
SIGMA: Refining Large Language Model Reasoning via Sibling-Guided Monte Carlo Augmentation

Anonymous Author(s)

Affiliation

Address

email

A Extended Results of Our Method on the 60K-Sample Dataset

In Appendix A, we present the extended evaluation of our SIGMA method on a larger 60K-sample dataset. We first analyze performance gains across three backbone architectures in Appendix A.1 including LLAMA3-8B[3], MISTRAL-7B[6], and DEEPSEEK MATH-7B[13], demonstrating consistent improvements over existing fine-tuning baselines. We then provide a benchmark level breakdown across six diverse mathematical reasoning tasks in Appendix A.2 (GSM8K[2], MATH[5], College[14], DeepMind[12], Olympiad[4], and Theorem[1]), highlighting SIGMA’s robustness and its ability to generalize across varying levels of problem difficulty.

A.1 Analysis of Performance Across Different Base Models

To further demonstrate the effectiveness of our proposed SIGMA method, we expanded the dataset from 30K to 60K samples. The experimental results are presented in Table 1. Notably, our model LLAMA3-SIGMA-8B-60K outperforms LLAMA3-8B-DART-MATH, which was fine-tuned on a much larger dataset of 590K examples, across the average performance on six benchmark datasets[15]. Below, we provide a detailed analysis based on different base models.

LLaMA3-8B. Given the same number of training samples, our model achieves an average score that is 1 point higher than MathFusion[10]. In particular, it surpasses MathFusion by 5.8 on the DeepMind benchmark and by 3.2 on GSM8K. Even when compared with models trained on significantly more data, our 60K-trained model surpasses all others, (including DART-Math-590K) by 0.5 in average score and achieves the highest score on the DeepMind benchmark (49.2) among all models.

Mistral-7B. When trained on the same dataset size (60K), our MISTRAL-SIGMA-7B-60K model outperforms all other models. It achieves a 0.8 point gain over MathFusion and a 4-point improvement over DART-Math. Notably, its score on the DeepMind benchmark reaches 46.1, which is even higher than models fine-tuned on the full 590K dataset, outperforming DART-Math by 1.8 .

DeepSeekMath-7B. Using an identical number of training examples, our DEEPSEEK MATH-SIGMA-7B-60K model outperforms MathFusion by 1.8 and DART-Math by 1.9. On the MATH dataset, it achieves a score of 56.5, the highest across all models regardless of training data size, exceeding DART-Math (590K) by 2.9 . Overall, the average performance of our model is very close to that of DART-Math (590K).

A.2 Analysis of Performance Across Six Different Benchmarks

To systematically quantify how each data construction technique impacts performance across our six benchmarks: GSM8K, MATH, College, DeepMind, Olympiad, and Theorem. We present different benchmarks’ analysis comparing SIGMA with MetaMath [16], MMIQC [8], RefAug [17],

Table 1: Performance comparison across base models and training strategies. Results are reported as exact-match accuracy under 0-shot greedy decoding(temperature = 0). All scores were obtained from the first attempt. Arrows indicate accuracy changes relative to the baseline in blue background. Some results are quoted from MathFusion [10] and some are quoted from DART-Math [15].

