

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 BOOSTSTEP: BOOSTING MATHEMATICAL CAPABILITY OF LARGE LANGUAGE MODELS VIA STEP-ALIGNED IN CONTEXT LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Large language models (LLMs) have demonstrated impressive ability in solving complex mathematical problems with multi-step reasoning and can be further enhanced with well-designed in-context learning (ICL) examples. However, this potential is often constrained by two major challenges in ICL: granularity mismatch and irrelevant information. We observe that while LLMs excel at decomposing mathematical problems, they often struggle with reasoning errors in fine-grained steps. Moreover, ICL examples retrieved at the question level may omit critical steps or even mislead the model with irrelevant details. To address this issue, we propose BoostStep, a method that enhances reasoning accuracy through step-aligned ICL, a novel mechanism that carefully aligns retrieved reference steps with the corresponding reasoning steps. Additionally, BoostStep incorporates an effective "first-try" strategy to retrieve for exemplars highly relevant to the current state of reasoning. BoostStep is a flexible and powerful method that integrates seamlessly with chain-of-thought (CoT) and tree search algorithms, refining both candidate selection and decision-making. Empirical results show that BoostStep improves GPT-4o's CoT performance by 4.6% across mathematical benchmarks, significantly surpassing traditional few-shot learning's 1.2%. Moreover, it can achieve an additional 7.5% gain combined with tree search. Surprisingly, it enhances state-of-the-art LLMs to solve challenging math problems using simpler examples. It improves DeepSeek-R1-671B and Qwen3-235B's performance on American Invitational Mathematics Examination (AIME) by 2.2% and 5.0% respectively, leveraging simple examples only from the MATH dataset.

1 INTRODUCTION

Mathematical reasoning is a crucial and challenging task in the development of artificial intelligence. It serves as an indicator of a model's ability to perform complex reasoning and has a wide range of applications, such as problem-solving, theorem proving, and scientific discovery.

When solving complex mathematical problems, cutting-edge LLMs often adopt a multi-step reasoning strategy. Specifically, they first decompose a complex problem into several simpler steps and then solve each single step independently.

Through the analysis of error cases, we found that current SOTA models are relatively correct in the step-dividing phase, that is, the model can know exactly what tasks should be completed in each step. However, there are still a lot of mistakes within each reasoning step, such as wrong formula use, wrong calculation, insufficient enumeration, etc. To quantitatively substantiate this observation, we provide GPT-4o-mini with a ground truth reasoning process to determine whether the error in another response was due to an overarching flawed reasoning approach or a deviation within a particular step. In less advanced models like LLaMA-3.1-8B (Dubey et al., 2024), 91.3% of errors originate from single-step reasoning. In more advanced models like GPT-4o, up to 99.2% of errors are ascribable to some particular steps. This exaggerated proportion suggests that the correctness of single-step reasoning is the bottleneck of reasoning capability.

Various approaches have been employed to improve reasoning correctness, such as producing chains of thought through prompt engineering (Kojima et al., 2022; Wei et al., 2022), fine-tuning with

054 mathematical data (Shao et al., 2024; Yang et al., 2024; Ying et al., 2024), or generating multiple
 055 candidate reasoning paths using Tree Search Methods (Zhang et al., 2024b;a; Wang et al., 2024b).
 056

057 Among those techniques, in-context learning is a particularly important one, which offers similar ex-
 058 amples to provide detailed guidance. However, the examples retrieved by traditional problem-level
 059 in-context learning are listed before the reasoning process, thereby lacking fine-grained guidance
 060 during the reasoning process. Moreover, since the example problem can't be identical to the new
 061 one, the irrelevant steps in those examples may even become a distraction from the current reason-
 062 ing, thus even negatively affecting the single-step reasoning capability for some specific steps.
 063

064 To this end, we refine in-context learning from problem-level to step-level granularity to offer similar
 065 example steps during an ongoing reasoning process for fine-grained step-aligned guidance. We also
 066 ensure that the introduced example is still relevant at the step level to avoid distractions.
 067

068 Firstly, we have constructed an example problem bank with step-level granularity based on reasoning
 069 content instead of commonly adopted grammatical separation. This ensures the steps in the problem
 070 bank are consistent with the actual reasoning steps, thereby providing more appropriate guidance.
 071

072 Building on the step-level granularity within the example problem bank, we propose an approach
 073 that incorporates in-context learning through a "first-try" format during an on-going reasoning pro-
 074 cess. For a given problem to be solved, we first break down the solving process into step-by-step
 075 reasoning paths. During the reasoning of a single step, we first allow the model to attempt a 'first
 076 try' to comprehend what the model currently needs to reason about. Based on this initial attempt, we
 077 then search the problem bank to find similar steps that can guide the model to accurately output the
 078 current step. This helps ensure a higher similarity between the retrieved examples and the current
 079 step so the distraction from irrelevant steps can be avoided and the guidance effect can be improved.
 080

081 Compared with traditional problem-level ICL, our method provides examples during the reasoning
 082 process directly based on the steps to be solved, thereby offering more relevant guidance. It demon-
 083 strates significant improvements over traditional few-shot learning across various benchmarks, with
 084 an average increase of 3.4% on GPT-4o.
 085

086 Moreover, our method also reduces the sensitivity to the similarity between the example and the
 087 target problem, as two different problems can still share similar steps. Consequently, dissimilar
 088 problems can still offer effective guidance. On multi-modal benchmarks with lower similarity to
 089 example problems, traditional few-shot learning has a detrimental effect, resulting in an accuracy
 090 reduction of 0.9% on GPT-4o. In contrast, our approach still achieves an improvement of 2.8%.

