficiency rankings across diverse hardware profiles. We conduct extensive experiments across five TIR benchmarks, quantify their PTE costs, and identify four inefficiency patterns that appear in TIR. We also discover that trajectories with higher PTE costs tend to have lower reasoning correctness, indicating that simply using more tools does not improve the quality of the answer. The code is available at <https://github.com/sqs-ustc/tool-reasoning-framework-PTE>.

1 Introduction

Large Language Models (LLMs) demonstrate remarkable capabilities in complex tasks via Tool-Integrated Reasoning (TIR)(Li et al., 2025b,c; Tongyi DeepResearch Team, 2025; Gou et al., 2024; Wang et al., 2023; Li et al., 2025d; Lin et al., 2024; Xiao et al., 2025). However, TIR efficiency

*Corresponding Author

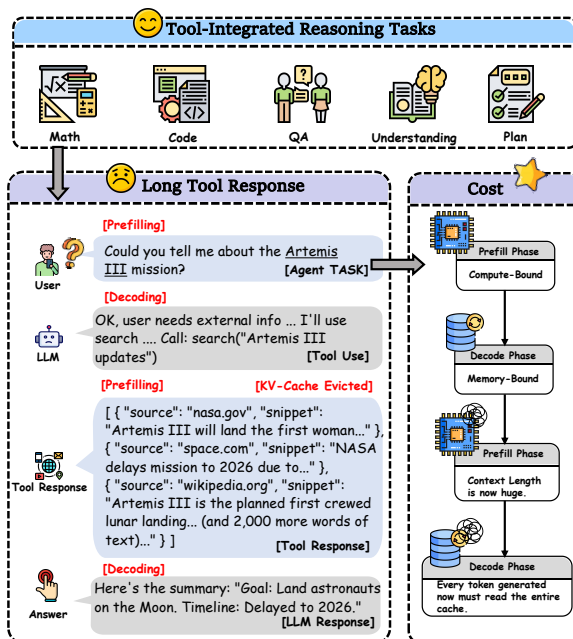


Figure 1: **Illustration of the asymmetric costs in Tool-Integrated Reasoning (TIR).** Toolcalls induce KV-Cache eviction, and long tool responses inflate the context length, significantly increasing the cost of subsequent decoding steps.

evaluation remains underexplored. In real-world scenarios, toolcalls create pauses between LLM requests that cause KV-Cache eviction(Li et al., 2025a; Pan et al., 2025), while long, unfiltered tool responses inflate the context length. This renders the memory-bound decode phase significantly more expensive, yet existing TIR benchmarks focus primarily on accuracy(Wei et al., 2025a; Wong et al., 2025), and efficiency metrics rely on naive token counts or toolcall counts(Gao et al., 2025; Wang et al., 2025a) that fail to reflect the true runtime drain.

Specifically, current efficiency metrics fail to capture the asymmetric costs of the compute-bound prefill phase and the memory-bound decode phase. This oversight proves costly in TIR, where toolcalls trigger KV-Cache eviction and long responses in-

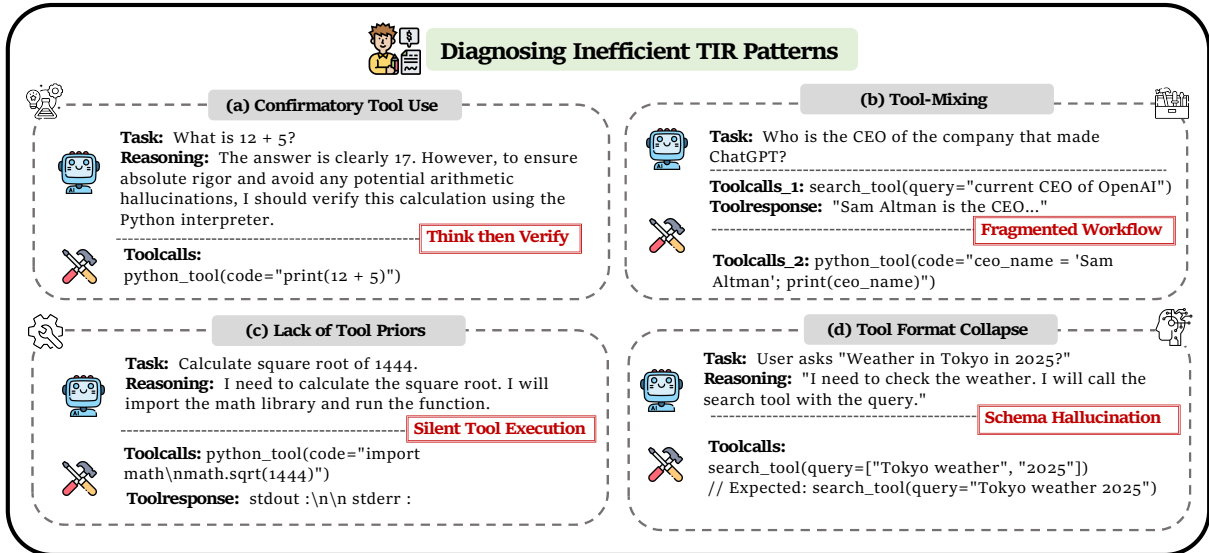


Figure 2: Overview of the four inefficiency patterns in Tool-Integrated Reasoning.

flate the context, directly increasing the HBM transfer overhead during decode for every subsequent reasoning token. Consequently, existing metrics lack a unified framework to weigh the true expense of internal reasoning against external tool use.

To address these challenges, we introduce a comprehensive evaluation framework grounded in the first principles of transformer inference. We propose PTE (Prefill Token Equivalents), a hardware-aware metric that explicitly models the asymmetric costs of the prefill and decode phases. PTE prices the memory-bound decode cost in units of one compute-bound prefill token, giving a single scale that unifies internal reasoning and external tool use while explicitly accounting for non-reusable KV-Cache and long-tool-response scenarios. Therefore, we provide PTE as an efficiency metric that offers an estimate closer to physical hardware behavior in TIR.

To validate our metric, we first demonstrate in a high-concurrency industrial setting that PTE aligns closely with wall-clock latency, significantly outperforming token-count metrics. We further confirm its robustness across various hardware profiles, showing that PTE maintains consistent model efficiency rankings regardless of the specific device.

With this foundation, we conduct extensive experiments across five TIR benchmarks, quantify PTE costs for thousands of trajectories, and dissect four TIR inefficiency patterns: Confirmatory Tool Usage, Tool-Mixing, Lack of Tool Priors, and Tool Format Collapse (as illustrated in Fig. 2). Crucially, contrasting PTE with standard token counts reveals

a divergence where token generation tends to front-load budgets (Fig. 8), whereas PTE exposes how hardware costs actually escalate in later steps due to context accumulation (Fig. 5). Consequently, we observe that SOTA models with near-identical accuracy can differ by orders of magnitude in PTE. We further identify a clear trade-off in thinking models, finding that while their heavy computational overhead pays off on complex reasoning tasks, it yields diminishing returns on simpler ones. Finally, statistical results indicate that trajectories with higher PTE costs tend to have lower reasoning correctness, suggesting that simply using more tools does not improve answer quality.

The main contributions of our work are summarized as follows:

- We introduce PTE, a hardware-aware metric that unifies reasoning and tool-use costs by explicitly modeling prefill–decode asymmetry in TIR.
- We validate PTE in a high-concurrency industrial setting, demonstrating that it aligns significantly better with wall-clock latency than token counts and maintains consistent model rankings across diverse hardware profiles.
- Through extensive experiments on five benchmarks, we quantify TIR costs, identify four distinct inefficiency patterns, and reveal a negative correlation between PTE costs and reasoning correctness.
- We release a high-concurrency, modular TIR framework featuring flexible tool customization and built-in evaluation.

2 Related Work

Benchmarks for Tool-Augmented LLMs Early TIR benchmarks focused on API selection and plan decomposition, often abstracting away execution. Examples include BFCL(Patil et al.), Webshop(Yao et al., 2022), ToolBench(Qin et al., 2023), and T-Eval(Chen et al., 2023), which primarily measure task success rate. Subsequent datasets like API-Bank(Li et al., 2023) and Critic-Tool(Huang et al., 2025) introduced execution-based metrics such as API success rate. Recently, Tool-Integrated Reasoning (TIR) benchmarks have evolved to cover complex, multi-step tasks. These include web-browsing (e.g., BrowseComp(Wei et al., 2025a), WideSearch(Wong et al., 2025), and GAIA(Mialon et al., 2023)) and domain-specific reasoning in math and code (e.g., GSM8K(Cobbe et al., 2021), MATH500(Hendrycks et al., 2021), SWE-Bench(Jimenez et al., 2024)). Concurrently, Serper(Dai et al., 2025) offers a principled metric for evaluating information-seeking efficiency in tool-augmented systems. While accuracy remains the primary metric, efficiency is often overlooked or measured simply by token or step counts (e.g., MCP-RADAR(Gao et al., 2025), ToolQA(Zhuang et al., 2023), CLASSIC(Xu et al., 2025)). Although recent works attempt to incorporate cost awareness, either through performance-aware cost(Zhao et al., 2025) or economic models for API pricing(Zellinger and Thomson, 2025), they lack a unified framework grounded in the physical latency of transformer inference.

Efficiency in Tool-Integrated Reasoning Research has identified inefficient behaviors in TIR, such as "cognitive offloading"(Wang et al., 2025a) and "over-tooluse"(Qian et al., 2025). To mitigate these, Reinforcement Learning (RL) has been widely adopted, primarily focusing on reward engineering and algorithmic innovation. In terms of reward engineering, approaches fall into two categories: (1) **Indirect optimization**, which implicitly improves efficiency by supervising reasoning quality, such as evaluating tool variety(Dong et al., 2025a), the contribution of individual steps(Yu et al., 2025b), or usage correctness(Singh et al., 2025); and (2) **Direct optimization**, which explicitly incorporates cost penalties. However, current penalties rely on naive metrics like toolcall counts(Wang et al., 2025a; Wei et al., 2025b) or token counts(Wang et al., 2025b; Liu et al., 2025), failing to capture hardware-level latency. Beyond

reward design, algorithmic innovations aim to optimize the inference process itself. These include entropy-based exploration strategies(Dong et al., 2025b; Chen et al., 2025d), dynamic routing between reasoning and tool-use(Chen et al., 2025c), and gradient-based stopping criteria(Yu et al., 2025a; Lin and Xu, 2025) to prune redundant steps.

3 PTE: A First-Principles Efficiency Metric

3.1 Background: The Physical Reality of LLM Inference Cost

Models	Active Params (N_{params})	Calculated γ
Open-Source Dense Models		
Qwen2.5-7B-Instruct	6.53B	0.00329
Qwen2.5-32B-Instruct	31.0B	0.00320
Qwen2.5-72B-Instruct	70.0B	0.00175
Qwen3-32B	31.2B	0.00200
Llama-3.1-8B-Instruct	8.0B	0.00625
Llama-3.1-70B-Instruct	70.6B	0.00175
Open-Source MoE Models		
Qwen3-30B-A3B	3.3B	0.00563
Qwen3-235B-A22B-Instruct	22B	0.00163
Qwen3-235B-A22B-Thinking	22B	0.00163
GLM-4.5-Air	12B	0.00200
GLM-4.5	32B	0.00183
DeepSeek-V3.1-Terminus	37B	0.00068
GPT-OSS-120B	5.1B	0.00388

Table 1: **Gamma Values of Different LLMs.** We evaluated a range of state-of-the-art open-source models, all of which were officially declared as having tool-calling capabilities.