Model	# Samples	In-Domain		Out-of-Domain			AVG	
		MATH	GSM8K	College	DM	Olympiad Theorem		
LLaMA3-8B (General Base Model)								
Llama3-8B-MetaMath	400K	32.5	77.3	20.6	35.0	5.5	13.8	30.8
Llama3-8B-RFT	590K	39.7	81.7	23.9	41.7	9.3	14.9	35.2
Llama3-8B-DART-Math	590K	46.6	81.1	28.8	48.0	14.5	19.4	39.7
Llama3-8B-MMIQC	2.3M	39.5	77.6	29.5	41.0	9.6	16.2	35.6
Llama3-SIGMA-8B-15K	15K	36.0	82.0 ^{↑4.1}	24.2	42.0	10.5	22.0 ^{↑5.0}	36.1 ^{↑0.5}
Llama3-SIGMA-8B-30K	30K	40.8 ^{↑2.0}	79.5 ^{↑1.6}	26.3 ^{↑0.8}	47.5 ^{↑5.5}	12.7 ^{↑0.1}	19.1 ^{↑2.1}	37.7 ^{↑2.1}
MathFusion (Sequential)	30K	38.8	77.9	25.1	42.0	12.6	17.0	35.6
MathFusion (Conditional)	30K	34.7	76.9	21.2	27.4	11.9	15.5	31.3
MathFusion (Parallel)	30K	38.1	75.4	25.5	41.9	11.9	18.9	35.3
Llama3-8B-MetaMat	60K	28.7	78.5	19.7	31.3	5.3	16.1	29.9
Llama3-8B-MMIQC	60K	24.4	69.7	13.4	30.9	5.2	10.6	25.7
Llama3-8B-RefAug	60K	20.3	68.6	15.5	29.1	5.5	13.0	25.3
Llama3-8B-DART-Math	60K	39.6	82.2	27.9	39.9	12.9	22.9	37.6
MathFusion-Llama3-8B	60K	46.5	79.2	27.9	43.4	17.2	20.0	39.0
Llama3-SIGMA-8B-60K	60K	44.9	82.4 ^{↑3.2}	28.1	49.2 ^{↑5.8}	15.3	21.3 ^{↑1.3}	40.2 ^{↑1.2}
Mistral-7B-v0.1 (General Base Model)								
Mistral-7B-MetaMath	400K	29.8	76.5	19.3	28.0	5.9	14.0	28.9
Mistral-7B-WizardMath-V1.1	418K	32.3	80.4	23.1	38.4	7.7	16.6	33.1
Mistral-7B-RFT	590K	38.7	82.3	24.2	35.6	8.7	16.2	34.3
Mistral-7B-DART-Math	590K	45.5	81.1	29.4	45.1	14.7	17.0	38.8
Mistral-SIGMA-7B-15K	15K	30.0	75.3 ^{↑1.4}	20.8 ^{↑1.9}	39.5 ^{↑10.2}	7.7	16.3 ^{↑0.8}	31.6 ^{↑1.7}
Mistral-SIGMA-7B-30K	30K	35.5 ^{↑2.8}	78.6 ^{↑4.7}	22.1 ^{↑3.2}	43.8 ^{↑14.5}	11.1 ^{↑1.8}	18.0 ^{↑2.5}	34.9 ^{↑5.0}
MathFusion (Sequential)	30K	32.7	73.9	18.9	29.3	9.3	15.5	29.9
MathFusion (Conditional)	30K	26.3	73.0	15.6	21.4	7.3	12.8	26.1
MathFusion (Parallel)	30K	30.9	75.1	20.9	26.5	11.0	15.2	29.9
Mistral-7B-MMIQC	60K	17.3	61.4	11.1	13.5	5.0	5.9	19.0
Mistral-7B-RefAug	60K	17.4	63.1	12.5	18.1	3.9	11.1	21.0
Mistral-7B-MetaMath	60K	22.7	70.8	14.1	27.2	5.0	12.2	25.3
Mistral-7B-DART-Math	60K	34.1	77.2	23.4	36.0	8.7	18.2	32.9
MathFusion-Mistral-7B	60K	41.6	79.8	24.3	39.2	13.6	18.1	36.1
Mistral-SIGMA-7B-60K	60K	40.3	79.2	24.1	46.1 ^{↑6.9}	12.3	19.2 ^{↑1.1}	36.9 ^{↑0.8}
DeepSeekMath-7B (Math-Specialized Base Model)								
DeepSeekMath-7B-RFT	590K	53.0	88.2	41.9	60.2	19.1	27.2	48.3
DeepSeekMath-7B-DART-Math	590K	53.6	86.8	40.7	61.6	21.7	32.2	49.4
DeepSeekMath-7B-Instruct	780K	46.9	82.7	37.1	52.2	14.2	28.1	43.5
DeepSeekMath-7B-MMIQC	2.3M	45.3	79.0	35.3	52.9	13.0	23.4	41.5
DeepSeekMath-SIGMA-7B-15K	15K	52.2 ^{↑2.3}	81.1 ^{↑4.5}	37.4	64.5	20.3	26.1 ^{↑3.3}	47.0 ^{↑1.3}
DeepSeekMath-SIGMA-7B-30K	30K	54.9 ^{↑5.0}	82.2 ^{↑5.6}	36.7	67.2 ^{↑2.6}	21.6	26.6 ^{↑3.8}	48.2 ^{↑2.5}
MathFusion (Sequential)	30K	49.9	76.6	38.8	64.6	21.6	22.8	45.7
MathFusion (Conditional)	30K	48.5	74.6	37.0	55.2	19.3	19.0	42.3
MathFusion (Parallel)	30K	50.9	76.7	38.9	62.2	19.0	23.8	45.3
DeepSeekMath-7B-MMIQC	60K	26.3	60.6	19.2	41.5	10.4	6.8	27.5
DeepSeekMath-7B-RefAug	60K	33.1	71.6	26.2	35.4	10.5	14.0	31.8
DeepSeekMath-7B-MetaMath	60K	40.0	79.0	33.2	45.9	9.5	18.9	37.8
DeepSeekMath-7B-DART-Math	60K	51.4	82.9	39.1	62.8	21.0	27.4	47.4
MathFusion-DeepSeekMath-7B	60K	53.4	77.9	39.8	65.8	23.3	24.6	47.5
DeepSeekMath-SIGMA-7B-60K	60K	56.5 ^{↑3.1}	81.7 ^{↑3.8}	37.2	68.4 ^{↑2.6}	22.5	29.3 ^{↑4.7}	49.3 ^{↑1.8}