091 Besides, BoostStep also shows a promising potential to improve the reasoning quality on harder
 092 problems with simpler examples. With examples from MATH (Hendrycks et al., 2021), it helps
 093 Deepseek-R1 and Qwen3-235B-Instruct-2507 (Yang et al., 2025) achieve an improvement of 2.2%
 094 and 5.0% respectively on the much more challenging American Invitational Mathematics Examina-
 095 tion (AIME) problems.
 096

097 Moreover, our method is also highly compatible with various current reasoning strategies that
 098 employ step-level tree search. Typically, a tree-search method requires a reason model to generate
 099 multiple step-level candidate reasoning paths and a critic model to evaluate the correctness of these
 100 candidates. Our approach can be integrated into both aspects. Specifically, when the reason model
 101 generates new candidate reasoning nodes, our method can introduce similar examples in the afore-
 102 mentioned 'first-try' manner to improve the accuracy of candidates. Additionally, it can aid the critic
 103 model by incorporating similar example steps into the evaluation of candidate reasoning processes
 104 to provide similar guidance. Experiments indicate that both applications contribute positively and
 105 bring about an improvement of 8.5% jointly on GPT-4o.
 106

107 2 RELATED WORKS

108 **Mathematical Reasoning.** Mathematical reasoning has long been a challenging task in artificial
 109 intelligence. Early methods (Feigenbaum et al., 1963; Fletcher, 1985) attempted to perform sim-
 110 ple mathematical reasoning through rule-based methods. With the advent of large language models
 111 with enhanced reasoning capabilities, contemporary approaches typically focus on enhancing per-
 112 formance during both the training and inference phases. The first category improves mathematical
 113 capability by fine-tuning with more high-quality mathematical data (Shao et al., 2024; Yang et al.,
 114

108 2024; Lewkowycz et al., 2022; Yue et al., 2023; Xu et al., 2024). However, it demands substantial
 109 high-quality mathematical data and computational resources. Consequently, more efforts have been
 110 put into exploring various techniques during inference to enhance mathematical reasoning perfor-
 111 mance. Some work (Wei et al., 2022; Kojima et al., 2022) involves prompt engineering to enable
 112 models to generate comprehensive chains of thought. Others (Madaan et al., 2024; Gou et al., 2023;
 113 Ke et al., 2024) use self-refinement techniques to revise the initial reasoning outputs.

114 **Step-level Mathematical Reasoning.** Recently, many studies have shifted the granularity of math-
 115 ematical reasoning from the problem level to the step level. This approach involves addressing each
 116 next step individually and completing small segments of reasoning within the overall task. These
 117 works often employ tree searching strategies like Tree of Thoughts (ToT) (Yao et al., 2024; Besta
 118 et al., 2024) or Monte Carlo Tree Search (Zhang et al., 2024b;a; Chen et al., 2024; Feng et al., 2023;
 119 Zhu et al., 2022), extending multiple steps to optimize step answers and ultimately obtain the op-
 120 timal solution. Additionally, Process-Supervised Models (Lightman et al., 2023; Luo et al., 2024)
 121 are frequently used to verify the correctness of new candidate steps in real-time and prune reasoning
 122 paths. This more detailed auxiliary strategy demonstrates greater potential.

123 **ICL in Mathematical Reasoning.** In-context learning can provide low-cost guidance through
 124 similar examples. However, research on in-context learning within mathematical reasoning tasks re-
 125 mains insufficient. Typically, this approach involves providing the model with similar problems and
 126 their ground truth solutions to offer a general strategy for solving new problems (Hendrycks et al.,
 127 2021; Wei et al., 2022). Some efforts have been made to improve the relevance of retrieved examples
 128 by designing better retrieval mechanisms (Liu et al., 2024b). Others try to provide high-level con-
 129 text instead to improve the generalizability (Wu et al., 2024). Some recent approaches (Dong et al.,
 130 2024) introduce ICL into an on-going reasoning. However, all these methods share a common lim-
 131 itation: the lack of fine-grained step-level guidance. They still perform ICL in problem granularity
 132 and thus may not offer effective guidance for single-step reasoning.

3 STEP-LEVEL IN-CONTEXT LEARNING

3.1 REVISITING IN-CONTEXT LEARNING FROM CONDITIONAL PROBABILITY

138 Current models often employ next-token prediction for training and inference, where the conditional
 139 probability is central to the model’s generation of the next token. Given a problem q , a model’s
 140 reasoning process can be represented by $r_{predict} = \arg \max_r P_{model}(r | q)$, where we train the
 141 model to get a better conditional probability P_{model} so that $r_{predict}$ can be closer to the ground truth
 142 answer $r_{gt} = \arg \max_r P_{gt}(r | q)$.

144 In-context learning provides the model with conditional probabilities similar to the ground truth
 145 answer for imitation without changing the probability model P_{model} . Specifically, an example
 146 problem q' and its corresponding correct solution r' is provided and it can be posited that the
 147 conditional probability $P(r' | q')$ is similar to the probability of the ground truth answer of
 148 the target problem $P(r_{gt} | q)$. Consequently, the model will imitate this similar example and
 149 $r'_{predict} = \arg \max_r P_{model}(r | q, q', r')$ will be closer to r_{gt} comparing to $r_{predict}$.

150 However, given that the actual reasoning process r can be highly complex, the complete reasoning
 151 process is often divided into multiple steps s_1, s_2, \dots . Step-level reasoning iteratively guides the
 152 model to generate the next step $s_{i+1}^{0-shot} = \arg \max_s P_{model}(s | q, s_1, s_2, \dots, s_i)$.

153 At the step granularity, examples retrieved based on the problem q are evidently insufficient for
 154 providing appropriate guidance. Similar problem q' may not necessarily contain the correspond-
 155 ing steps to guide the reasoning for the new problem q . Moreover, irrelevant steps may provide
 156 dissimilar conditional probabilities, thereby distracting the model’s reasoning process.

158 To this end, we propose step-aligned in-context learning and a first-try strategy to provide detailed
 159 and relevant example steps when in step-level reasoning. Specifically, when generating new steps
 160 s_{i+1} based on previous reasoning steps s_i, s_{i-1}, \dots, s_1 and question q , we first utilize a first-try
 161 strategy to obtain an approximate estimate of s_{i+1}^{first} . Then, we use this s_{i+1}^{first} to retrieve a similar step
 s'_{i+1} along with the corresponding $q', s'_1, s'_2, \dots, s'_n$. Since these two steps are similar, a very rea-