As shown in Fig. 1, LLM inference costs are phase-dependent as the prefill and decode phases exhibit different bottlenecks.

Prefill Phase (Compute-Bound): Input tokens are processed in parallel, rendering this phase primarily limited by the GPU’s computational throughput (FLOPs).

Decode Phase (Memory-Bound): Output tokens are generated sequentially. Performance here is bottlenecked by the HBM bandwidth required to load model weights and the Key-Value (KV) Cache. Notably, the cost of retrieving the KV-Cache grows linearly with the cumulative context length (L_{seq}).

KV-Cache Eviction: In the context of TIR, we define eviction broadly to encompass any scenario where the cache state is invalidated during the tool-execution pause, ranging from Time-to-Live (TTL)

expiration in continuum serving systems to the architectural necessity of full re-computation in stateless pipelines.

3.2 Formal Definition of PTE

To bridge this gap, we introduce PTE (Prefill Token Equivalents), a metric that unifies both compute-bound and memory-bound costs into a single, physically meaningful unit: the equivalent cost of processing one input (prefill) token. For a reasoning trajectory with k turns, the total cost is:

$$PTE = \sum_{i=1}^k (D_{prefill_i} + \gamma \cdot L_{seq_i} \cdot D_{decode_i}), \quad (1)$$

where:

- $D_{prefill_i}$: The total number of context tokens up to turn i . This represents the compute-bound prefill cost.
- D_{decode_i} : The number of tokens generated by the model in turn i .
- L_{seq_i} : The cumulative sequence length (total context) before the decode phase begins in turn i .
- γ : A dimensionless coefficient representing the relative cost of a memory-bound operation to a compute-bound one.

3.3 Cost Modeling

To quantify the inference cost, we model the computational overhead of the two distinct phases.

Compute-Bound Prefill Cost. The prefill phase processes input tokens in parallel. Its cost is dominated by matrix multiplications, widely approximated as proportional to the model size (Kaplan et al., 2020; Zhong et al., 2024):

$$C_{prefill} \approx 2 \cdot N_{params} \text{ [FLOPs]}, \quad (2)$$

where N_{params} denotes the number of active parameters.

Memory-Bound Decode Cost. The decode phase is limited by the memory bandwidth required to load the KV cache. For a model with n_{layers} layers and hidden dimension d_{model} using FP16 precision (2 bytes), the memory access volume is $S_{KV} = 4 \cdot n_{layers} \cdot d_{model}$ [Bytes]. To unify dimensions with the prefill cost, we convert this memory volume into equivalent computational cost using the Hardware Operational Intensity (HOI) (Williams et al., 2009; Yuan et al., 2024; Peng et al., 2025):

$$C_{decode}^{eq} = S_{KV} \cdot HOI \text{ [FLOPs]}. \quad (3)$$

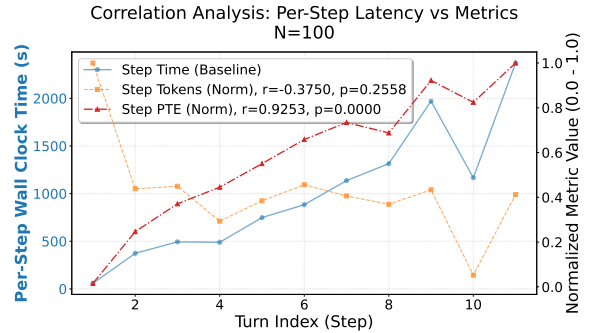


Figure 3: **Correlation analysis between real-world latency, PTE, and output token counts** ($N = 100$). Metrics are normalized for visual comparison. PTE demonstrates a strong linear correlation ($r = 0.9253$) with wall-clock latency, whereas output token counts show limited correlation ($r = -0.3750$).

The γ Coefficient. We define γ as the ratio of the equivalent decode cost to the prefill cost, representing the relative penalty of memory-bound operations:

$$\gamma = \frac{C_{decode}^{eq}}{C_{prefill}} = \frac{2 \cdot n_{layers} \cdot d_{model} \cdot HOI}{N_{params}}. \quad (4)$$

We refine γ to accommodate modern architectural optimizations: for **Grouped Query Attention (GQA)**, we scale γ by the KV-to-Query head ratio (H_{kv}/H_q); for **Multi-Head Latent Attention (MLA)**, we substitute d_{model} with the compressed dimensions ($d_{latent} + d_{rope}$).

γ serves as a static property of a model-hardware pair. Table 1 lists the results for all evaluated models, while the full statistics are provided in Appendix A.1. For a deeper interpretation of γ as a measure of computational cost scaling efficiency, please refer to Appendix A.4.

4 Validation: Fidelity and Robustness

To ensure PTE serves as a reliable efficiency metric, we validate it against two criteria: physical fidelity to real-world latency and robustness across hardware architectures.

4.1 Validation against Wall-clock Latency

We conducted high-concurrency experiments using DeepSeek-V3.2 on an $8 \times H200$ cluster to simulate industrial TIR scenario. The results in Fig. 3 reveal a divergence: the raw token counts show limited correlation with actual runtime ($r = -0.3750$, $N = 100$, $p = 0.2558$). In contrast, PTE demonstrates a strong positive correlation

Hardware	HOI (FLOPs/Byte)	γ Scaling Factor (α)	Rank Corr. (ρ)
H100 (Ref)	756.5	1.00×	1.000
H200	348.1	0.46×	0.995
A100	322.5	0.43×	0.989
RTX 4090	327.4	0.43×	0.989
V100	138.9	0.18×	0.956

Table 2: **Robustness of PTE across Hardware Profiles.** High correlations ($\rho > 0.95$) demonstrate that PTE maintains consistent model efficiency rankings despite significant variations in hardware specifications.

($r = 0.9253$, $N = 100$, $p < 10^{-4}$) with wall-clock time, significantly outperforming both naive token counts and commercial API pricing metrics. (See Appendix B.3 and G).¹

4.2 Robustness across Hardware Profiles

We performed a sensitivity analysis of the coefficient γ across devices with varying memory-to-compute ratios. Table 2 summarizes the derived Hardware Operational Intensity (HOI) and the resulting scaling factor $\alpha = \gamma/\gamma_{\text{base}}$ for each device (See Appendix A.3 for detailed hardware specifications). Despite significant variations in hardware characteristics (where α ranges from $0.18\times$ to $1.00\times$), the relative efficiency rankings of models remain highly consistent, with Spearman’s rank correlation coefficients (ρ) consistently exceeding 0.95. This confirms that PTE captures intrinsic efficiency characteristics independent of the specific deployment platform.

5 Experimental Setup

To evaluate TIR efficiency differences using PTE costs, we designed a comprehensive experimental framework covering diverse TIR tasks and models.

5.1 Benchmarks Formulation

We evaluate models on five benchmarks targeting distinct TIR capabilities. For mathematical reasoning, we use the full test sets of **MATH500**(Hendrycks et al., 2021) (math problems) and **AIME 2024/2025** (high-difficulty competition problems), equipping agents with a Python tool. For information seeking, we evaluate on randomly sampled subsets of 500 examples from **SimpleQA**(Wei et al., 2024) (factual QA requiring spe-

¹Note that step-level statistics are reported for descriptive purposes only, since steps within a trajectory are serially dependent, p-values should not be interpreted as formal hypothesis tests.

cific fact retrieval) and **WebInstruct-Verified**(Ma et al., 2025) (complex, multi-disciplinary tasks requiring both retrieval and computation). The former utilizes Search and Visit tools, while the latter adds Python to the toolset. Detailed descriptions of each benchmark are provided in Appendix B.1.

5.2 Models and Agent Framework

We evaluated the tool-capable open-source models listed in Table 3 using vLLM. To ensure a fair comparison of intrinsic capabilities, we utilized identical system prompts and tool definitions across all models within our unified framework.

The framework was equipped with three specific tools: (1) a Search tool utilizing Serper API, (2) a Visit tool for webpage content retrieval powered by Jina API, and (3) a Python tool implemented via an open-source Python sandbox.

Our framework orchestrates the iterative TIR process and logs token-level statistics for evaluation. Specifically, at each turn i , it records the pre-fill tokens D_{prefill_i} , decoded tokens D_{decode_i} , and cumulative sequence length L_{seq_i} (see Appendix B.2 for implementation details).

6 Results and Analysis

6.1 Main Results

Efficiency vs. Accuracy Landscape. We visualize the main results in Fig. 4, which plots the accuracy of the model against PTE. In this visualization, the bottom-right zone represents better performance in both metrics. More detailed experimental results are available in the Appendix C.1. From the results we conclude three findings about TIR.

- **Vast Cost gap:** Fig. 4 shows that while many models cluster at high accuracy (e.g, $\sim 70\%$ on AIME24), their corresponding PTE values can span an order of magnitude or more.
- **Task-Specific TIR Abilities:** Our data shows that TIR capability is not a general skill but is highly specialized by task and tool type. For example, Qwen2.5-72B excels as a web agent (SimpleQA) but performs poorly as a Python reasoner (MATH500 & AIME). This suggests that TIR agent evaluations may be task-specific.
- **Different TIR Behavioral Patterns across Models:** We identify five types of different model TIR patterns from the results.

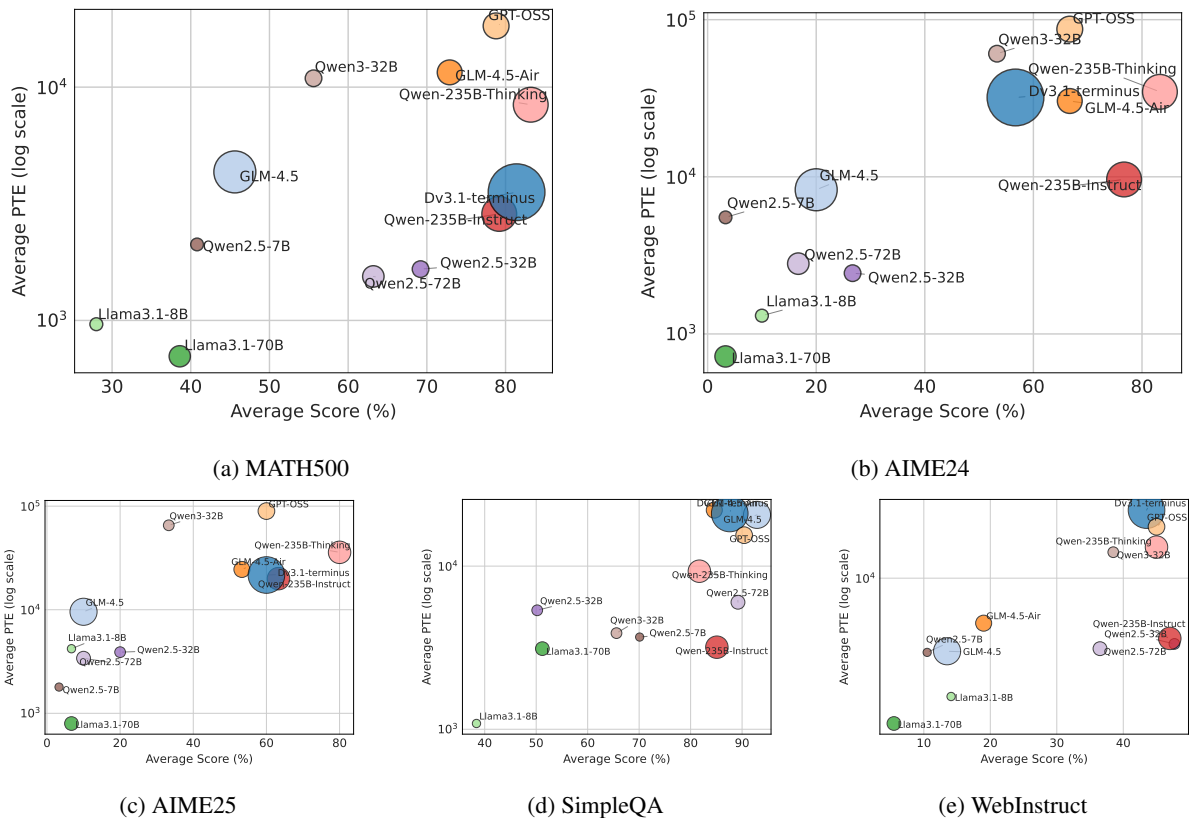


Figure 4: **PTE (Prefill Token Equivalents) versus Average Score on five benchmarks.** The bubble size represents the scale of active parameters. Models in the bottom-right region exhibit better trade-offs between efficiency and accuracy. Note the logarithmic scale on the y-axis.