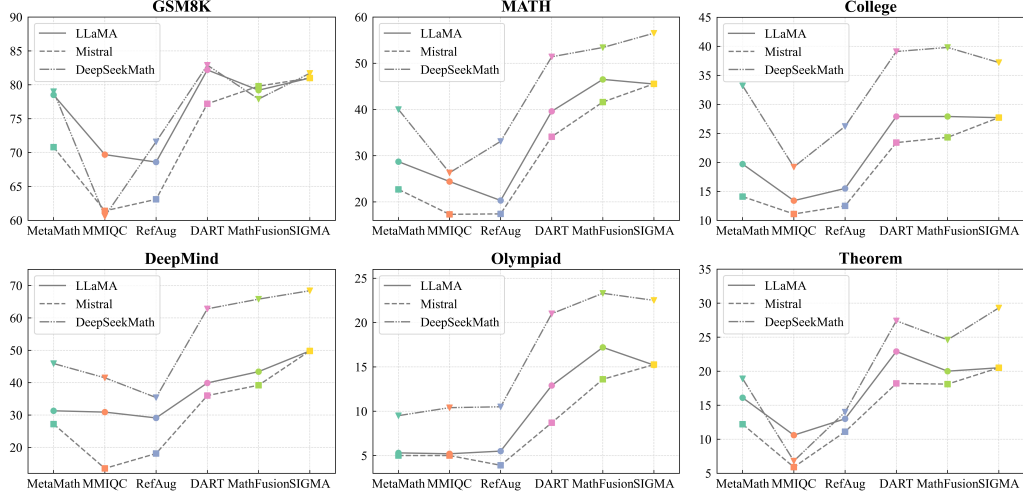


Figure 1: Performance comparison of models fine-tuned on **60k**-sample datasets generated by different methods, evaluated across six benchmark tasks. Different colored dots represent different methods.

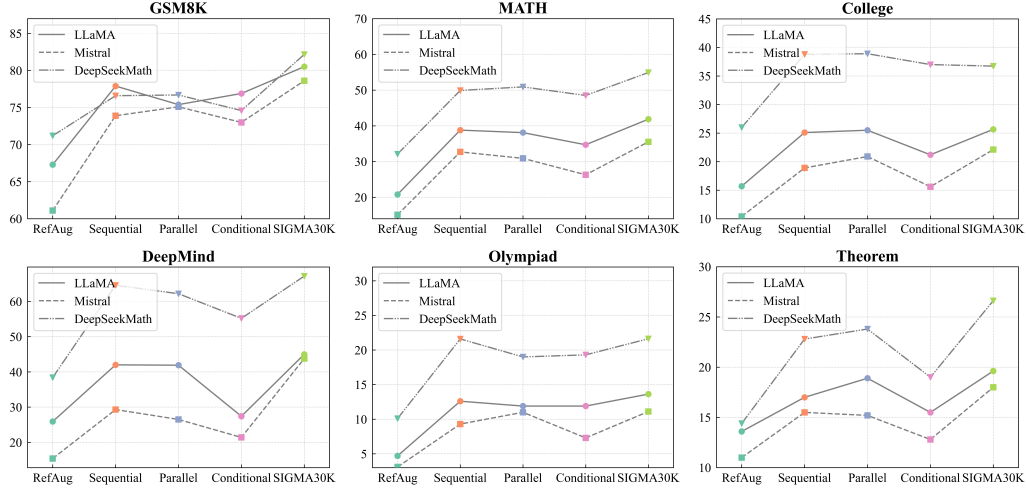


Figure 2: Performance comparison of models fine-tuned on **30k**-sample datasets generated by different methods, evaluated across six benchmark tasks. Different colored dots represent different methods.

34 DART-Math, and MathFusion (Figure 1). We highlight both absolute and relative gains, as well
 35 as consistency across the three backbone models, to demonstrate SIGMA’s robustness and general
 36 applicability.

37 **60K-Sampled Datasets Comparison Across Six Benchmarks.** *GSM8K*: SIGMA consistently
 38 outperforms MathFusion across all base architectures. Furthermore, it matches the performance
 39 of DART-Math, and remarkably, all three models converge to nearly identical accuracies when
 40 fine-tuned with our method. *MATH*: On the more demanding MATH benchmark, SIGMA deliv-
 41 ers substantial gains over DART-Math for every backbone. In particular, it also exceeds MathFu-
 42 sion’s results on both LLaMA and DeepSeekMath, underscoring its robustness on complex arith-
 43 metic reasoning. *College*: In the College evaluation, SIGMA yields pronounced improvements
 44 for Mistral and DeepSeekMath. This demonstrates the strength of our data generation in enhanc-
 45 ing performance on intermediate-difficulty problems. *DeepMind*: SIGMA outstrips all competing
 46 approaches—including DART-Math and MathFusion—by a wide margin on DeepMind. The consist-
 47 ent uplift across every model highlights its effectiveness on advanced reasoning tasks. *Olympiad*:

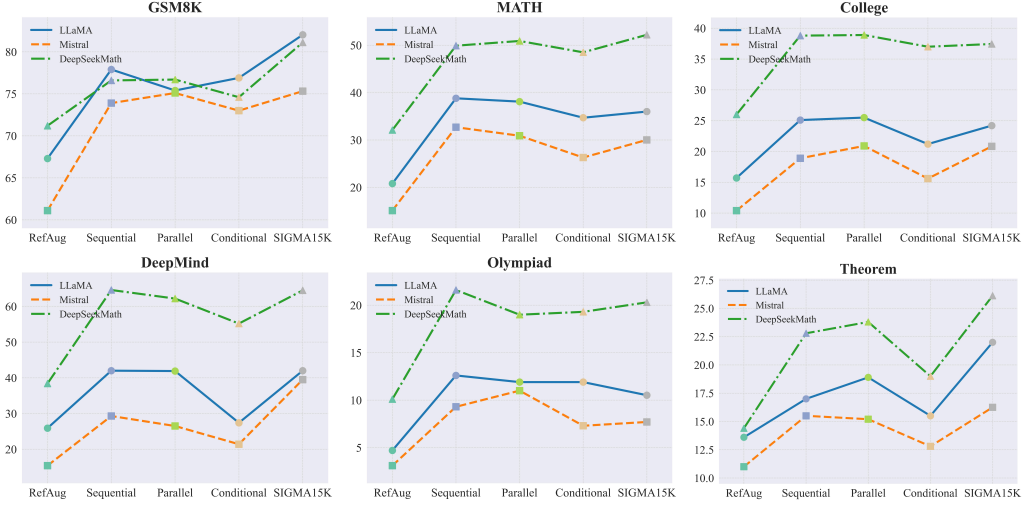


Figure 3: Performance comparison of models fine-tuned on **15k**-sample datasets generated by our SIGMA15k and **30k**-sample datasets generated by other methods, evaluated across six benchmark tasks. Different colored dots represent different methods.

For the Olympiad benchmark, SIGMA falls slightly short of MathFusion on LLaMA and Mistral but outperforms it on DeepSeekMath. Crucially, it still achieves clear, uniform improvements over DART-Math across all three architectures. *Theorem*: On Theorem, SIGMA secures a clear advantage over MathFusion, confirming its ability to generate data that strengthens formal and symbolic reasoning.

30K-Sampled Datasets Comparison Across Six Benchmarks. As shown in Figure 2, *GSM8K*: SIGMA surpasses all baselines on GSM8K when fine-tuned with 30 K samples, demonstrating consistent superiority across multiple base models. *MATH*: On MATH benchmark, SIGMA delivers marked gains over every competing method, showing stable uplift across different base models. *College*: In College evaluations, SIGMA yields notable improvements for each architecture, demonstrating that our data construction effectively enhances general-purpose models’ ability for reasoning challenging mathematical problems. *DeepMind*: SIGMA delivers substantial gains over MathFusion on DeepMind tasks, underscoring its strong adaptability to DeepMind dataset. *Olympiad*: For Olympiad problems, SIGMA maintains consistent gains over DART-Math across all backbones and narrows the gap with MathFusion, confirming robust performance on competition-level questions. *Theorem*: On the Theorem dataset, SIGMA achieves uniform improvements over every baseline, illustrating that its advantages in theorem reasoning persist regardless of dataset size.

SIGMA-15K vs. 30K-Scale Methods: Performance Across Six Benchmarks. As shown in Figure 3, *GSM8K*: With 15K training samples, SIGMA still leads all augmentation strategies on GSM8K, delivering uniformly higher accuracies across LLaMA, Mistral and DeepSeekMath and narrowing performance disparities among them. *MATH*: Even at reduced scale, SIGMA achieves clear improvements on MATH, outperforming each baseline and preserving its edge on all three backbones. *College*: In the College evaluation, SIGMA demonstrates a pronounced advantage, boosting each architecture’s capability to tackle intermediate-difficulty tasks despite the limited training data. *DeepMind*: SIGMA adapts effectively to the DeepMind dataset, surpassing MathFusion and other approaches across all backbones, which underscores its resilience on complex reasoning challenges. *Olympiad*: On Olympiad questions, SIGMA outperforms DART-Math for every base model and approaches MathFusion’s standard, confirming its steady competitiveness on high-difficulty problems even at reduced scale. *Theorem*: For the Theorem benchmark SIGMA sustains its lead over all competing strategies, demonstrating that its strengths in reasoning persist when trained on only fifteen thousand instances.

79 B Training Setup

80 All fine-tuning experiments were carried out on four NVIDIA H100 GPUs with DeepSpeed2 ZeRO
81 [11] optimizations, operating in mixed precision [9] (FP16) to maximize memory efficiency and
82 throughput.

83 The computation sequence token length was fixed at 4096 to capture long range mathematical rea-
84 soning. We set `per_device_train_batch_size=8` and used `gradient_accumulation_steps=4`
85 to accumulate gradients.

86 We used the AdamW [7] optimizer with weight decay of 0.01. A cosine decay learning rate
87 scheduler was applied with a linear warmup over the first 3 percent of total training steps
88 (`warmup_ratio=0.03`) to prevent instability. All models were trained for 3 epochs.