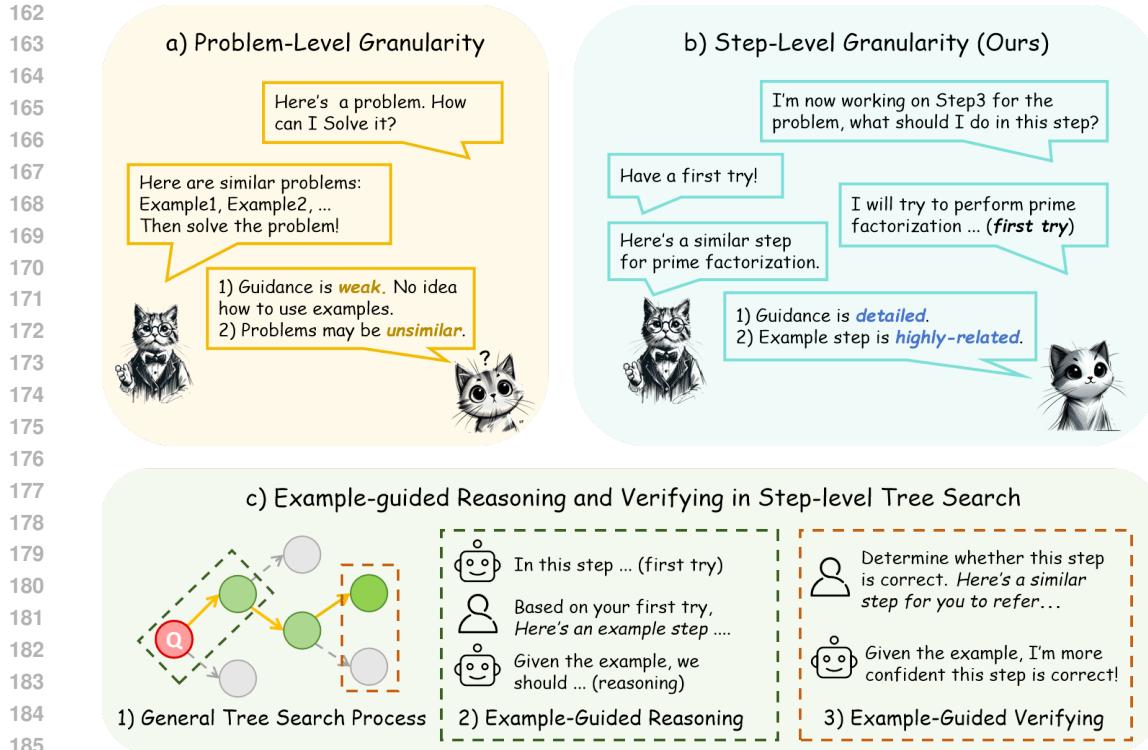


Figure 1: Our strategy refines in-context learning from problem-level granularity (fig.a) to step-level granularity (fig.b) to provide more real-time fine-grained guidance. Moreover, our strategy can guide the reasoning and verifying process in tree-searching strategies by introducing examples.

sonable assumption is that $P(s'_{n+1} | q', s'_1, \dots, s'_n)$ closely approximates $P(s_{gt_{i+1}} | q, s_1, \dots, s_i)$. Therefore, the generated step $s_{i+1} = \arg \max_s P_{model}(s | q, s_1, \dots, s_i, q', s'_1, \dots, s'_n, s'_{n+1})$ will be more closed to $s_{gt_{i+1}}$ comparing to s_{i+1}^{0-shot} . Details about our step-level in-context learning and first-try strategy will be explained in Sec. 3.3

3.2 STEP-LEVEL EXAMPLE PROBLEM BANK

Current open-source mathematical data no longer consist solely of problems and their final answers to determine whether the final answer obtained is correct or not. Instead, they also provide detailed solution processes to provide more fine-grained measurements. However, most current open-source mathematical data still do not break down the solution processes to the step level.

A major advantage of decomposing the question example bank into individual steps is that it facilitates step-level retrieval and guidance, which is of significant importance. As illustrated in Fig. 2, two distinctly different problems may contain similar key steps. Traditional problem-level in-context learning often overlooks such examples, whereas step-level in-context learning can effectively recall these steps, thereby providing fine-grained guidance to the ongoing reasoning process.

How to derive different steps from a complete solution is of great importance. Some approaches (Lightman et al., 2023) proposed using a clear semantic delimiter like the period '.' or a new line to segment steps. This allows for the quick decomposition without any additional assistance. However, this simple decomposition mode is obviously unreliable. A single reasoning step should have a consistent target and a complete thought process, making it the atomic granularity of reasoning. Using semantic delimiter may disrupt this atomicity. For example, it may split a complete enumeration for the same objective into multiple steps.

Therefore, we suggest that the most appropriate method for step segmentation is to allow the reason model itself to autonomously decompose the process. This approach ensures that the granularity of

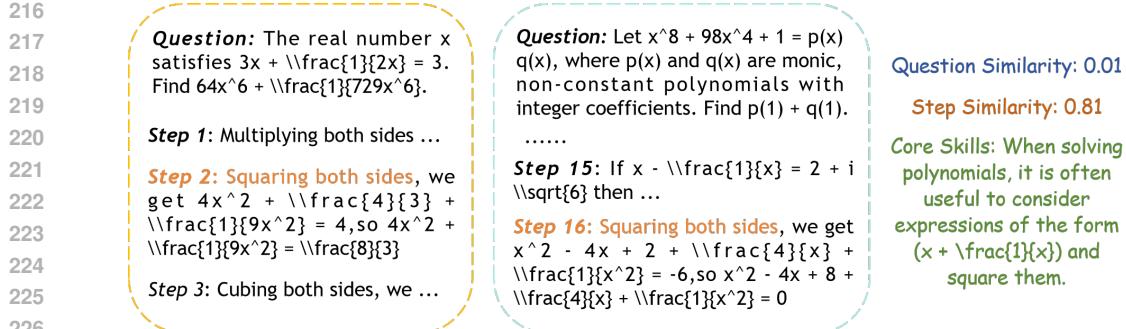


Figure 2: Different problems may contain similar steps. Problem-level in-context learning will ignore this example due to low problem similarity. In contrast, our step-level in-context learning strategy can introduce the core skills by step-level retrieval and guidance.

the decomposed steps in example problem bank aligns with that of the real-time reasoning steps. Specifically, we define the concept of a step through prompts, which encapsulate a complete and simple inference. This guides GPT-4o in decomposing the answer at the step level.

We demonstrate some specific examples of different step-dividing strategies in Sec. 4.2.

3.3 STEP-LEVEL ICL WITH FIRST-TRY STRATEGY

The core challenge of in-context learning lies in how to effectively retrieve relevant problems or steps for effective guidance. This is contingent upon both the similarity between the problem database and the target problem, as well as the retrieval strategy employed. Traditional problem-level in-context learning involves retrieving similar problems based solely on the problem statement, which is straightforward but effective, as similar problems typically encompass similar reasoning processes.