- **Llama-3.1 series** are pure instruct models. Internal thinking tokens are negligible on every task, yielding high efficiency and medium accuracy.
- **Qwen-2.5 series** achieve high efficiency with moderate accuracy, exhibiting a task-adaptive reasoning strategy. On SimpleQA, the models suppress reasoning to trigger early toolcalls. Conversely, on the four more challenging datasets, they allocate a significant token budget to the initial step, thereby delaying the first tool invocation (See Fig. 8 in Appendix). We label this phenomenon the “first-step effect” and further analyze it in Sec. 6.2.
- **Qwen3-235B-Thinking** and **Qwen3-32B** (default config) are models with thinking inference mode enabled. They show comparatively high token costs and PTE costs compared to those under standard inference mode (Fig. 8d and 5d). On AIME25, the Qwen3-235B-Thinking model yields a +16.7% accuracy gain over Qwen3-235B-Instruct at only 1.8x PTE cost. However,

on the easy SimpleQA task the same model drops 3.4% in accuracy while PTE grows by 4.2x, indicating severe over-thinking. **Thinking mode is therefore beneficial only when the task difficulty justifies the extra compute.**

- **Deepseek-V3.1-Terminus**, **GPT-OSS** and **Qwen3-235B-Instruct** are frontier models that exhibit relatively high accuracy across most domains. However, the first two are mainly due to their powerful consistent tool-use capacity, whose lengthy and multi-round tool responses noticeably harm efficiency. This suggests that only when agent infrastructures resolve the tooluse KV-Cache eviction problem can these powerful models fully unleash their potential.
- **Tongyi-Deeprersearch** is omitted from Fig. 5 and Appendix Fig. 8 due to recurrent tool format collapse, resulting in an extremely high PTE. This case is discussed in detail in Sec. 6.2.

out obvious accuracy gain over its peers(See Fig. 4e and Fig. 6). We identify tool-mixing phenomenon as a kind of **inefficient over-tooluse**.

- **Lack of Tool Priors** Fig. 4 shows that on MATH and AIME the **Qwen-2.5 series** models tend to lose accuracy and incur higher PTE costs once the Python tool is enabled. We suspect that this reflects limited prior exposure to Python tool-use during pre-training. An example trajectory in Appendix D.3 illustrates how the model often forgets to print the final value, so no output is returned, suggesting that a tool is most efficient when the model has already learned how to use it.

- **Tool Format Collapse** Tongyi-Deeprersearch carries one of the highest PTEs we recorded due to recurrent tool format collapse, as shown in Table 5 and Table 6. It was fine-tuned on a fixed schema and fails when the wording changes even slightly: renaming the tool from “**search**” to “**google_search_tool**”, or sending a single query instead of the expected list queries, is enough to break the call. These syntactic mismatches lead to significant failures in both accuracy and efficiency. Example trajectories are provided in Appendix D.4. This is a clear manifestation of losing generalization in TIR, where performance is not robust to non-semantic changes in the input format.

We observe the same tool formats sensitivity in the **Qwen-2.5 series**. Because it was trained on Python blocks wrapped in triple backticks (`python ...`), it occasionally reverts to this style instead of using the required `<tool_call>` tag. This leads to parsing failures from the vLLM engine.

Quantitative breakdowns of these patterns including their detection heuristics, occurrence frequencies, and PTE cost multipliers are detailed in Appendix E.

7 From Efficiency to Intelligence

As shown in Fig. 7, trajectories reaching the correct answer consistently exhibit lower PTE values than those finishing with an error across most models and benchmarks. We observe that erroneous trajectories frequently involve prolonged interaction patterns, including repeated toolcalls and extended intermediate reasoning steps, which coincide with

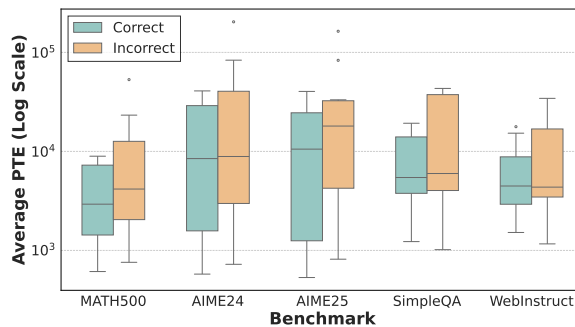


Figure 7: **Distribution of PTEs for correct and incorrect trajectories across five benchmarks.** Incorrect trajectories(right bars) consistently exhibit higher PTE compared to correct ones(left bars). Note the logarithmic scale on the y-axis.

elevated PTE cost. Detailed statistical breakdowns for each model and benchmark are provided in Appendix C.2, while Appendix F presents difficulty-stratified analyses confirming that these efficiency gaps persist independently of problem hardness.

Qualitative inspection suggests that unsuccessful trajectories often include toolcalls that return unhelpful or irrelevant information, followed by additional steps that attempt to reconcile or reinterpret prior outputs. Such behaviors tend to lengthen the effective sequence length (L_{seq}), expand the KV-cache, and are therefore accompanied by higher PTE values. In contrast, many correct trajectories exhibit more compact interaction patterns, with fewer detours and shorter contexts.

Importantly, these observations are correlational. Rather than implying that higher PTE causes errors, our results indicate that elevated PTE frequently co-occurs with reasoning processes marked by uncertainty, redundancy, or ineffective tool utilization. From this perspective, PTE may serve as a coarse-grained diagnostic signal for identifying potentially inefficient reasoning trajectories.

8 Conclusion

In this paper, we introduce PTE, a hardware-aware metric unifying reasoning and tool-use costs by modeling the prefill-decode asymmetry. Experiments across models and tasks reveal that TIR costs span orders of magnitude, largely driven by four specific inefficiency patterns. We also find a correlation between low PTE and high accuracy of trajectory. Hopefully, these findings may provide a unified view of TIR efficiency and offer a new perspective for future work in this field.

9 Limitations

Several limitations exist. Firstly, PTE measures transformer computation, omitting real-world costs like API latency. Secondly, the γ parameter represents a simplified abstraction of architectural efficiency; while useful, it may not fully capture all nuances of specific hardware optimizations or runtime dynamics. Finally, our empirical validation was limited to specific tasks and models; the generalizability of our findings, particularly the link between low PTE and high-quality reasoning, requires further exploration across broader domains and architectures.

10 Ethical Considerations

This work introduces a new metric to evaluate the efficiency of Large Language Models in tool-integrated reasoning scenarios. We believe that our proposed metric, PTE, contributes positively to the community by enabling researchers and practitioners to identify computational inefficiencies, potentially leading to reduced energy consumption and a lower carbon footprint for deploying LLM agents.

The datasets used in our experiments (MATH500, AIME, SimpleQA, etc.) are all publicly available and standard in the field. We adhered to the terms of use for all APIs and models employed during our evaluation. We do not foresee any immediate negative ethical consequences or societal harm resulting from the proposed metric or the analysis of inefficiency patterns.

Acknowledgments

This work was supported by the Anhui Provincial Natural Science Foundation under Grant 2108085UD12. We acknowledge the support of GPU cluster built by MCC Lab of Information Science and Technology Institution, USTC. The AI-driven experiments, simulations and model training were performed on the robotic AI-Scientist platform of Chinese Academy of Sciences.

References

Kedi Chen, Zhikai Lei, Xu Guo, Xuecheng Wu, Siyuan Zeng, Jianghao Yin, Yinqi Zhang, Qin Chen, Jie Zhou, Liang He, Qipeng Guo, Kai Chen, and Wei Zhang. 2025a. [Code-driven number sequence calculation: Enhancing the inductive reasoning abilities of large language models](#). *Preprint*, arXiv:2510.14620.

Kedi Chen, Dezhao Ruan, Yuhao Dan, Yaoting Wang, Siyu Yan, Xuecheng Wu, Yinqi Zhang, Qin Chen, Jie Zhou, Liang He, Biqing Qi, Linyang Li, Qipeng Guo, Xiaoming Shi, and Wei Zhang. 2025b. [A survey of inductive reasoning for large language models](#). *Preprint*, arXiv:2510.10182.

Qianben Chen, Jingyi Cao, Jiayu Zhang, Tianrui Qin, Xiaowan Li, King Zhu, Dingfeng Shi, He Zhu, Minghao Liu, Xiaobo Liang, Xin Gui, Ge Zhang, Jian Yang, Yuchen Eleanor Jiang, and Wangchunshu Zhou. 2025c. [A²fm: An adaptive agent foundation model for tool-aware hybrid reasoning](#). *Preprint*, arXiv:2510.12838.

Yifei Chen, Guanting Dong, and Zhicheng Dou. 2025d. [Toward effective tool-integrated reasoning via self-evolved preference learning](#). *Preprint*, arXiv:2509.23285.

Zehui Chen, Weihua Du, Wenwei Zhang, Kuikun Liu, Jiangning Liu, Miao Zheng, Jingming Zhuo, Songyang Zhang, Dahua Lin, Kai Chen, and 1 others. 2023. [T-eval: Evaluating the tool utilization capability step by step](#). *CoRR*.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.

Lu Dai, Yijie Xu, Jinhui Ye, Hao Liu, and Hui Xiong. 2025. [Seper: Measure retrieval utility through the lens of semantic perplexity reduction](#). *Preprint*, arXiv:2503.01478.

Guanting Dong, Yifei Chen, Xiaoxi Li, Jiajie Jin, Hongjin Qian, Yutao Zhu, Hangyu Mao, Guorui Zhou, Zhicheng Dou, and Ji-Rong Wen. 2025a. [Tool-star: Empowering LLM-brained multi-tool reasoner via reinforcement learning](#). *Preprint*, arXiv:2505.16410.

Guanting Dong, Hangyu Mao, Kai Ma, Licheng Bao, Yifei Chen, Zhongyuan Wang, Zhongxia Chen, Jiazhen Du, Huiyang Wang, Fuzheng Zhang, Guorui Zhou, Yutao Zhu, Ji-Rong Wen, and Zhicheng Dou. 2025b. [Agentic reinforced policy optimization](#). *Preprint*, arXiv:2507.19849.

Xuanqi Gao, Siyi Xie, Juan Zhai, Shiqing Ma, and Chao Shen. 2025. [Mcp-radar: A multi-dimensional benchmark for evaluating tool use capabilities in large language models](#). *Preprint*, arXiv:2505.16700.

Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujie Yang, Minlie Huang, Nan Duan, and Weizhu Chen. 2024. [Tora: A tool-integrated reasoning agent for mathematical problem solving](#). *Preprint*, arXiv:2309.17452.