89 Initial learning rates were tuned per model:

- 90 • DeepSeekMath-7B: 5×10^{-5}
- 91 • Mistral-7B: 4×10^{-6}
- 92 • LLaMA3-8B: 1×10^{-5}

93 We logged metrics at regular intervals during training and validation. Random seeds were fixed and
94 model checkpoints were saved after each epoch to ensure reproducibility.

95 C Examples


96 In Appendix C, we present detailed examples and Q&A examples to provide deeper insight into
97 the mechanism by which the SIGMA method refines reasoning data. Appendix C.1 illustrates how
98 SIGMA refines reasoning steps on both geometry and algebra problems. Appendix C.2 presents
99 additional SIGMA-refined reasoning paths for questions from MATH and GSM8K datasets.

100 C.1 Detailed examples

Example 1. Geometry

Question: The altitude to the hypotenuse of a triangle with angles of 30° and 60° is 3 units. What is the area of the triangle, in square units? Express your answer in simplest radical form.

#Picture below is compiled from Asymptote code in question line.



Answer: The area is $6\sqrt{3}$ square units.

MCTS Selected Path:

(0-1): Step: Calculate the length of the hypotenuse:

$$\text{hypotenuse} = 2 \cdot 3$$

(0-1-2): Step: Calculate the length of the base:

$$\text{base} = \frac{\text{hypotenuse}}{2}$$

(0-1-2-1): Step: Calculate the length of the altitude:

$$\text{altitude} = 3$$

(0-1-2-1-1): Step: Calculate the area of the triangle:

$$\text{area} = 0.5 \cdot \text{base} \cdot \text{altitude}$$

The answer is: 4.5

Step-by-Step Sibling Comparison and Critique:

Step 1:

- **Sibling(0-2):**Step:Calculate the length of the hypotenuse using the altitude and the 60-degree angle.We use the relationship in a 30° - 60° - 90° triangle.Since the altitude from the right angle to the hypotenuse is 3 units, the hypotenuse is:
 $\text{hypotenuse} = 2\sqrt{3}$.
- **Sibling(0-0):**Step:Calculate the length of the legs of the triangle.Using triangle ratios again, the longer leg (opposite the 60° angle) is: $\text{leg length} = 3\sqrt{3}$.
- **Evaluation:** Step 0-0 incorrectly assumes a leg length without relating it to the given altitude; step 0-2 correctly uses the 30° - 60° - 90° property to link the altitude and hypotenuse; step 0-1 computed hypotenuse but resulted in wrong value.

Step 2:

- **Sibling(0-1-0):**Step:Calculate the length of the side opposite the 30-degree angle: $\text{side}_{30} = \text{hypotenuse} / 2$
- **Sibling(0-1-1):**Step:Compute the base using the Pythagorean theorem: $\text{base} = \sqrt{\text{hypotenuse}^2 - \text{altitude}^2}$
- **Evaluation:** Step 0-1-0 conflates the 30° as opposite side with the base without justification; step 0-1-1 applies Pythagorean theorem correctly but adds unnecessary complexity; step 0-1-2 halves the hypotenuse that only holds in an isosceles right triangle.

Step 3:

- **Sibling(0-1-2-0):**Step:Calculate the length of the height: $\text{height}=3$
- **Sibling(0-1-2-2):**Step:Compute the altitude using triangle legs: $\text{altitude}=(\text{leg1} * \text{leg2}) / \text{hypotenuse}$
- **Evaluation:**Step 0-1-2-0 merely restates the given;step 0-1-2-1 correctly extracts the altitude;step 0-1-2-2 applies a geometric formula but assumes leg values.

Step 4:

- **Sibling(0-1-2-1-0):**Step:Calculate square: $\text{base} * \text{altitude} / 2 = 6$
- **Sibling(0-1-2-1-2):**Step:Calculate answer: $0.5 * \text{base} * \text{altitude} = 6$
- **Evaluation:** Step 0-1-2-1-0 calculate the square; step 0-1-2-1-1 calculate the area of the triangle; step 0-1-2-1-2 compute the answer.

Key points:

- For step 0-1, explicitly derive the hypotenuse length from the altitude to the hypotenuse using the correct geometric relationship in a 30-60-90 triangle.
- For step 0-1-2, form and solve the altitude equation for that parameter.
- For step 0-1-2-1, substitute back to obtain all side lengths.
- For step 0-1-2-1-1, keep the square computing logic still.

Optimized Output:**Step 1:** Assign variables to the sides of the 30-60-90 triangle.Let the shortest leg (opposite 30°) be x . Then the hypotenuse is $2x$, and the longer leg (opposite 60°) is $x\sqrt{3}$. Thus, the legs are x and $x\sqrt{3}$, and the hypotenuse is $2x$.**Step 2:** Use the given altitude length to solve for x .