At the more granular step level, however, the situation becomes much more complex. A simple strategy is to perform retrieval using the given problem and all preceding reasoning steps $s_{i-1}, s_{i-2}, \dots, s_1, q$. The clear drawback of this method is the excessive length of the retrieval content, which diminishes the emphasis on the uniqueness of the current step. Another strategy is to use the previous step s_{i-1} to retrieve s'_{i-1} from a step-level database, thereby guiding the reasoning of s_i through the correct resolution of s'_{i-1} . However, this approach is rather crude, as it models step-level reasoning as a Markov process, which is evidently unreasonable. Similar steps can be applicable to different reasoning tasks, and therefore similarity in the previous step does not necessarily indicate that the retrieved subsequent step will provide valuable guidance for the reasoning in the current step.

To this end, we propose a straightforward and effective "first-try" strategy to enhance the similarity of search steps. Our premise is that the most accurate way to estimate the next step is to actually allow the model to attempt the reasoning for the next step. Specifically, given a problem q and all preceding reasoning steps $s_{i-1}, s_{i-2}, \dots, s_1$, we first instruct the model to attempt continuing the reasoning process to arrive at a tentative step s_i^{try} without the aid of any examples. Subsequently, we use s_i^{try} to retrieve similar steps s'_j along with their corresponding problem q' and preceding steps s'_1, \dots, s'_{j-1} from a step-level database. Finally, we feed the retrieved similar steps back to the model, enabling it to deduce the final step s_i . Besides, we add a widely accepted strategy reference rejection. Specifically, if the similarity of the retrieved most similar example remains below a certain threshold, we consider that there are no sufficiently similar examples available for reference and we do not provide any examples to avoid the negative effects associated with incoherent in-context learning. This "try-retrieve-reason" strategy significantly enhances retrieval relevance, thereby improving reasoning effectiveness. Experiments in Sec. 4.3 compare our method with several other retrieval strategies, demonstrating our superiority.

3.4 STEP-LEVEL GUIDANCE IN TREE SEARCH

Our step-level in-context learning can significantly enhance the model's single-step reasoning capability, which makes it easily integrated into common step-level tree-search strategies.

270
 271 **Table 1: BoostStep generalizes across models and benchmarks.** Comparison of different ICL
 272 strategies on different benchmarks on GPT-4o, Qwen2.5-Math-72B-Instruct and Qwen3-32B.

Model	Method	MATH	AMC12	AMC10	AQUA	MathBench(C)	MathBench(H)	Olympiad	Avg
GPT-4o	0-shot	73.4	53.6	55.8	81.1	80.0	77.3	40.6	66.0
	few-shot	73.8	56.5	56.7	83.9	80.7	79.3	39.3	67.2 (+1.2)
	Ours	76.4	63.0	60.4	85.4	82.0	84.0	43.3	70.6 (+4.6)
Qwen2.5 Math-72B	0-shot	83.0	67.4	67.7	84.6	80.6	82.0	49.7	73.6
	few-shot	83.8	67.4	66.8	85.0	81.3	82.7	49.9	73.8 (+0.2)
	Ours	85.2	69.2	69.6	86.6	82.7	84.7	52.7	75.8 (+2.2)
Qwen3-32B	0-shot	85.4	66.6	66.8	83.9	83.3	84.7	56.6	75.3
	few-shot	86.0	64.5	66.3	85.1	86.7	80.0	53.9	74.6 (-0.7)
	Ours	87.6	68.9	69.1	85.1	90.7	87.3	57.0	78.0 (+2.7)

281
 282 Generally, tree search methods necessitate two key components: a reason model that generates
 283 step-level reasoning and a Process-Supervised Reward Model (PRM) that continuously evaluates
 284 the current reasoning step in real time. Our method is beneficial for both of these components. It
 285 enhances the step-level reasoning performed by the reason model and improves the effectiveness of
 286 the PRM in evaluating current reasoning steps.

287 For the reason model, tree search methods inherently require step-by-step reasoning expansion.
 288 When expanding at node s_i , we can apply the previously mentioned strategy: the model performs
 289 n first tries and retrieve for n example steps. For each example, the model then completes the
 290 reasoning to generate n child nodes $s_{i+1}^1, \dots, s_{i+1}^n$ with the help of these examples. Similarly, our
 291 strategy can improve the accuracy of individual nodes s_{i+1}^j .

292 Evidently, judgment ability is closely related to reasoning ability. Therefore, since our strategy can
 293 enhance the accuracy of single-step reasoning, a reasonable assumption is that introducing appropriate
 294 example steps can improve the PRM’s ability to assess the correctness of the current reasoning
 295 process. In particular, when evaluating the correctness of an inference step candidate s_i^j , we re-
 296 trieve similar steps s_k' along with their corresponding preceding steps s_{k-1}', \dots, s_1' and question q'
 297 from the step-level example bank. Similarly, the probability distributions $P(s_k'|s_{k-1}', \dots, s_1', q')$
 298 and $P(s_{gt_i}|s_{i-1}, \dots, s_1, q)$ exhibit similarities. This resemblance aids in assessing the discrepancy
 299 between s_i^j and s_{gt_i} , thereby enhancing the accuracy of the critic model’s evaluations.

300 Detailed ablation experiments in Sec. 4.4 demonstrate that both strategies contribute positively.

305 4 EXPERIMENTS

308 4.1 EXPERIMENT SETTING

310 **Reasoning Model.** We conducted experiments on different reasoning models including GPT-
 311 4o (Hurst et al., 2024), which is our primary model, Qwen2.5-Math-72B-Instruct (Yang et al.,
 312 2024), Qwen3-32B (Yang et al., 2025) and SOTA rasoning models Qwen-QwQ-32B (Team, 2024),
 313 DeepSeek-R1-671B (Guo et al., 2025) and Qwen3-235B-A22B-Instruct-2507 (Yang et al., 2025).

314 **Evaluation Benchmark.** We tested our approach on several challenging open-source mathematical
 315 benchmarks. More details are listed in the appendix.

317 **Example Problem Bank.** The problems and the solutions are obtained from PRM800K (Lightman
 318 et al., 2023). Then we use our step-dividing strategy discussed above to divide example steps.