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. *Preprint*, arXiv:2103.03874.
- Shiting Huang, Zhen Fang, Zehui Chen, Siyu Yuan, Junjie Ye, Yu Zeng, Lin Chen, Qi Mao, and Feng Zhao. 2025. **Criticoool: Evaluating self-critique capabilities of large language models in tool-calling error scenarios.** *Preprint*, arXiv:2506.13977.
- Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. 2024. **Swe-bench: Can language models resolve real-world github issues?** *Preprint*, arXiv:2310.06770.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. **Scaling laws for neural language models.** *Preprint*, arXiv:2001.08361.
- Hanchen Li, Qiuyang Mang, Runyuan He, Qizheng Zhang, Huanzhi Mao, Xiaokun Chen, Alvin Cheung, Joseph Gonzalez, and Ion Stoica. 2025a. **Continuum: Efficient and robust multi-turn LLM agent scheduling with kv cache time-to-live.** *Preprint*, arXiv:2511.02230.
- Kuan Li, Zhongwang Zhang, Huifeng Yin, Rui Ye, Yida Zhao, Liwen Zhang, Litu Ou, Dingchu Zhang, Xixi Wu, Jialong Wu, Xinyu Wang, Zile Qiao, Zhen Zhang, Yong Jiang, Pengjun Xie, Fei Huang, and Jingren Zhou. 2025b. **Websailor-v2: Bridging the chasm to proprietary agents via synthetic data and scalable reinforcement learning.** *Preprint*, arXiv:2509.13305.
- Kuan Li, Zhongwang Zhang, Huifeng Yin, Liwen Zhang, Litu Ou, Jialong Wu, Wenbiao Yin, Baixuan Li, Zhengwei Tao, Xinyu Wang, Weizhou Shen, Junkai Zhang, Dingchu Zhang, Xixi Wu, Yong Jiang, Ming Yan, Pengjun Xie, Fei Huang, and Jingren Zhou. 2025c. **Websailor: Navigating super-human reasoning for web agent.** *Preprint*, arXiv:2507.02592.
- Minghao Li, Yingxiu Zhao, Bowen Yu, Feifan Song, Hangyu Li, Haiyang Yu, Zhoujun Li, Fei Huang, and Yongbin Li. 2023. **Api-bank: A comprehensive benchmark for tool-augmented LLMs.** *Preprint*, arXiv:2304.08244.
- Xuefeng Li, Haoyang Zou, and Pengfei Liu. 2025d. **ToRL: Scaling tool-integrated RL.** *Preprint*, arXiv:2503.23383.
- Heng Lin and Zhongwen Xu. 2025. **Understanding tool-integrated reasoning.** *Preprint*, arXiv:2508.19201.
- Shuhang Lin, Wenyue Hua, Lingyao Li, Che-Jui Chang, Lizhou Fan, Jianchao Ji, Hang Hua, Mingyu Jin, Jiebo Luo, and Yongfeng Zhang. 2024. **Battleagent: Multi-modal dynamic emulation on historical battles to complement historical analysis.** *Preprint*, arXiv:2404.15532.
- Jiayu Liu, Cheng Qian, Zhaochen Su, Qing Zong, Shijue Huang, Bingxiang He, and Yi R. Fung. 2025. **Cost-bench: Evaluating multi-turn cost-optimal planning and adaptation in dynamic environments for LLM tool-use agents.** *Preprint*, arXiv:2511.02734.
- Yi Lu, Jianing Wang, Linsen Guo, Wei He, Hongyin Tang, Tao Gui, Xuanjing Huang, Xuezhi Cao, Wei Wang, and Xunliang Cai. 2025. **R-horizon: How far can your large reasoning model really go in breadth and depth?** *Preprint*, arXiv:2510.08189.
- Xueguang Ma, Qian Liu, Dongfu Jiang, Ge Zhang, Zejun Ma, and Wenhua Chen. 2025. **General-reasoner: Advancing LLM reasoning across all domains.** *Preprint*, arXiv:2505.14652.
- Grégoire Mialon, Clémentine Fourier, Craig Swift, Thomas Wolf, Yann LeCun, and Thomas Scialom. 2023. **Gaia: a benchmark for general ai assistants.** *Preprint*, arXiv:2311.12983.
- Zaifeng Pan, Ajikumar Patel, Zhengding Hu, Yipeng Shen, Yue Guan, Wan-Lu Li, Lianhui Qin, Yida Wang, and Yufei Ding. 2025. **Kvflow: Efficient prefix caching for accelerating LLM-based multi-agent workflows.** *Preprint*, arXiv:2507.07400.
- Shishir G Patil, Huanzhi Mao, Fanjia Yan, Charlie Cheng-Jie Ji, Vishnu Suresh, Ion Stoica, and Joseph E Gonzalez. The berkeley function calling leaderboard (bfcl): From tool use to agentic evaluation of large language models. In *Forty-second International Conference on Machine Learning*.
- Han Peng, Peiyu Liu, Zican Dong, Daixuan Cheng, Junyi Li, Yiru Tang, Shuo Wang, and Wayne Xin Zhao. 2025. **How efficient are diffusion language models? a critical examination of efficiency evaluation practices.** *Preprint*, arXiv:2510.18480.
- Cheng Qian, Emre Can Acikgoz, Hongru Wang, Xiusi Chen, Avirup Sil, Dilek Hakkani-Tür, Gokhan Tur, and Heng Ji. 2025. **Smart: Self-aware agent for tool overuse mitigation.** *Preprint*, arXiv:2502.11435.
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2023. **ToolLLM: Facilitating large language models to master 16000+ real-world apis.** *Preprint*, arXiv:2307.16789.
- Joykirat Singh, Raghav Magazine, Yash Pandya, and Akshay Nambi. 2025. **Agentic reasoning and tool integration for LLMs via reinforcement learning.** *Preprint*, arXiv:2505.01441.
- Tongyi DeepResearch Team. 2025. **Tongyi deepresearch technical report.** *Preprint*, arXiv:2510.24701.
- Hongru Wang, Cheng Qian, Wanjun Zhong, Xiusi Chen, Jiahao Qiu, Shijue Huang, Bowen Jin, Mengdi Wang, Kam-Fai Wong, and Heng Ji. 2025a. **Acting less is reasoning more! teaching model to act efficiently.** *Preprint*, arXiv:2504.14870.

- Ke Wang, Houxing Ren, Aojun Zhou, Zimu Lu, Sichun Luo, Weikang Shi, Renrui Zhang, Linqi Song, Mingjie Zhan, and Hongsheng Li. 2023. [Mathcoder: Seamless code integration in LLMs for enhanced mathematical reasoning](#). *Preprint*, arXiv:2310.03731.
- Ningning Wang, Xavier Hu, Pai Liu, He Zhu, Yue Hou, Heyuan Huang, Shengyu Zhang, Jian Yang, Jiaheng Liu, Ge Zhang, Changwang Zhang, Jun Wang, Yuchen Eleanor Jiang, and Wangchunshu Zhou. 2025b. [Efficient agents: Building effective agents while reducing cost](#). *Preprint*, arXiv:2508.02694.
- Jason Wei, Nguyen Karina, Hyung Won Chung, Yunxin Joy Jiao, Spencer Papay, Amelia Glaese, John Schulman, and William Fedus. 2024. [Measuring short-form factuality in large language models](#). *Preprint*, arXiv:2411.04368.
- Jason Wei, Zhiqing Sun, Spencer Papay, Scott McKinney, Jeffrey Han, Isa Fulford, Hyung Won Chung, Alex Tachard Passos, William Fedus, and Amelia Glaese. 2025a. [Browsecomp: A simple yet challenging benchmark for browsing agents](#). *Preprint*, arXiv:2504.12516.
- Yifan Wei, Xiaoyan Yu, Yixuan Weng, Tengfei Pan, Angsheng Li, and Li Du. 2025b. [Autotir: Autonomous tools integrated reasoning via reinforcement learning](#). *Preprint*, arXiv:2507.21836.
- Samuel Williams, Andrew Waterman, and David Patterson. 2009. [Roofline: an insightful visual performance model for multicore architectures](#). *Commun. ACM*, 52(4):65–76.
- Ryan Wong, Jiawei Wang, Junjie Zhao, Li Chen, Yan Gao, Long Zhang, Xuan Zhou, Zuo Wang, Kai Xi-ang, Ge Zhang, Wenhao Huang, Yang Wang, and Ke Wang. 2025. [Widesearch: Benchmarking agentic broad info-seeking](#). *Preprint*, arXiv:2508.07999.
- Yang Xiao, Mohan Jiang, Jie Sun, Keyu Li, Jifan Lin, Yumin Zhuang, Ji Zeng, Shijie Xia, Qishuo Hua, Xuefeng Li, Xiaojie Cai, Tongyu Wang, Yue Zhang, Liming Liu, Xia Wu, Jinlong Hou, Yuan Cheng, Wenjie Li, Xiang Wang, and 2 others. 2025. [Limi: Less is more for agency](#). *Preprint*, arXiv:2509.17567.
- Frank F. Xu, Yufan Song, Boxuan Li, Yuxuan Tang, Kritanjali Jain, Mengxue Bao, Zora Z. Wang, Xuhui Zhou, Zhitong Guo, Murong Cao, Mingyang Yang, Hao Yang Lu, Amaad Martin, Zhe Su, Leander Maben, Raj Mehta, Wayne Chi, Lawrence Jang, Yiqing Xie, and 2 others. 2025. [Theagentcompany: Benchmarking LLM agents on consequential real world tasks](#). *Preprint*, arXiv:2412.14161.
- Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. 2022. [Webshop: Towards scalable real-world web interaction with grounded language agents](#). *Advances in Neural Information Processing Systems*, 35:20744–20757.
- Dian Yu, Yulai Zhao, Kishan Panaganti, Linfeng Song, Haitao Mi, and Dong Yu. 2025a. [Every question has its own value: Reinforcement learning with explicit human values](#). *Preprint*, arXiv:2510.20187.
- Yuanqing Yu, Zhefan Wang, Weizhi Ma, Shuai Wang, Chuhan Wu, Zhiqiang Guo, and Min Zhang. 2025b. [Steptool: Enhancing multi-step tool usage in LLMs via step-grained reinforcement learning](#). In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management, CIKM '25*, page 3952–3962. ACM.
- Zhihang Yuan, Yuzhang Shang, Yang Zhou, Zhen Dong, Zhe Zhou, Chenhao Xue, Bingzhe Wu, Zhikai Li, Qingyi Gu, Yong Jae Lee, Yan Yan, Beidi Chen, Guangyu Sun, and Kurt Keutzer. 2024. [LLM inference unveiled: Survey and roofline model insights](#). *Preprint*, arXiv:2402.16363.
- Michael J. Zellinger and Matt Thomson. 2025. [Economic evaluation of LLMs](#). *Preprint*, arXiv:2507.03834.
- Yufeng Zhao, Junnan Liu, Hongwei Liu, Dongsheng Zhu, Yuan Shen, Songyang Zhang, and Kai Chen. 2025. [Dissecting tool-integrated reasoning: An empirical study and analysis](#). *Preprint*, arXiv:2508.15754.
- Yinmin Zhong, Shengyu Liu, Junda Chen, Jianbo Hu, Yibo Zhu, Xuanzhe Liu, Xin Jin, and Hao Zhang. 2024. [{DistServe}: Disaggregating prefill and decoding for goodput-optimized large language model serving](#). In *18th USENIX Symposium on Operating Systems Design and Implementation (OSDI 24)*, pages 193–210.
- Yuchen Zhuang, Yue Yu, Kuan Wang, Haotian Sun, and Chao Zhang. 2023. [Toolqa: A dataset for LLM question answering with external tools](#). *Preprint*, arXiv:2306.13304.