We are told the altitude to the hypotenuse is 3 units, so:

$$3 = \frac{x\sqrt{3}}{2} \implies x = \frac{3 \times 2}{\sqrt{3}} = \frac{6}{\sqrt{3}}.$$

Rationalizing the denominator:

$$x = \frac{6}{\sqrt{3}} \times \frac{\sqrt{3}}{\sqrt{3}} = \frac{6\sqrt{3}}{3} = 2\sqrt{3}.$$

Step 3: Calculate the lengths of all sides using $x = 2\sqrt{3}$.- Short leg (opposite 30°): $x = 2\sqrt{3}$.- Longer leg (opposite 60°): $x\sqrt{3} = 2\sqrt{3} \times \sqrt{3} = 2 \times 3 = 6$.- Hypotenuse: $2x = 2 \times 2\sqrt{3} = 4\sqrt{3}$.**Step 4:** Calculate the area of the triangle.

The area of a right triangle is half the product of its legs:

$$\text{Area} = \frac{1}{2} \times (\text{short leg}) \times (\text{long leg}) = \frac{1}{2} \times 2\sqrt{3} \times 6 = \frac{1}{2} \times 12\sqrt{3} = 6\sqrt{3}.$$

The answer is:

103

Example 2. Algebra**Question:** Four positive integers A , B , C and D have a sum of 36. If $A + 2 = B - 2 = C \times 2 = D \div 2$, what is the value of the product $A \times B \times C \times D$?**Answer:** The correct product is .**MCTS Selected Path:****(0-2):** Step: Define the variables A , B , C , and D .**(0-2-2):** Step: Write the system of equations:

$$\begin{cases} A + B + C + D = 36, \\ A + 2 = B - 2, \\ B = C, \\ C = \frac{D}{2}. \end{cases}$$

(0-2-2-1): Step: Solve the system to obtain

$$A = 8, \quad B = 8, \quad C = 8, \quad D = 16.$$

(0-2-2-1-0): Step: Compute the product:

$$A \times B \times C \times D = 8 \times 8 \times 8 \times 16 = 4096.$$

The answer is: (Incorrect)**Step-by-Step Sibling Comparison and Critique:****Step 1:**

104

- **Sibling(0-0):**Step:Define variables: $A + B + C + D = 36$
- **Sibling(0-1):**Step:Define sum: $SUM = 36$
- **Evaluation:** Step 0-1 correctly begins by defining the total sum constraint explicitly; step 0-1 gives clear definition of total sum; step 0-2 omits in this step.

Step 2:

- **Sibling(0-2-0):**Step:Define equations: $A = C * 2 - 2$; $B = C * 2 + 2$;
 $D = C * 2 * 2$
- **Sibling(0-2-1):**Step:Define relationships: $A = B - 4$; $B = C * 2 + 2$;
 $D = A * 2 + 4$
- **Evaluation:** Step 0-2-0 expresses all variables in terms of a single variable C , which provides a clearer structure; step 0-2-1 give recurrent definition of relative variables;step 0-2-2 instead defines relationships in isolated pairs.

Step 3:

- **Sibling(0-2-2-0):**Step:Solve equations: $A + B + C + D = 36$; $A + 2 = B - 2$; $C * 2 = D / 2$
- **Sibling(0-2-2-2):**Step:Solve equations: $5 * C + 4 = 36$
- **Evaluation:** Step 0-2-2-0 captures the equality logic via $C * 2 = D/2$ correctly; step 0-2-2-2 express all variables using C ; step 0-2-2-1 misrepresents this logic using $B = C$ and $C = D/2$.

Step 4:

- **Sibling(0-2-2-1-1):**Step:Calculate requirement: $A * B * C * D = 4096$
- **Sibling(0-2-2-1-2):**Step:Calculate answer: $A * B * C * D = 4096$
- **Evaluation:** Step 0-2-2-1-1 calculate the required output; step 0-2-2-1-2 calculate the answer ; step 0-2-2-1-0 compute the product.

Key Points:

- For step 0-2, explicitly state constraints in initial setup.
- For step 0-2-2, express all variables using a single reference variable.
- For step 0-2-2-1, correctly encode and solve the system with logical relationships.
- For step 0-2-2-1-0, keep the computing logic.

Optimized Output:

Step 1: Introduce a common value k so that

$$A + 2 = B - 2 = 2C = \frac{D}{2} = k \quad \text{and} \quad A + B + C + D = 36.$$

Step 2: Express each variable in terms of k :

$$A = k - 2, \quad B = k + 2, \quad C = \frac{k}{2}, \quad D = 2k,$$

then substitute into the sum:

$$(k - 2) + (k + 2) + \frac{k}{2} + 2k = 36 \implies \frac{9k}{2} = 36 \implies k = 8.$$

Step 3: Determine the four integers:

$$A = 6, \quad B = 10, \quad C = 4, \quad D = 16,$$

which indeed satisfy both the sum and the four-way equality.

Step 4: Compute the product:

$$A \times B \times C \times D = 6 \times 10 \times 4 \times 16 = 3840.$$

The answer is: 3840

106 **C.2 Q&A examples**

Example 1

Question: Mari made 4 more than five times as many decorated buttons as Kendra. Sue made half as many as Kendra. Sue made 6 buttons. How many did Mari make?