319 **Retriever.** We utilized the TF-IDF strategy as the retriever. The TF-IDF weight matrix is derived
 320 from the example problem bank because the impact of the newly generated step is negligible.

322 **Hyper-Parameters.** The temperature value is 0 in all the experiments except 0.3 at step-level tree
 323 search. The reference rejection threshold is 0.7. The shot number for traditional ICL is 4.

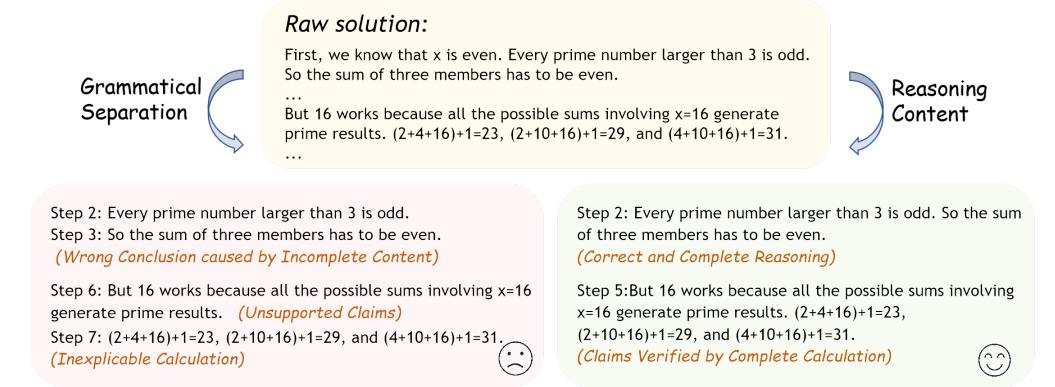
323 **Prompt.** The specific prompts are listed in the appendix.

324
 325 **Table 2: BoostStep enables “simple-aids-complex”.**
 326 Simpler examples from PRM800K can guide LLMs on
 327 much more challenging AIME.

Model	Method	AIME23	AIME24
QwQ-32B	0-shot	38.9	43.3
	few-shot	33.3 (-5.6)	38.9 (-4.4)
	Ours	41.1 (+2.2)	47.8 (+4.5)
DeepSeek-R1	0-shot	75.6	80.0
	few-shot	65.6 (-10.0)	70.0 (-10.0)
	Ours	77.8 (+2.2)	82.2 (+2.2)
Qwen3-235B Instruct-2507	0-shot	70.0	70.0
	few-shot	66.7 (-3.3)	66.7 (-6.6)
	Ours	73.3 (+3.3)	76.7 (+6.7)

328
 329 **Table 3: BoostStep generalizes across**
 330 **modality.** Plain-text examples from
 331 PRM800K can provide effective guid-
 332 ance for MLLMs when solving multi-
 333 modal mathematical problems from
 334 MathVision and MathVerse, while tra-
 335 ditional few-shot learning may even
 336 have negative impact.

Method	MathVision	MathVerse
0-shot	30.6	53.2
few-shot	28.7 (-1.9)	53.2 (0.0)
Ours	35.2 (+4.6)	54.2 (+1.0)



351 Figure 3: Our step division can provide complete and clear example steps as we divide the steps
 352 based on the reasoning content. By contrast, previous grammar-based dividing strategy may result
 353 in meaningless or even incorrect steps.

4.2 MAIN RESULTS

355 **BoostStep generalizes across models.** We evaluate BoostStep as a model-agnostic strategy across
 356 both general-purpose and math-specialized LLMs. As shown in Tab. 1, BoostStep brings consistent
 357 improvements to both GPT-4o and Qwen-Math, demonstrating its broad applicability. The per-
 358 formance gains across all benchmarks confirm that BoostStep does not rely on model-specific tuning
 359 or architectural assumptions, but instead offers universally beneficial step-level guidance.

360 **BoostStep generalizes to out-of-bank benchmarks.** Notably, the example bank used by BoostStep
 361 is constructed solely from the PRM800K training set (i.e., MATH500), yet the method achieves
 362 strong performance across a wide range of math benchmarks. This out-of-distribution generaliza-
 363 tion highlights the robustness of the step-level design: even when benchmark distributions differ,
 364 BoostStep increases the likelihood of identifying transferable intermediate steps, enabling effective
 365 reasoning in unfamiliar problem domains.

366 **BoostStep enables “simple-aids-complex”.** As shown in Tab. 1 and Tab. 2, although the exam-
 367 ple bank is constructed from PRM800K—an easier dataset than the evaluated benchmarks—the
 368 step-level design still provides valuable guidance for solving specific steps within complex math
 369 problems. This demonstrates that even a relatively simple bank can effectively enhance model per-
 370 formance on more challenging tasks by offering transferable step-wise insights.

371 **SOTA reasoning LLMs also benefit from BoostStep.** While SOTA reasoning LLMs already ex-
 372 hibit strong performance on complex mathematical problems, BoostStep can still bring gains by
 373 providing accurate step-level guidance, as evidenced in Tab. 2. In contrast, traditional few-shot
 374 learning often fails to deliver effective support, highlighting the unique advantages of BoostStep’s
 375 structured, step-wise assistance even for advanced models.

376 **BoostStep generalizes across modality.** Despite differences in input modality, multi-modal mathe-
 377 matical reasoning still follows a step-by-step logical process. BoostStep leverages this shared

378
379
380

Table 4: Comparing Booststep with some step-level reasoning strategies without In-context learning.

Method	MATH	AMC12	AMC10	Avg
ToT	77.8	58.7	59.0	65.17
Self-Refine	73.0	51.4	54.8	59.73
Ours	76.4	63.0	60.4	66.60

Table 6: **Robustness toward bank quality.** Experiments on the sensitivity of the similarity between questions and examples. R_t indicates that the examples are the t_{th} similar without any rejection.

Method	Math-level5	AMC12	AMC10
0-shot	50.7	53.6	55.8
few-shot R ₁	52.2 (+1.5)	56.5 (+2.9)	56.7 (+0.9)
few-shot R ₄	46.3 (-4.4)	52.2 (-1.4)	53.7 (-2.1)
Ours R ₁	56.0 (+5.3)	62.3 (+8.7)	60.4 (+4.6)
Ours R ₄	52.2 (+1.5)	61.6 (+8.0)	58.1 (+2.3)

389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1

432
 433 Table 8: Comparison of different step-
 434 level example problem Bank construction
 435 methods. 'GS' represents using Gram-
 436 matical Separation ' as delimiter while
 437 our strategy use the reasoning content to
 438 divide the steps.