A Metric Formulation and Hardware Grounding

A.1 Model Architectural Details

Due to space constraints, the main text presents a simplified view of the model specifications. Table 3 provides the complete architectural details for all evaluated models, including the number of layers (n_{layers}), hidden dimensions (d_{model}), and the architectural factor (H_{kv}/H_q) used to calculate the γ coefficient.

A.2 Hardware Operational Intensity Details

To translate memory size into a computational cost, we employ an important concept from the field of high-performance computing (HPC): Hardware Operational Intensity (HOI) (Williams et al., 2009; Yuan et al., 2024; Peng et al., 2025). HOI is an architectural constant that defines the ridge point in the hardware Roofline model. It represents the ratio between the hardware’s peak computational throughput (FLOPs/s) and its peak memory bandwidth (Bytes/s).

$$HOI = \frac{\text{Peak FLOPs}}{\text{Peak Memory Bandwidth}}. \quad (5)$$

For consistency, we utilize the NVIDIA H100 80GB PCIe (Hopper architecture) as our reference hardware. These numerator and denominator values are architectural constants provided in official vendor documentation, such as the **NVIDIA H100 Tensor Core GPU Datasheet**.

Peak Memory Bandwidth: The official specifications for the H100 80GB PCIe (HBM2e) version indicate a peak memory bandwidth of 2.0 TB/s (i.e., 2.0×10^{12} Bytes/s).

Peak Computational Performance: For the FP16/BF16 precision, the H100 specification sheet indicates a peak theoretical throughput of approximately 1,513 TFLOPS, a figure that accounts for architectural features such as the Transformer Engine. As the Roofline model’s definition of HOI relies on the architecture’s theoretical performance ceiling, we utilize this peak value (i.e., $1,513 \times 10^{12}$ FLOPs/s) for the numerator.

HOI Calculation:

$$\begin{aligned} HOI &= \frac{1,513 \times 10^{12} \text{ FLOPs/s}}{2.0 \times 10^{12} \text{ Bytes/s}} \\ &= \mathbf{756.5 \text{ FLOPs/Byte}}. \end{aligned} \quad (6)$$

Thus, we use the value 756.5 to represent the theoretical peak operational intensity for the H100

architecture at FP16/BF16 precision. We utilize this architectural constant to anchor our γ coefficient.

A.3 Sensitivity Analysis of γ Across Hardware Architectures

To assess the robustness of the PTE metric across different hardware environments, we conducted a sensitivity analysis of the parameter γ . As defined before, γ is linearly proportional to the Hardware Operational Intensity (HOI):

$$\gamma \propto HOI = \frac{\text{Peak FLOPs}}{\text{Peak Memory Bandwidth}}. \quad (7)$$

We define a scaling factor $\alpha = \gamma/\gamma_{\text{base}}$, using the NVIDIA H100 as the baseline. Table 4 presents the comparative analysis across various architectures, including the H200, A100, V100, and RTX 4090. The four devices are selected to give the sweep a footprint as wide as possible in the real-world hardware landscape. The variation in α reflects the evolving battle against the "Memory Wall." However, the newer H200 significantly reduces γ ($0.46\times$) by more than doubling the memory bandwidth, demonstrating specific hardware optimizations targeting memory bottlenecks.

To verify the stability of our efficiency metric under these varying conditions, we re-evaluated the model rankings on the WebInstruct-Verified dataset trajectories by simulating PTE costs across different hardware profiles. Despite significant variations in hardware characteristics (where α ranges from $0.18\times$ to $1.00\times$), the relative efficiency rankings of the models remain highly consistent. This is evidenced by the Spearman’s rank correlation coefficients (ρ) reported in Table 4, which consistently exceed 0.95 across all devices.

In summary, while absolute PTE costs fluctuate with hardware specifications, the PTE metric can still capture some intrinsic efficiency characteristics of TIR models, indicating its potential to be an indicator independent of the specific deployment platform.

A.4 Interpreting γ : Computational Cost Scaling Efficiency

Our analysis suggests that γ should be interpreted as a measure of computational cost scaling efficiency rather than absolute per-token latency. A low γ may not indicate a faster model in terms of

Models	Active Params (N_{params})	Layers (n_{layers})	Hidden Dim (d_{model})	Architecture Factor (H_{kv}/H_q)	Calculated γ
Open-Source Dense Models					
Qwen2.5-7B-Instruct	6.53B	28	3584	1/7	0.00329
Qwen2.5-32B-Instruct	31.0B	64	5120	1/2	0.00320
Qwen2.5-72B-Instruct	70.0B	80	8192	1/8	0.00175
Qwen3-32B	31.2B	64	5120	1/8	0.00200
Llama-3.1-8B-Instruct	8.0B	32	4096	1/4	0.00625
Llama-3.1-70B-Instruct	70.6B	80	8192	1/8	0.00175
Open-Source MoE Models					
Qwen3-30B-A3B	3.3B	48	2048	1/8	0.00563
Qwen3-235B-A22B-Instruct	22B	94	4096	1/16	0.00163
Qwen3-235B-A22B-Thinking	22B	94	4096	1/16	0.00163
GLM-4.5-Air	12B	46	4096	1/12	0.00200
GLM-4.5	32B	92	5120	1/12	0.00183
DeepSeek-V3.1-Terminus	37B	61	576	–	0.00068
GPT-OSS-120B	5.1B	36	2880	1/8	0.00388

¹ All architectural parameters are sourced from the models’ public config.json files.

² Llama 3.1 series does not report active parameters. We assume active parameters equals to total parameters.

³ For DeepSeek-V3, the d_{model} is substituted by $d_{latent} + d_{rope}$, reflecting the significant KV-Cache optimization achieved by its Multi-Head Latent Attention (MLA) architecture.

³ We use the NVIDIA H100 PCIe as our reference hardware for consistency, which has a Hardware Operational Intensity (HOI) of approximately 756.5 FLOPs/Byte, as detailed in Appendix A.2.

Table 3: **Gamma Values of Different LLMs.** We evaluated a range of state-of-the-art open-source models, all of which were officially declared as having tool-calling capabilities.

tokens-per-second, but rather that the model’s total infer cost grows more slowly as the sequence length (L_{seq}) increases. This architectural trade-off is more evident when comparing different Mixture-of-Experts (MoE) models from Table 3.

For instance, **GPT-OSS-120B** exhibits a relatively high γ (≈ 0.00388), which seems to be correlated with its computation-light architecture (5.1B activated parameters). This design appears to be more sensitive to the memory overhead of the KV cache, leading to less efficient scaling. Conversely, a computation-heavy model such as **DeepSeek-V3.1-Terminus** (37B activated parameters) presents a much lower γ (≈ 0.00068), suggesting its cost scales more slowly. This implies that its massive computational cost significantly outweighs its memory cost at longer context lengths. This comparison indicates that γ may be a useful metric to identify architectures that are more efficient in non-reusable KV-Cache and long-context scenarios.

B Experimental Setup

B.1 Benchmark Details

Our benchmark suite was selected to assess LLMs’ TIR capabilities, requiring both internal reasoning and external tool use. Recent surveys on inductive reasoning abilities of LLMs (Chen et al., 2025b)

provide comprehensive taxonomies of reasoning patterns that underpin these benchmark designs. We evaluated the common LLMs on five distinct benchmarks, each targeting different capability aspects:

- **MATH500**(Hendrycks et al., 2021): A benchmark focused on **mathematical reasoning**, consisting of a collection of math word problems that require numerical computation. We provide Python tool for this benchmark.
- **AIME24, AIME25**: A **competition math** benchmark featuring high-difficulty problems from the American Invitational Mathematics Examination. We provide Python tool for these benchmarks.
- **SimpleQA**(Wei et al., 2024): A benchmark for **factual question answering**. It contains questions that require the retrieval of specific facts from external knowledge sources. We provide Search and Visit tools for this benchmark. For our evaluation, we use a randomly sampled subset of 500 examples to ensure efficient yet representative testing.
- **Webinstruct-Verified**(Ma et al., 2025): A complex, **multi-disciplinary QA** benchmark. It is designed to test models on tasks that may require a

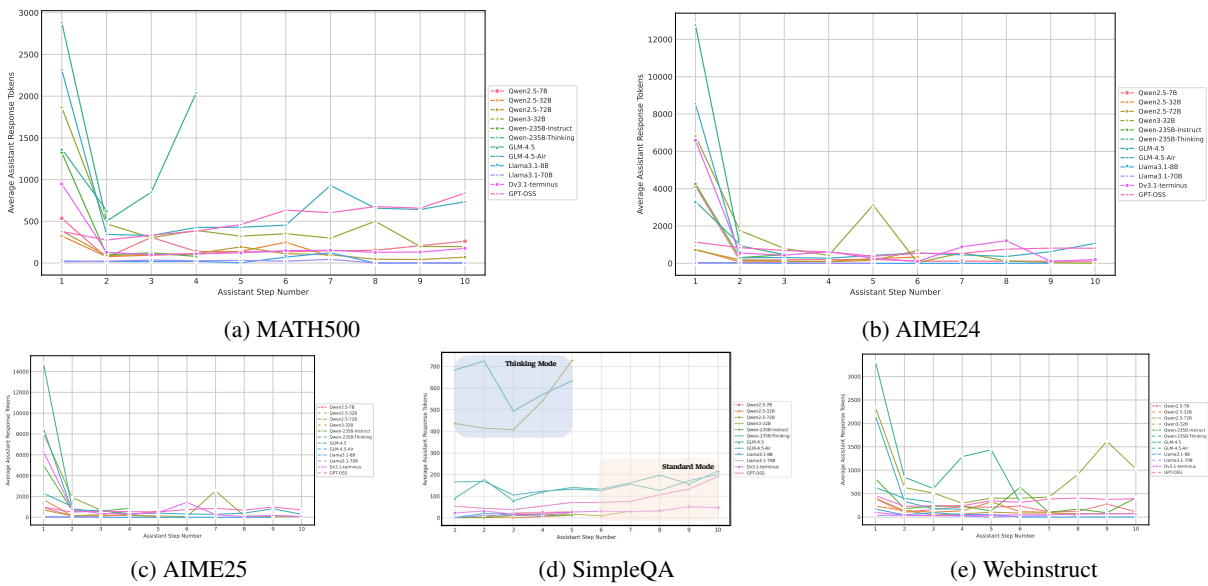


Figure 8: **Distribution of Average Assistant Response Tokens Across Reasoning Steps.** The figure illustrates the response length for each step in a reasoning trajectory. The phenomenon is described as "**first-step effect**" where models tend to "front-load" their computational budget.

System Prompt For Webinstruct-Verified

You are an **Expert Problem Solver**. Your mission is to accurately solve problems by skillfully integrating reasoning with your available tools: `python_tool`, `google_search_tool`, and `visit_tool`. As you proceed, adhere to the following principles:

- Strategic Tool Selection:** For math/data, you **MUST** use `python_tool`. For facts/real-world info, you **MUST** use `google_search_tool` and `visit_tool`. For complex problems, use web tools to gather data first, then `python_tool` to compute.
- Rigorous Verification:** Before giving a Final Answer, you **MUST** use `python_tool` to verify all calculations and use your web tools to cross-check all information for accuracy.
- Effective Tool Use:** Write clean, executable Python code, remembering to use `print()` for output. Formulate precise search queries for `google_search_tool` and specify a clear goal for `visit_tool`.
- Persistence:** Do not give up. Engage in as many tool calls as necessary to thoroughly explore the topic and find a complete, accurate answer.
- Final Answer Format:** After completing all reasoning and verification, provide the final answer and nothing else, wrapped in `<ANSWER>answer</ANSWER>`.