Answer:

Step 1: Sue made 6 buttons, so Kendra made $2 \times 6 = 12$.

Step 2: Mari made 4 more than $5 \times$ Kendra's count: $5 \times 12 + 4$.

Step 3: Compute $5 \times 12 = 60$, then add 4 to get 64.

Step 4: Therefore, Mari made 64 buttons.

The answer is: 64

107

Example 2

Question: When three standard dice are tossed, the numbers a, b, c are obtained. Find the probability that

$$(a - 1)(b - 1)(c - 1) \neq 0$$

Answer:

Step 1: Total outcomes $= 6^3 = 216$.

Step 2: $(a - 1)(b - 1)(c - 1) = 0$ iff at least one die shows 1.

Step 3: Favorable: none is 1 $\implies 5^3 = 125$ outcomes.

Step 4: Probability $= \frac{125}{216}$.

Step 5: Hence the probability is $\frac{125}{216}$.

The answer is: $\frac{125}{216}$

108

Example 3

Question: What is the minimum value of y if $y = 3x^2 + 6x + 9$?

Answer:

Step 1: Rewrite or note vertex of parabola $3(x^2 + 2x) + 9 = 3(x + 1)^2 + 6$.

Step 2: Since $(x + 1)^2 \geq 0$, minimum occurs at $x = -1$.

Step 3: Substitute $x = -1$: $y = 3 \cdot 0 + 6 = 6$.

Step 4: Therefore, the minimum value is 6.

The answer is: 6

109

Example 4

Question: You are given a sequence of 58 terms $P + n$ where P is the product of all primes ≤ 61 and $n = 2, 3, \dots, 59$. Let N be the number of primes in this sequence. Find N .

Answer:

Step 1: If n is prime (17 values), $P + n \equiv 0 \pmod{n}$ so composite.

Step 2: Remaining $58 - 17 = 41$ terms with composite n may be prime.

Step 3: Known result for Euclid-type numbers gives exactly 16 primes among them.

Step 4: Thus $N = 16$.

The answer is: 16

110

Example 5

Question: A rectangle has length 40 m and width 20 m. A similar smaller rectangle has half the length and width. What is the area of the smaller rectangle?

111

Answer:

Step 1: Scale factor = $\frac{1}{2}$, so smaller sides are 20 m and 10 m.

Step 2: Area = $20 \times 10 = 200 \text{ m}^2$.

Step 3: Alternatively area factor $(\frac{1}{2})^2 = \frac{1}{4}$ of 800 gives 200.

Step 4: Therefore, the area is 200.

The answer is: 200

112

Example 6

Question: Last year, Isabella took 7 math tests with distinct integer scores 91–100. Each running average after n tests was an integer. She scored 95 on the 7th test. What was her 6th test score?

Answer:

Step 1: Let S_n be sum after n tests; $S_n \equiv 0 \pmod{n}$.

Step 2: From $S_7 = S_6 + 95 \equiv 0 \pmod{7} \implies S_6 \equiv 3 \pmod{7}$.

Step 3: Also $S_6 \equiv 0 \pmod{6}$, so $S_6 \equiv 24 \pmod{42}$.

Step 4: In allowable range $546 \leq S_6 \leq 600$, the valid sum is 570, giving $s_6 = 570 - S_5 = 100$.

Step 5: Therefore, the 6th score is 100.

The answer is: 100

113

Example 7

Question: Let p, q be distinct solutions to $\frac{4x - 12}{x^2 + 2x - 15} = x + 2$. If $p > q$, find $p - q$.

Answer:

Step 1: Domain excludes $x = -5, 3$. Multiply by $x^2 + 2x - 15$ to get $x^3 + 4x^2 - 15x - 18 = 0$.

Step 2: Rational root $x = 3$ is extraneous. Factor gives $(x - 3)(x^2 + 7x + 6) = 0$.

Step 3: Solve $x^2 + 7x + 6 = 0 \implies x = -1, -6$.

Step 4: Thus $p = -1, q = -6$, so $p - q = 5$.

Step 5: Therefore, the difference is 5.

The answer is: 5

114

Example 8

Question: A vampire needs 7 gallons of blood per week, sucking 2 pints per person. How many people per day must he feed on?

Answer:

Step 1: $7 \text{ gal} \times 8 \text{ pt/gal} = 56 \text{ pints/week}$.

Step 2: At 2 pt/person, he needs $56/2 = 28 \text{ people/week}$.

Step 3: Dividing by 7 days gives $28/7 = 4 \text{ people/day}$.

Step 4: Therefore, he must feed on 4 people each day.

The answer is: 4

115

Example 9

Question: Marla is mixing lilac paint 70% blue, 20% red, rest white. If she adds 140 oz of blue, how many ounces of white does she add?

Answer:

Step 1: White% = $100 - (70 + 20) = 10\%$.

Step 2: Total T : $0.7T = 140 \implies T = 200 \text{ oz}$.

Step 3: White amount = $0.1 \times 200 = 20 \text{ oz}$.

Step 4: Therefore, she adds 20 ounces of white.