Strategy	AMC12	AMC10	MATH
GS	56.5	58.1	74.8
Ours	63.0	60.4	76.4

442
 443 Table 9: Detailed ablation on incorporating retrieving
 444 similar example steps during the reasoning and verify-
 445 ing phases of step-level tree search methods.

Reason	Verify	AMC12	AMC10	MATH	Avg
w/o tree-search		53.6	55.8	73.4	60.9
✗	✗	58.7	59.0	77.8	65.2 (+4.3)
✓	✗	64.4	62.2	79.2	68.6 (+7.7)
✗	✓	61.6	60.4	78.2	66.7 (+5.8)
✓	✓	65.2	63.6	79.4	69.4 (+8.5)

446 in Sec.3.3, retrieving by the entire reasoning path $s_{i-1}, s_{i-2}, \dots, s_1, q$ or only by the immediately
 447 preceding step s_{i-1} . Tab. 7 presents the detailed result. Our method significantly outperforms the
 448 other two strategies, better anticipating the content that needs to be inferred in the current step.

449
 450 **Efficiency.** To provide appropriate examples at the step granularity, we have introduced more so-
 451 phisticated reasoning and retrieval mechanisms. Therefore, the additional time cost also warrants
 452 discussion. The extra cost is primarily attributable to the per-step retrieval and the first-try strategy.
 453 Fortunately, owing to the adoption of appropriate strategies, neither introduces significant time costs.

454 For example retrieval, since the TF-IDF vectors of the example bank can be precomputed, what
 455 needs to be encoded and computed during real-time reasoning is actually minimal, resulting in a
 456 negligible time cost. Quantitatively, a single retrieval takes only a few milliseconds, which accounts
 457 for less than 1% of the time required by any model.

458 For the first-try part, a rejection strategy is adopted: if the similarity of the most similar example step
 459 is still below a certain threshold, we directly use the first-try as the inference content. This strategy
 460 ensures the quality of the provided examples while also improving the efficiency of our approach.
 461 Quantitatively, the first-try attempt will only add 30% time cost to our reasoning.

462 **Case Study.** We provide a specific example of how Booststep improves single-step reasoning
 463 through example steps in Sec. C in the appendix.

464 EXTENDING BOOSTSTEP TO TREE SEARCH

465 The reasoning capability of the reason model and the verifying capability of the critic model are two
 466 core factors of step-level tree search methods, and our strategy can bring benefits in both ways. On
 467 one hand, it can improve the accuracy of generating candidate nodes using the previously mentioned
 468 first-try strategy when reasoning nodes are generated. On the other hand, it can increase the accuracy
 469 of evaluation by introducing similar examples during critic model assessments and therefore ensures
 470 that the correct reasoning nodes are more likely to be preserved. These can be decoupled, allowing
 471 us to demonstrate the effectiveness of each component through ablation studies.

472 We utilize GPT-4o as the reason model, GPT-4o-mini as the PRM and adopt the Pairwise Preference
 473 Reward Model (PPRM) configuration (Zhang et al., 2024b). Detailed settings of our tree search
 474 method will be listed in the appendix.

475 Tab. 9 presents the results of integrating in-context learning into the reasoning and evaluation phases
 476 of Tree Search methods. The results of this ablation study indicate that introducing example steps
 477 can enhance both the reasoning and verifying capabilities of tree search methods. Therefore both
 478 approaches contribute to the improvement of overall reasoning performance.

479 5 CONCLUSION

480 We propose BoostStep, addressing two critical challenges in previous In-Context Learning strate-
 481 gies: granularity mismatch and irrelevant information. BoostStep can provide highly-related exam-
 482 ples at the step granularity, thereby providing fine-grained guidance during an on-going reasoning.
 483 BoostStep is a strong and general approach which can enhance the model's reasoning capabilities
 484 and reduce the sensitivity of the examples. It can break through the limitations of traditional ICL
 485 like achieving 'simple-aids-complex' and cross-modal guidance. Moreover, it can be applied in tree
 486 search methods to enhance the reasoning and verifying capability.

486 REFERENCES
487

488 Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michal Podstawski, Lukas Gian-
489 inazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczek, et al. Graph of
490 thoughts: Solving elaborate problems with large language models. In *Proceedings of the AAAI
491 Conference on Artificial Intelligence*, volume 38, pp. 17682–17690, 2024.

492 Guoxin Chen, Minpeng Liao, Chengxi Li, and Kai Fan. Step-level value preference optimization
493 for mathematical reasoning. *arXiv preprint arXiv:2406.10858*, 2024.

494 OpenCompass Contributors. Opencompass: A universal evaluation platform for foundation models.
495 <https://github.com/open-compass/opencompass>, 2023.

496 Guanting Dong, Chenghao Zhang, Mengjie Deng, Yutao Zhu, Zhicheng Dou, and Ji-Rong Wen.
497 Progressive multimodal reasoning via active retrieval. *arXiv preprint arXiv:2412.14835*, 2024.

498 Haodong Duan, Junming Yang, Yuxuan Qiao, Xinyu Fang, Lin Chen, Yuan Liu, Xiaoyi Dong,
499 Yuhang Zang, Pan Zhang, Jiaqi Wang, et al. Vlmevalkit: An open-source toolkit for evalua-
500 tating large multi-modality models. In *Proceedings of the 32nd ACM international conference on
501 multimedia*, pp. 11198–11201, 2024.

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

505 Edward A Feigenbaum, Julian Feldman, et al. *Computers and thought*, volume 37. New York
506 McGraw-Hill, 1963.

507 Xidong Feng, Ziyu Wan, Muning Wen, Stephen Marcus McAleer, Ying Wen, Weinan Zhang, and
508 Jun Wang. Alphazero-like tree-search can guide large language model decoding and training.
509 *arXiv preprint arXiv:2309.17179*, 2023.

510 Charles R Fletcher. Understanding and solving arithmetic word problems: A computer simulation.
511 *Behavior Research Methods, Instruments, & Computers*, 17(5):565–571, 1985.