Figure 9: System Prompt Example

System Prompt For LLM Judge

Predicted Answer: predicted answer
 Expected Answer: ground truth
 Determine if the predicted answer is semantically correct (case-insensitive).
 Reply **only** a JSON object in the format: `{"correct": 0}` or `{"correct": 1}`
 Do not include any additional text or markdown formatting.

Figure 10: Answer Verify Prompt

combination of information retrieval and computation to answer. We provide Python and Search and Visit tools for this benchmark. For our evaluation, we use a randomly sampled subset of 500 examples to ensure efficient yet representative testing.

B.2 Implementation Details

We used the paid Serper and Jina APIs for the search and visit tools, and adopted the open-source SandboxFusion for python tool execution. All tool-call generation and parsing were handled by the vLLM inference engine, on which every tested model was deployed.

B.3 Latency Validation Setup

To rigorously evaluate the correlation between the proposed PTE metric and real-world latency, We conducted a high-throughput evaluation experiment to mimic industrial Tool-Integrated Reasoning (TIR) workloads. We deployed DeepSeek-V3.2 on an 8×H200 GPU node via vLLM (TP=8), simulating a high-concurrency scenario with 256 parallel requests. The evaluation employed a synthetic dataset of complex tool-requiring QAs derived from Wikidata, requiring the model to utilize Search, Visit, and Python. Crucially, we recorded the pure model generation latency, explicitly excluding time consumed by tool execution and network transmission. We then calculated the Pearson correlation between this generation time and both PTE and token counts. Sample data containing reasoning traces and timing statistics are released in our repository.

B.4 Prompt Details

An example system prompt that combines the three tools in the Webinstruct-Verified dataset is given in Fig. 9. At the end of each rollout trajectory, DeepSeek-V3 served as a LLM judge to check answer correctness. Its prompt is shown in Fig. 10.

C Comprehensive Empirical Results

C.1 Full Benchmark Performance

This section presents the full experimental results for all evaluated models across the five benchmarks. We report the detailed metrics including Top-1 Accuracy, average uncached token counts (sum of prefill and decode), tool-usage times, and the overall Average PTE. These values correspond to the aggregated performance analysis discussed in Sec.

6.1. The complete numerical results are listed in Table 5 and Table 6.

C.2 PTE Analysis by Reasoning Correctness

This section supplements the efficiency-accuracy analysis presented in Sec. 7. While Fig. 7 provides an aggregated visualization of the cost distribution, Fig. 11 presents the granular performance metrics for individual models.

D Qualitative Case Studies

D.1 Pattern I: Confirmatory Tool Usage

This example illustrates the Confirmatory Tool Usage behavior discussed in Sec. 6.2, where Qwen3-235B-Thinking use python tool to verify its internal thoughts rather than using it as direct solvers. See Fig. 12.

D.2 Pattern II: Tool-Mixing Example Trajectory

This example illustrates the "tool-mixing" behavior discussed in Sec. 6.2, where DeepSeek-V3.1-Terminus alternates between search and Python tools within a single trajectory on the Webinstruct-Verified benchmark. See Fig. 13.

D.3 Pattern III: Lack of Tool Priors

This example supports the analysis in Sec. 6.2. It shows a trajectory from Qwen-2.5-7B-Instruct on AIME24 benchmark, where the provision of a Python tool (for which it lacks strong pre-training) leads to an inefficient and incorrect reasoning path. This contrasts with specialized code-reasoning methods (Chen et al., 2025a) that optimize the inductive reasoning abilities of LLMs through dedicated code-driven training, suggesting that a tool is most efficient when the model has already learned how to use it. See Fig. 14.

D.4 Pattern IV: Tool Format Collapse

This example illustrates the sensitivity to tool format phenomenon discussed in Sec. 6.2. It shows a SimpleQA trajectory where Tongyi-Deepresearch fails to correctly invoke a tool when presented with a semantically identical but syntactically different tool definition from what it was trained on. See Fig. 15.

E Quantitative Analysis of Inefficiency Patterns

We measured the exact PTE cost of each inefficiency pattern. Although these patterns can co-occur, they appear in distinct model-task settings (Sec. 6.2) and do not always occur together. We

Hardware	Peak TFLOPS (FP16)	Memory BW (TB/s)	HOI (FLOPs/Byte)	Scaling Factor (α)	Spearman’s Corr (ρ)
NVIDIA H100 (Ref)	1,513	2.00	756.5	1.0 \times	1.000000
NVIDIA H200	1,617	4.80	348.1	0.46 \times	0.994505
NVIDIA A100	624	1.93	322.5	0.43 \times	0.989011
NVIDIA V100	125	0.90	138.9	0.18 \times	0.956044
NVIDIA RTX 4090	330	1.00	327.4	0.43 \times	0.989011

Table 4: **Hardware Performance and Scaling Comparison on WebInstruct-Verified Trajectories**

Models	MATH500				AIME24				AIME25			
	acc \uparrow	tokens \downarrow	tooluse \downarrow	PTE \downarrow	acc \uparrow	tokens \downarrow	tooluse \downarrow	PTE \downarrow	acc \uparrow	tokens \downarrow	tooluse \downarrow	PTE \downarrow
<i>Open-Source Dense Models</i>												
Qwen2.5-7B-Instruct	40.8	1808	1.178	2117	3.3	5086	0.500	5519	3.3	1461	0.133	1793
Qwen2.5-32B-Instruct	69.2	1517	1.840	1659	26.7	2040	1.267	2436	20.0	3382	2.467	3894
Qwen2.5-72B-Instruct	63.2	1668	1.264	1544	16.7	2855	1.267	2801	10.0	3237	1.367	3402
Qwen3-32B	55.6	7852	2.064	10912	53.3	27290	2.367	60374	33.3	34853	2.567	65356
Llama-3.1-8B-Instruct	28.0	844	0.922	963	10.0	1239	1.033	1307	10.0	3459	1.333	4203
Llama-3.1-70B-Instruct	38.6	718	1.162	702	3.3	735	1.100	719	6.7	815	1.000	797
<i>Open-Source MoE Models</i>												
Qwen3-235B-Instruct	79.2	3266	1.024	2861	76.7	9624	1.033	9618	63.3	13494	1.100	19956
Qwen3-235B-Thinking	83.2	7023	1.022	8406	83.3	25261	1.000	34735	80.0	26811	0.933	35829
GLM-4.5-Air	72.9	7982	1.766	11569	66.7	22717	2.500	30359	53.3	17836	1.633	24489
GLM-4.5	45.6	3621	0.960	4331	20.0	6478	0.833	8270	10.0	6182	0.967	9555
DeepSeek-V3.1-Term	81.4	4210	1.908	28203	56.7	31982	2.800	292724	60.0	21492	2.933	21639
GPT-OSS-120B	78.8	6363	2.604	18306	24016	5439	5.867	87049	60.0	25958	6.333	89627
Tongyi-Deeprersearch	77.6	7925	1.242	27387	60.0	28315	1.133	143937	56.7	23912	0.900	104796

¹ GPT-OSS-120B uses default medium reasoning level.

² Qwen3-32B uses under default generate config, which enables thinking mode.

Table 5: **Accuracy (%), average uncached total tokens(prefil and decode), tooluse times, and average PTE comparison (Part 1: MATH & AIME)**

therefore quantified each within its primary occurrence setting (model-task pair), where per-pattern overhead remains clearly measurable.

We defined heuristic rules to detect each pattern and measured the **Frequency** (% of trajectories exhibiting the pattern) and **Cost Multiplier** (PTE increase compared to pattern-free trajectories in the same setting). As shown in Table 7, these patterns are not merely qualitative quirks but significant drivers of computational inefficiency, with cost multipliers ranging from 1.77 \times to 2.42 \times .

Detection Heuristics. Confirmatory: Detected when the model generates an answer or final value in its reasoning text before invoking the Python tool for verification. **Tool-Mixing:** Detected when a trajectory uses more than one distinct tool type (e.g., both search and Python). **Format Collapse:** Detected when the output contains JSON parsing errors or schema violations. **Lack of Priors:** Detected when a tool call returns empty output or execution error due to missing print statements or incorrect syntax.

F PTE reflects reasoning inefficiency beyond query difficulty

To verify that high PTE reflects reasoning inefficiency rather than inherent problem difficulty, we conduct three stratified analyses on MATH500 (dataset-provided levels 1–5) and AIME25 (GPT-4o-annotated levels 4–5).

F.1 Intra-Level Analysis: Efficiency gaps persist when difficulty is controlled

If PTE primarily reflected question difficulty, success and failure at the same difficulty level would be expected to yield similar PTE distributions. Contradicting this hypothesis, incorrect trajectories exhibit significantly higher PTE than correct ones within every difficulty level (See Table 8).

By holding difficulty constant, these gaps suggest that reasoning efficiency remains a major source of PTE variation beyond difficulty. Notably, the largest occur at Level 1 (up to 793%), where failures are often associated with reasoning inefficiencies, as detailed in Section 6.2 and Appendix D.1.

Models	SimpleQA				WebInstruct-Verified			
	acc \uparrow	tokens \downarrow	tooluse \downarrow	PTE \downarrow	acc \uparrow	tokens \downarrow	tooluse \downarrow	PTE \downarrow
Open-Source Dense Models								
Qwen2.5-7B-Instruct	70.1	3533	2.409	3674	10.5	2589	1.485	3368
Qwen2.5-32B-Instruct	50.2	5300	4.020	5361	47.7	2953	1.799	3813
Qwen2.5-72B-Instruct	89.2	5909	2.380	6006	36.5	3407	1.775	3557
Qwen3-32B	65.6	2903	0.985	15573	38.5	3883	1.885	14635
Llama-3.1-8B-Instruct	38.4	1024	1.140	1085	14.1	1373	1.291	1764
Llama-3.1-70B-Instruct	51.2	3083	1.720	3120	5.5	1154	1.120	1187
Open-Source MoE Models								
Qwen3-235B-A22B-Instruct	85.1	3166	2.325	3184	47.0	3831	1.400	15772
Qwen3-235B-A22B-Thinking	81.7	6325	1.787	9306	45.0	9787	1.180	63154
GLM-4.5-Air	84.6	17992	4.860	22045	19.0	4854	1.065	5170
GLM-4.5	92.9	17122	4.455	20617	13.5	2964	1.005	3424
DeepSeek-V3.1-Terminus	87.6	20783	5.235	21023	43.6	26583	6.840	27137
GPT-OSS-120B	90.4	12641	3.480	15474	45.0	10640	3.740	21256
Tongyi-Deepresearch	4.8	14155	6.940	45677	21.5	16509	3.830	63154

¹ GPT-OSS-120B uses default medium reasoning level.

² Qwen3-32B uses under default generate config, which enables thinking mode.

Table 6: **Accuracy (%), average uncached total tokens(prefil and decode), tooluse times, and average PTE comparison (Part 2: SimpleQA & Webinstruct)**

Pattern	Dominant Setting	Heuristic Definition	Freq.	Cost
Confirmatory	Qwen3-235B-Instruct (MATH500)	Answer token generated pre-tool	81%	1.77 \times
Tool-Mixing	DeepSeek-V3.1-Term (WebInstruct)	Unique tool types > 1	59%	2.42 \times
Format Collapse	Tongyi-DeepResearch (SimpleQA)	Schema/JSON Error in output	100%	N/A
Lack of Priors	Qwen2.5-72B-Instruct (AIME24)	Empty or Error tool return	33%	2.15 \times

Table 7: **Quantitative breakdown of inefficiency patterns.** Frequency and cost multiplier measured within dominant model-task settings where each pattern occurs most distinctly.