116

The answer is: 20

117

Example 10

Question: By partial fractions,

$$\frac{1}{x(x+1)(x+2)(x+3)(x+4)} = \frac{A}{x} + \frac{B}{x+1} + \frac{C}{x+2} + \frac{D}{x+3} + \frac{E}{x+4}$$

for some constants A, B, C, D, E . Find $A + B + C + D + E$.

Answer:

Step 1: Multiply both sides by $x(x+1)(x+2)(x+3)(x+4)$ to get

$$\begin{aligned} 1 &= A(x+1)(x+2)(x+3)(x+4) \\ &\quad + Bx(x+2)(x+3)(x+4) \\ &\quad + Cx(x+1)(x+3)(x+4) \\ &\quad + Dx(x+1)(x+2)(x+4) \\ &\quad + Ex(x+1)(x+2)(x+3). \end{aligned}$$

Step 2: Evaluate at the roots:

$$\begin{aligned} x = 0 : \quad 1 &= 24A \implies A = \frac{1}{24}, \\ x = -1 : \quad 1 &= -6B \implies B = -\frac{1}{6}, \\ x = -2 : \quad 1 &= 4C \implies C = \frac{1}{4}, \\ x = -3 : \quad 1 &= -6D \implies D = -\frac{1}{6}, \\ x = -4 : \quad 1 &= 24E \implies E = \frac{1}{24}. \end{aligned}$$

Step 3: Sum the constants:

$$A + B + C + D + E = \frac{1}{24} - \frac{4}{24} + \frac{6}{24} - \frac{4}{24} + \frac{1}{24} = \frac{1-4+6-4+1}{24} = \frac{0}{24} = 0.$$

Step 4: Therefore, the sum is zero.

The answer is: 0

118

Example 11

Question: Let P be a point outside of circle O . A segment is drawn from P such that it is tangent to circle O at T . Meanwhile, a secant from P intersects O at A and B , with $PA < PB$. If $PA = 3$ and $PT = AB - PA$, then what is PB ?

Answer:

Step 1: By the Power of a Point, $PT^2 = PA \cdot PB$.

Step 2: Since $AB = PB - PA$, the condition $PT = AB - PA$ gives $PT = (PB - PA) - PA = PB - 2PA$, so with $PA = 3$ we have $PT = PB - 6$.

Step 3: Substitute into the power relation: $(PB - 6)^2 = 3PB$.

Step 4: Expand: $PB^2 - 12PB + 36 = 3PB$, hence $PB^2 - 15PB + 36 = 0$.

Step 5: Solve via the quadratic formula:

$$PB = \frac{15 \pm \sqrt{225 - 144}}{2} = \frac{15 \pm 9}{2},$$

yielding $PB = 12$ or $PB = 3$.

Step 6: Discard $PB = 3$ since $PB > PA = 3$, so $PB = 12$.

The answer is: 12

119

References

- [1] Wenhu Chen, Ming Yin, Max Ku, Pan Lu, Yixin Wan, Xueguang Ma, Jianyu Xu, Xinyi Wang, and Tony Xia. TheoremQA: A theorem-driven question answering dataset. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023.
- [2] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- [3] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- [4] Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun. Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. pages 3828–3850, 2024.
- [5] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [6] Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. Mistral 7b, 2023.
- [7] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [8] Haoxiong Liu, Yifan Zhang, Yifan Luo, and Andrew C Yao. Augmenting math word problems via iterative question composing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 24605–24613, 2025.
- [9] Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, et al. Mixed precision training. In *International Conference on Learning Representations*, 2018.
- [10] Qizhi Pei, Lijun Wu, Zhuoshi Pan, Yu Li, Honglin Lin, Chenlin Ming, Xin Gao, Conghui He, and Rui Yan. Mathfusion: Enhancing mathematic problem-solving of llm through instruction fusion. *arXiv preprint arXiv:2503.16212*, 2025.
- [11] Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 3505–3506, 2020.
- [12] David Saxton, Edward Grefenstette, Felix Hill, and Pushmeet Kohli. Analysing mathematical reasoning abilities of neural models. In *International Conference on Learning Representations*, 2019.
- [13] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- [14] Zhengyang Tang, Xingxing Zhang, Benyou Wang, and Furu Wei. Mathscale: Scaling instruction tuning for mathematical reasoning. In *Forty-first International Conference on Machine Learning*, 2024.
- [15] Yuxuan Tong, Xiwen Zhang, Rui Wang, Ruidong Wu, and Junxian He. DART-math: Difficulty-aware rejection tuning for mathematical problem-solving. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [16] Longhui Yu, Weisen Jiang, Han Shi, Jincheng YU, Zhengying Liu, Yu Zhang, James Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical questions for large language models. In *The Twelfth International Conference on Learning Representations*, 2024.
- [17] Lunjun Zhang, Arian Hosseini, Hritik Bansal, Mehran Kazemi, Aviral Kumar, and Rishabh Agarwal. Generative verifiers: Reward modeling as next-token prediction. In *The Thirteenth International Conference on Learning Representations*, 2025.