512 Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujiu Yang, Nan Duan, and Weizhu Chen.
513 Critic: Large language models can self-correct with tool-interactive critiquing. *arXiv preprint
514 arXiv:2305.11738*, 2023.

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

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

522 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
523 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv
524 preprint arXiv:2103.03874*, 2021.

525 Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-
526 trow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint
527 arXiv:2410.21276*, 2024.

528 Pei Ke, Bosi Wen, Andrew Feng, Xiao Liu, Xuanyu Lei, Jiale Cheng, Shengyuan Wang, Aohan
529 Zeng, Yuxiao Dong, Hongning Wang, et al. Critiquellm: Towards an informative critique
530 generation model for evaluation of large language model generation. In *Proceedings of the 62nd
531 Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp.
532 13034–13054, 2024.

533 Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large
534 language models are zero-shot reasoners. *Advances in neural information processing systems*,
535 35:22199–22213, 2022.

540 Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ra-
 541 masesh, Ambrose Sloane, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. Solving quantitative
 542 reasoning problems with language models. *Advances in Neural Information Processing Systems*,
 543 35:3843–3857, 2022.

544 Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan
 545 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. *arXiv preprint*
 546 *arXiv:2305.20050*, 2023.

547 Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. Program induction by rationale gener-
 548 ation: Learning to solve and explain algebraic word problems. *arXiv preprint arXiv:1705.04146*,
 549 2017.

550 Hongwei Liu, Zilong Zheng, Yuxuan Qiao, Haodong Duan, Zhiwei Fei, Fengzhe Zhou, Wen-
 551 wei Zhang, Songyang Zhang, Dahua Lin, and Kai Chen. Mathbench: Evaluating the theory
 552 and application proficiency of llms with a hierarchical mathematics benchmark. *arXiv preprint*
 553 *arXiv:2405.12209*, 2024a.

554 Jiayu Liu, Zhenya Huang, Chaokun Wang, Xunpeng Huang, Chengxiang Zhai, and Enhong Chen.
 555 What makes in-context learning effective for mathematical reasoning: A theoretical analysis.
 556 *arXiv preprint arXiv:2412.12157*, 2024b.

557 Liangchen Luo, Yinxiao Liu, Rosanne Liu, Samrat Phatale, Harsh Lara, Yunxuan Li, Lei Shu, Yun
 558 Zhu, Lei Meng, Jiao Sun, et al. Improve mathematical reasoning in language models by automated
 559 process supervision. *arXiv preprint arXiv:2406.06592*, 2024.

560 Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri
 561 Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement
 562 with self-feedback. *Advances in Neural Information Processing Systems*, 36, 2024.

563 Chengwei Qin, Aston Zhang, Chen Chen, Anirudh Dagar, and Wenming Ye. In-context learning
 564 with iterative demonstration selection. *arXiv preprint arXiv:2310.09881*, 2023.

565 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 566 Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical
 567 reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

568 Qwen Team. Qwq: Reflect deeply on the boundaries of the unknown, November 2024. URL
 569 <https://qwenlm.github.io/blog/qwq-32b-preview/>.

570 Ke Wang, Junting Pan, Weikang Shi, Zimu Lu, Mingjie Zhan, and Hongsheng Li. Measuring
 571 multimodal mathematical reasoning with math-vision dataset. *arXiv preprint arXiv:2402.14804*,
 572 2024a.

573 Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang
 574 Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. In *Pro-
 575 ceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume
 576 1: Long Papers)*, pp. 9426–9439, 2024b.

577 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 578 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in*
 579 *neural information processing systems*, 35:24824–24837, 2022.

580 Jinyang Wu, Mingkuan Feng, Shuai Zhang, Feihu Che, Zengqi Wen, and Jianhua Tao. Beyond ex-
 581 amples: High-level automated reasoning paradigm in in-context learning via mcts. *arXiv preprint*
 582 *arXiv:2411.18478*, 2024.

583 Yifan Xu, Xiao Liu, Xinghan Liu, Zhenyu Hou, Yueyan Li, Xiaohan Zhang, Zihan Wang, Aohan
 584 Zeng, Zhengxiao Du, Wenyi Zhao, et al. Chatglm-math: Improving math problem-solving in
 585 large language models with a self-critique pipeline. *arXiv preprint arXiv:2404.02893*, 2024.

586 An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jian-
 587 hong Tu, Jingren Zhou, Junyang Lin, et al. Qwen2. 5-math technical report: Toward mathematical
 588 expert model via self-improvement. *arXiv preprint arXiv:2409.12122*, 2024.

594 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu,
 595 Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint*
 596 *arXiv:2505.09388*, 2025.

597

598 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik
 599 Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Ad-*
 600 *vances in Neural Information Processing Systems*, 36, 2024.

601

602 Huaiyuan Ying, Shuo Zhang, Linyang Li, Zhejian Zhou, Yunfan Shao, Zhaoye Fei, Yichuan Ma,
 603 Jiawei Hong, Kuikun Liu, Ziyi Wang, et al. Internlm-math: Open math large language models
 604 toward verifiable reasoning. *arXiv preprint arXiv:2402.06332*, 2024.

605

606 Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhui Chen.
 607 Mammoth: Building math generalist models through hybrid instruction tuning. *arXiv preprint*
 608 *arXiv:2309.05653*, 2023.

609

610 Di Zhang, Xiaoshui Huang, Dongzhan Zhou, Yuqiang Li, and Wanli Ouyang. Accessing gpt-4 level
 611 mathematical olympiad solutions via monte carlo tree self-refine with llama-3 8b. *arXiv preprint*
 612 *arXiv:2406.07394*, 2024a.

613

614 Di Zhang, Jianbo Wu, Jingdi Lei, Tong Che, Jiatong Li, Tong Xie, Xiaoshui Huang, Shufei Zhang,
 615 Marco Pavone, Yuqiang Li, et al. Llama-berry: Pairwise optimization for o1-like olympiad-level
 616 mathematical reasoning. *arXiv preprint arXiv:2410.02884*, 2024b.

617

618 Renrui Zhang, Dongzhi Jiang, Yichi Zhang, Haokun Lin, Ziyu Guo, Pengshuo Qiu, Aojun Zhou,
 619 Pan Lu, Kai-Wei Chang, Yu Qiao, et al. Mathverse: Does your multi-modal llm truly see the
 620 diagrams in visual math problems? In *European Conference on Computer Vision*, pp. 169–186.
 621 Springer, 2025.