F.2 Inter-Level Analysis: Efficiency dominates difficulty in determining cost

If PTE reflected only difficulty, same-level success and failure would show similar PTEs. Instead, failures consistently exhibit higher PTE across all levels (See Table 9).

Beyond aggregate trends, this holds within individual models: GPT-OSS uses 75% less PTE to correctly solve a Level-5 problem (24,302) than to fail on a Level-4 problem (97,183). This shows that reasoning efficiency can frequently outweigh intrinsic problem difficulty in determining overall PTE.

F.3 Statistical Isolation: PTE predicts accuracy independently of difficulty

Partial correlation (controlling for difficulty) reveals a significant negative PTE-accuracy association ($r = -0.040$, $p = 0.002$), indicating the relationship persists beyond difficulty effects. In conclusion, PTE reflects reasoning inefficiency be-

yond query difficulty.

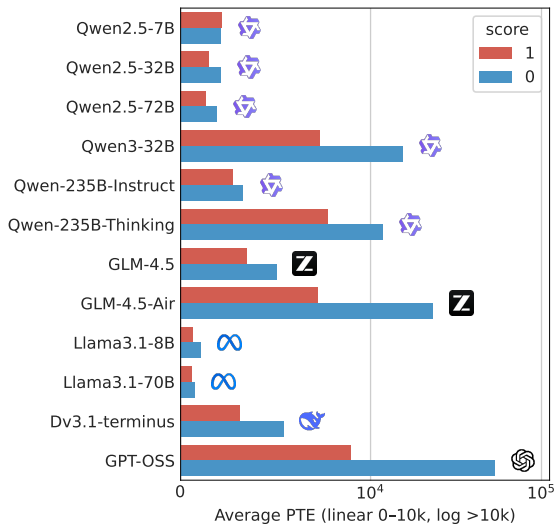
G Comparison with Commercial API Pricing

We compare PTE against commercial API pricing metrics to demonstrate its superiority in capturing non-linear hardware bottlenecks. We calculated API Cost using the standard formula:

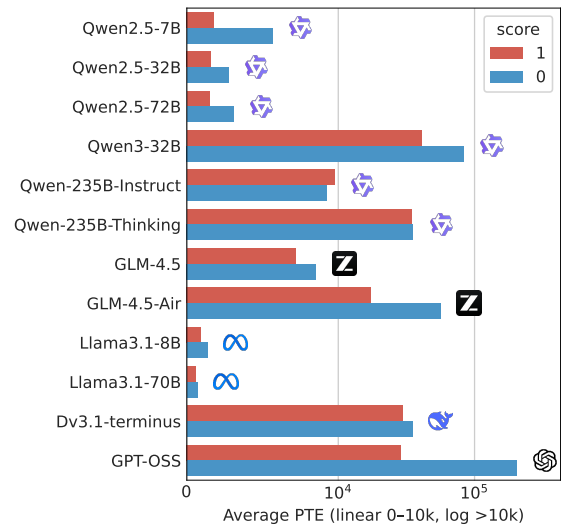
$$\text{Cost} = \text{Token}_{\text{in}} \times P_{\text{in}} + \text{Token}_{\text{out}} \times P_{\text{out}}, \quad (8)$$

where input tokens include the full context history (assuming cache miss) to reflect real-world context accumulation costs. We selected three representative industry price ratios ($P_{\text{in}} : P_{\text{out}}$): 1:1.5 (DeepSeek-V3.2), 1:3 (Standard), and 1:4 (GPT-4o/Claude 3.5).

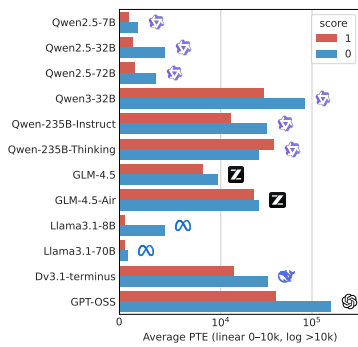
As shown in Table 10, PTE achieves the highest correlation with wall-clock latency ($r = 0.925$) because it models actual hardware bottlenecks (memory bandwidth constraints), not arbitrary commer-



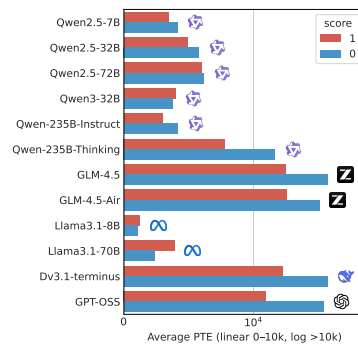
(a) MATH500



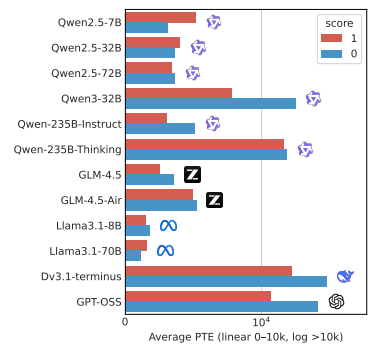
(b) AIME24



(c) AIME25



(d) SimpleQA



(e) Webinstruct

Figure 11: Comparison of PTE between Correct (red) and Incorrect (blue) Reasoning Trajectories. Incorrect reasoning is consistently associated with higher PTE across models and tasks.

cial pricing strategies. Its hardware-adaptive design ensures robustness across any deployment scenario.

Question: Jen enters a lottery by picking 4 distinct numbers from $S = \{1, 2, 3, \dots, 9, 10\}$. 4 numbers are randomly chosen from S . She wins a prize if at least two of her numbers were 2 of the randomly chosen numbers, and wins the grand prize if all four of her numbers were the randomly chosen numbers. The probability of her winning the grand prize given that she won a prize is $\frac{m}{n}$ where m and n are relatively prime positive integers. Find $m + n$.

Groundtruth(invisible to model): 116

Assistant: (use python tool to verify its internal thoughts rather than using it as direct solvers)

```

Okay, let's try to solve this problem step by step. First, let's make sure I understand the
problem correctly.
Jen picks 4 distinct numbers from the set S = {1, 2, ..., 10}. Then, 4
numbers are randomly chosen from S (let's call this the winning set). She wins a prize if at least
two of her numbers match the winning set. She wins the grand prize if all four match. We need to
find the probability of her winning the grand prize given that she won a prize. That is, P(grand
prize | won a prize) = P(grand prize) / P(won a prize), since grand prize is a subset of winning a
prize.
First, let's recall that in probability, the conditional probability P(A|B) = P(A ∩ B) / P(B).

.....

Exactly 0 matches: C(2,0)*C(3,2)=1*3=3.
Prize-winning sets (at least 1 match):
6+1=7.
Conditional probability: 1/7. Which makes sense. If we compute it as (1/10)/(7/10)=1/7.
Correct.
So the method is sound.
Back to the original problem: 115 is the number of
prize-winning sets, 1 is grand prize. So 1/115.
Wait, but let's compute C(10,4) to be sure:
10*9*8*7/(4*3*2*1)=210. Correct.
C(4,2)=6, C(6,2)=15, 6*15=90. Correct.
C(4,3)=4, C(6,1)=6,
4*6=24. Correct.
C(4,4)=1, C(6,0)=1. Correct.
90+24+1=115. Correct.
So the conditional
probability is 1/115. Therefore, m=1, n=115, m+n=116.
I think that's it. Let me just write a
Python code to verify the counts.
We can compute:
ntotal = comb(10, 4)
grand = comb(4, 4)
prize = comb(4, 2)*comb(6, 2) + comb(4, 3)*comb(6, 1) + comb(4, 4)*comb(6,
0)
probability = grand / prize
But since we need the fraction, it's grand / prize = 1 /
115.
Let's code it using math.comb (since Python 3.10+ has math.comb, but if not, we can define
it).
Wait, in Python, math.comb is available from 3.8 onwards. Let's assume the environment
supports it.
Code:
import math
Total number of winning sets (not needed for the ratio,
but for verification)
ntotal = math.comb(10, 4) # 210
Number of grand prize sets
grand = math.comb(4, 4) # 1
Number of prize-winning sets (at least 2
matches)
prize_2 = math.comb(4, 2) * math.comb(6, 2) # 6*15=90
prize_3 = math.comb(4, 3) *
math.comb(6, 1) # 4*6=24
prize_4 = grand # 1
prize_total = prize_2 + prize_3 + prize_4 #
90+24+1=115
Conditional probability
prob = grand / prize_total
But we need the fraction
m/n where m and n are coprime
from fractions import Fraction
fraction = Fraction(grand,
prize_total)
nm = fraction.numerator
n = fraction.denominator
print(f"m: {m}, n: {n}, m+n:
{m+n}")
This should output m=1, n=115, m+n=116.
Let me check if math.comb(6,2) is 15:
6*5/2=15, yes. math.comb(4,2)=6, correct.
Yes, this code should confirm the counts.

"tool_calls": [{"arguments": {"code": "\nimport math\n\nfrom fractions import Fraction\n\nntotal
= math.comb(10, 4)\ngrand = math.comb(4, 4) * math.comb(6, 0)\nprize_2 = math.comb(4, 2) *
math.comb(6, 2)\nprize_3 = math.comb(4, 3) * math.comb(6, 1)\nprize_4 = grand\nprize_total =
prize_2 + prize_3 + prize_4\n\nfraction = Fraction(grand, prize_total)\n\nm = fraction.numerator\nn
= fraction.denominator\n\nprint(f\"m: {m}, n: {n}, m+n: {m+n}\")\n"}], "name": "python_tool"},
"type": "function"}]

```

Tool: (tool response)

```

"stdout": "m: 1, n: 115, m+n: 116\n\nstderr:"

```

Figure 12: **Pattern I: Confirmatory Tool Usage.** An example from AIME24 where Qwen3-235B-Thinking solves the problem internally first, arriving at the correct answer. However, it subsequently invokes the Python tool solely to verify this known result. This "Think-then-Verify" behavior unnecessarily inflates the context length and PTE cost without contributing new information to the solution.

Model	Level 1 Gap	Level 3 Gap	Level 5 Gap
GPT-OSS	+5.1%	+158.2%	+371.1%
GLM-4.5-Air	+419.9%	+82.7%	+259.9%
Qwen3-32B	+793.2%	+69.6%	+100.2%
Qwen2.5-72B	+43.0%	+32.3%	+16.9%
Dv3.1-terminus	+341.8%	-0.2%	+46.3%

Table 8: **Percentage PTE increase of incorrect trajectories over correct ones within the same difficulty level.**

Trajectory Example of DeepSeek-V3.1-Terminus on Webinstruct-Verified

Question: Organic agriculture implies farming techniques that use no ____.

Groundtruth(invisible to model): Insecticides. Inorganic or synthetic fertilizers. Herbicides.

Assistant: (use search tool)

I'll help you understand what organic agriculture excludes by searching for information about organic farming techniques.

```
"tool_calls": {"arguments": {"query": ["organic agriculture definition what is not used", "organic farming prohibited substances techniques", "what does organic agriculture exclude"]}, "name": "google_search_tool"}
```

Tool: (tool response)

"A Google search for 'organic agriculture definition what is not used' found 10 results:\n\n# Web Results

- [5 Reasons Why Organic Farming Is Not Sustainable] (<https://www.edengreen.com/blog-collection/organic-farming-sustainability>)
Organic produce is grown without synthetic pesticides, fertilizers, or GMOs. It must meet strict USDA certification standards. These standards include using ...See more
- [What Organic Farming (and Processing) Doesn't Allow] (<https://www.usda.gov/about-usda/news/blog/organic-101-what-organic-farming-and-processing...>)
Methods like irradiation, sewage sludge, and genetic engineering are all expressly prohibited from being used when growing or processing organic foods.See more
- [Organic vs. Conventional Farming]..."