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648

A DETAILED EXPERIMENT SETTING

649

650

651

A.1 PROMPT

652

653

Prompt for 0-shot COT: You are a professional math problem solver. Solve the problem step by step and output the final answer within $\boxed{\text{ }}$.

655

656

657

Prompt for problem-level few-shot learning: You are a professional math problem solver. Solve the problem step by step and output the final answer within $\boxed{\text{ }}$. In case you don't know how to solve it, I will give you example problems with their full solutions which you can refer to.

659

660

Example i:

661

Problem: xxx

662

Solution: xxx

663

664

Prompt for first-try in step-level COT: You are a professional math problem solver. I will give you a math problem and part of its solution. And you need to only output the next step of the solution, starting with 'Step i :', where i is the step number. If you think that the final step is derived, put the answer within $\boxed{\text{ }}$.

665

666

667

Prompt for step-level few-shot learning: You are a professional math problem solver. I will give you a math problem and part of its solution. And you need to only output the next step of the solution, starting with 'Step i :', where i is the step number. In case you don't know how to derive the correct content, an example with 'Key Step' will be given. You need to learn how 'Key Step' is derived, and implement similar strategy in your derivation procedure. If you think that the final step is derived, put the answer within $\boxed{\text{ }}$.

668

669

Example Problem: xxx

670

Example Solution: Step 1: xxx, Step 2: xxx, ..., Step i (Key Step): xxx.

671

672

673

674

675

A.2 DETAILS OF GRADING AND METRICS

676

We follow the setting of Opencompass (Contributors, 2023) and VLMEvalKit (Duan et al., 2024). Specifically, we first require the model to put the final answer within $\boxed{\text{ }}$. Then, we use GPT-4o-mini as the judge model to compare the final answer with the ground truth answer. Compared to string matching, this approach can eliminate some false negative evaluations because the same mathematical expression can be expressed in many forms. If the model fails to follow the the expected format in the prompt and the rule-based extraction fails, the solution is directly judged as inconsistent with ground truth.

677

678

679

680

A.3 BENCHMARKS

681

682

We tested our approach on several mathematical benchmarks, including MATH500 (Hendrycks et al., 2021), AQuA (Ling et al., 2017), OlympiadBench-TO (He et al., 2024) and MATHBench (Liu et al., 2024a) and the real questions of American Invitational Mathematics Examination (AIME) in year 2023 and 2024. Specifically, we use the Olympiad-TO (text-only) subset of OlympiadBench and the application problems in college-level and high-level difficulty of MATHBench.

683

684

685

In addition, we collected problems from the AMC-10 and AMC-12 competitions to serve as more challenging benchmarks. They include real questions from the 2018-2023 AMC exams.

686

687

688

689

For multi-modal math benchmarks, we use MathVision-Mini (Wang et al., 2024a) and vision-dominant version of problems in MathVerse-Mini (Zhang et al., 2025).

702
703
704

Question: For some value of x , $0 < x < 180$, $\tan 53^\circ \tan 81^\circ \tan x^\circ = \tan 53^\circ + \tan 81^\circ + \tan x^\circ$. Find x .

705
706
707
708
709

First-Try: Use the identity for the tangent of the sum of angles, $\tan(A + B + C) = \tan A + \tan B + \tan C - \tan A \tan B \tan C$ (**Wrong Equation**) and compare it with the given equation.

710
711
712
713
714
715

Example Step: From the addition formula for tangent, $\tan(A + B + C) = \frac{\tan A + \tan B + \tan C - \tan A \tan B \tan C}{1 - (\tan A \tan B + \tan A \tan C + \tan B \tan C)}$ (**Key Example Equation**). Since $A + B + C = 180^\circ$, this is 0. Hence, $\tan A + \tan B + \tan C = \tan A \tan B \tan C$.

716
717
718
719
720
721

Final-Reasoning: Notice that the equation resembles the identity for the tangent of a sum of angles. We can use the identity: $\tan(A + B + C) = \frac{\tan A + \tan B + \tan C - \tan A \tan B \tan C}{1 - (\tan A \tan B + \tan A \tan C + \tan B \tan C)}$ (**Correct Equation**)

722
723
724
725

Next step: Since the equation is symmetric in terms of the tangents, we can assume that the angles 53° , 81° , and x° sum to 180° . Therefore, we have: $53^\circ + 81^\circ + x^\circ = 180^\circ$. (**Correct Conclusion**)

726
727
728
729
730

Figure 4: A specific example of adjusting reasoning during real-time inference through step-level in-context learning. The first try uses a wrong equation while the retrieving example step guides the model to use the correct equation and get the correct conclusion.

B DETAILED SETUP FOR EXAMPLE-GUIDED STEP-LEVEL TREE SEARCH

733
734
735
736
737
738

In the setup for tree search methods, we utilize GPT-4o as the reason model and employ GPT-4o-mini as the Process-supervised Reward Model (PRM). For the PRM, we adopted the Pairwise Preference Reward Model (PPRM) configuration (Zhang et al., 2024b). Specifically, PPRM transforms the absolute rewards calculation into preference predictions between solutions to calculate rewards. This approach reduces the variability associated with scoring characteristics and thus leads to a more robust and consistent evaluation of different solutions.

739
740
741
742
743
744
745
746

The complete reasoning process in our experiment is as follows: we start with the target problem as the root node and obtain two initial solution steps through sampling to serve as the two initial parent nodes. In each step-level reasoning phase, we expand these two parent nodes through sampling, generating four candidate child nodes. Using the PPRM, we select the two child nodes with higher confidence to become the parent nodes for the next step of reasoning. This process continues until both candidate nodes have completed their reasoning paths, resulting in the final answers. Finally, PPRM is used to select the ultimate answer from these two reasoning paths.

747
748

C CASE STUDY

Here we demonstrate a specific example of how our step-level in-context learning boosts step-level reasoning. Given the question, we first let the model have a first try on step one. Unfortunately, because the model is unfamiliar with trigonometric functions, it makes an error on the tangent sum formula, therefore leading to a wrong step. However, we can get a rough idea of what the model wants to calculate at this step according to the first try. Then, we find a similar step that correctly leverages the tangent sum formula in the step-level example problem bank. Therefore, with the guidance provided, the model correctly applied the tangent sum formula during the second reasoning attempt and arrived at the correct answer.