Assistant: (use visit tool)

Now let me visit the most relevant USDA page to get detailed information about what organic agriculture excludes:

```
"tool_calls": {"arguments": {"url": "https://www.usda.gov/about-usda/news/blog/organic-101-what-organic-farming-and-processing...", "goal": "Find specific information about what substances and methods are prohibited in organic agriculture"}, "name": "visit_tool"}
```

Tool: (tool response)

The useful information in <https://www.usda.gov/about-usda/news/blog/organic-101-what-organic-farming-and-processing...> for user goal Find specific information about what substances and methods are prohibited in organic agriculture as follows: \n\nEvidence in page: \nA very important part of the process-based regulatory framework is the prohibition of certain methods in organic production and handling...

Assistant:

.....

Tool:

.....

Assistant:(use python tool)

Based on my research, I can now provide a comprehensive answer about what organic agriculture excludes. Let me use Python to verify and organize the key points:

```
"tool_calls": {"arguments": {"code": "\n# Organizing the key prohibited substances and methods in organic agriculture\n\nprohibited_items = [\n    \"synthetic fertilizers\", \n    \"synthetic pesticides\", \n    \"genetically modified organisms (GMOs)\", \n    \"irradiation\", \n    \"sewage sludge\", \n    \"artificial colors\", \n    \"artificial preservatives\", \n    \"artificial flavors\", \n    \"genetic engineering methods (cell fusion, recombinant DNA, etc.)\"]\n\nprint(\"Organic agriculture excludes the use of:\")\nfor i, item in enumerate(prohibited_items, 1):\n    print(f\"{i}. {item}\")\n\nprint(\"The basic principle is: 'allow natural substances and prohibit synthetic substances'\")\nprint(\"Source: USDA National Organic Program standards and EPA guidelines\")"}}, "name": "python_tool"}
```

Tool: (tool response)

```
"stdout": "\nOrganic agriculture excludes the use of:\n1. synthetic fertilizers\n2. synthetic pesticides\n3. genetically modified organisms (GMOs)\n4. irradiation\n5. sewage sludge\n6. artificial colors\n7. artificial preservatives\n8. artificial flavors\n9. genetic engineering methods (cell fusion, recombinant DNA, etc.)\n\nThe basic principle is: 'allow natural substances and prohibit synthetic substances'\n\nSource: USDA National Organic Program standards and EPA guidelines\n\nstderr:"
```

Assistant:(final answer ✓)

```
<ANSWER>Organic agriculture implies farming techniques that use no synthetic fertilizers, synthetic pesticides, genetically modified organisms (GMOs), irradiation, sewage sludge, artificial colors, artificial preservatives, artificial flavors, or genetic engineering methods. The fundamental principle is to allow natural substances while prohibiting synthetic ones.</ANSWER>
```

Figure 13: **Pattern II: Tool-Mixing.** An example from WebInstruct-Verified where DeepSeek-V3.1-Terminus fragmentally alternates between Search and Python tools. This behavior accumulates context of intermediate outputs and inflates the PTE cost, yet yields no obvious accuracy improvement compared to single-toolset strategies (as evidenced in Fig. 4e).

Trajectory Example of Qwen-2.5-7B-Instruct on AIME24

Question: Find the largest possible real part of $[(75 + 117i)z + \frac{96+144i}{z}]$ where z is a complex number with $|z| = 4$.

Groundtruth(invisible to model): 540

Assistant: (python tool invoked, but no print function used)

To find the largest possible real part of the expression $((75+117i)z + \frac{96+144i}{z})$ where (z) is a complex number with $(|z|=4)$, we can start by expressing (z) in its polar form. Since $(|z|=4)$, we can write $(z = 4e^{i\theta})$ for some angle (θ) . Let's denote $(z = 4(\cos\theta + i\sin\theta))$. Then we have:
 $(75+117i)z = (75+117i)4(\cos\theta + i\sin\theta) = 4(75\cos\theta - 117\sin\theta + i(75\sin\theta + 117\cos\theta))$
 $\frac{96+144i}{z} = \frac{96+144i}{4(\cos\theta + i\sin\theta)} = \frac{1}{4}(96\cos\theta - 144\sin\theta + i(96\sin\theta + 144\cos\theta))$
Adding these two expressions together, we get:
 $(75+117i)z + \frac{96+144i}{z} = 4(75\cos\theta - 117\sin\theta + i(75\sin\theta + 117\cos\theta)) + \frac{1}{4}(96\cos\theta - 144\sin\theta + i(96\sin\theta + 144\cos\theta))$
The real part of this expression is:
 $4(75\cos\theta - 117\sin\theta) + \frac{1}{4}(96\cos\theta - 144\sin\theta) = 300\cos\theta - 468\sin\theta + 24\cos\theta - 36\sin\theta = 324\cos\theta - 504\sin\theta$
We need to maximize the expression $(324\cos\theta - 504\sin\theta)$. This is a linear combination of sine and cosine, and its maximum value is the amplitude of the function, which is $(\sqrt{324^2 + 504^2})$. Let's calculate this amplitude:
 $(\sqrt{324^2 + 504^2}) = \sqrt{104976 + 254016} = \sqrt{358992} = 36\sqrt{279}$
So the maximum value of the real part is $(36\sqrt{279})$. We can simplify $(\sqrt{279})$ as $(\sqrt{9 \cdot 31}) = 3\sqrt{31}$, so the maximum value is:
 $(36 \cdot 3 \cdot \sqrt{31}) = 108\sqrt{31}$. Let's verify this using Python code.

```
"tool_calls": {"arguments": "{\"code\": \"import math\n# Calculate the maximum value\nmax_value = 36 * 3 * math.sqrt(31)\nmax_value\""}, "name": "python_tool"}
```

Tool: (tool response)

```
"stdout:\n\nstderr:"
```

Assistant:(final answer~~X~~)

The maximum value of the real part of the expression $((75+117i)z + \frac{96+144i}{z})$ where (z) is a complex number with $(|z|=4)$ is $(108\sqrt{31})$. Thus, the final answer is: $(108\sqrt{31})$

Figure 14: **Pattern III: Lack of Tool Priors.** An example from AIME24 where Qwen-2.5-7B-Instruct fails to utilize the Python tool effectively. The model invokes the code interpreter but forgets to include a print statement, resulting in an empty output. This suggests a lack of pretraining exposure to the tool environment, leading to wasted inference steps.

Dataset	L5 Correct vs L1 Wrong	L5 Correct vs L2 Wrong	L5 Correct vs L3 Wrong	L5 Correct vs L4 Wrong
MATH500	57.6% (n=109)	38.4% (n=290)	57.7% (n=475)	68.4% (n=605)
AIME25	N/A	N/A	N/A	34.4% (n=126)

AIME25 problems were only annotated as Level 4–5 by GPT-4o, despite the 1–5 scale.

Table 9: **Percentage of Level-5 correct trajectories with lower PTE than average failures at easier levels.**

Metric	Ratio (In:Out)	Pearson r	Notes
Naive Token Count	1:1	0.625	Ignores context accumulation
Pricing (DeepSeek-V3.2)	1:1.5	0.681	Fixed I/O price ratio
Pricing (Standard)	1:3	0.734	Fixed I/O price ratio
Pricing (GPT-4o/Claude-3.5)	1:4	0.758	Fixed I/O price ratio
PTE (Ours)	Hardware-Adaptive	0.925	Captures physical bandwidth bottlenecks

Table 10: **Correlation with wall-clock latency: PTE vs. token-count and commercial pricing baselines.**

Trajectory Example of Tongyi-Deepresearch on SimpleQA

Question: What is the British-American kickboxer Andrew Tate's kickboxing name?

Groundtruth(invisible to model): King cobra

Tool Config:(Notice that we define `google_search_tool` and `visit_tool`)

```
"tools":[{"type": "function", "function": {"name": "google_search_tool", "description": "Performs
batched web searches: supply a list of queries; the tool retrieves the top 10 results for each
query in one call, each including a title, URL, and abstract. Use these abstracts to evaluate the
relevance of each result and determine which URL to visit for more detailed information based on
the original query.", "parameters": {"type": "object", "properties": {"query": {"type": "array",
"items": {"type": "string"}, "description": "Array of query strings. Include multiple
complementary search queries in a single call."}}, "required": ["query"]}},

{"type": "function", "function": {"name": "visit_tool", "description": "Visit a single webpage URL
and return a concise summary of the content tailored to the goal provided. The goal should be
directly derived from the original question to ensure the summary extracts only the most relevant
information. The summary should include key facts, dates, names, or events as applicable.",
"parameters": {"type": "object", "properties": {"url": {"type": "string", "description": "The URL
of the webpage to visit."}, "goal": {"type": "string", "description": "The goal of the visit for
the webpage."}}, "required": ["url", "goal"]}}}]
```

Assistant: (use `search` tool, which is not defined in our framework but provided in the Tongyi opensource framework)

```
The user asks: "What is the British-American kickboxer Andrew Tate's kickboxing name?" They
refer to Andrew Tate as a British-American kickboxer. His kickboxing name is "User83"? Or
something? Let's recall. Andrew Tate, also known as "Croatian" or "Andrew Tate" is known as
"Cairo"? I'm not sure. Let's recall specifics: Andrew Tate, born 1 December 1986, is a
British-American internet personality, businessman, and former professional kickboxer. He has a
nickname in kickboxing: "Croatian Wolf"? Actually a known nickname is "Crock"? Let's
confirm. I recall various references: Andrew Tate's actual nickname in kickboxing is "Crock"?
I might be mixing with other things. Let's recall known facts: Andrew Tate's kickboxing
nickname is "Crockett"? I've seen his social media profile uses the handle "Crock"? Let's
search."
```

```
"tool_calls": {"arguments": [{"query": ["Andrew Tate kickboxing name nickname", "Andrew Tate
kickboxing moniker", "What is Andrew Tate's kickboxing name?", "Andrew Tate fight name",
"Andrew Tate kickboxing record nickname"]}], "name": "search", "type": "function"}}
```

Tool: (tool response)

```
"Error: tool 'search' not registered."
```

Assistant: (use `search` tool again, fail to be parsed by engine due to extra linefeed, final answer~~X~~)

```
"<think>\n\nThe tool name is "search" under the "functions" namespace. Usage: we need to provide
a JSON with key "name": "search", "arguments": ...? The conversation format for calling the
tool is as per the instructions: we need to output:\n\n<tool_call>\n{"name": "search",
"arguments": {"query": [...]}}\n\nMake sure the name matches the function
defined: search. So do that.\n\n</think>\n\n<tool_call>\n{"name": "search", "arguments":
{"query": ["Andrew Tate kickboxing name nickname", "Andrew Tate kickboxing moniker", "What
is Andrew Tate's kickboxing name?", "Andrew Tate fight name", "Andrew Tate kickboxing record
nickname"]}}\n\n</tool_call>"
```

Figure 15: **Pattern IV: Tool Format Collapse.** An example from SimpleQA illustrating model brittleness. Tongyi-Deepresearch fails to adhere to the predefined tool schema, hallucinating a search tool format that differs syntactically from the system prompt (likely reverting to its training data format). This mismatch causes immediate parsing error and recurring failures in the following steps despite the semantic intent being correct.