

000 001 002 003 004 005 AGENTICMATH: ENHANCING LLM REASONING VIA 006 AGENTIC-BASED MATH DATA GENERATION 007 008 009

010 **Anonymous authors**
011 Paper under double-blind review
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

ABSTRACT

The creation of high-quality datasets to improve Large Language Model (LLM) reasoning remains a significant challenge, as current methods often suffer from generating low-quality/incorrect answers and limited information richness from available data sources. To address this, we propose **AgenticMath**, a novel agentic pipeline for generating high-quality mathematical question-answer pairs to enhance the supervised fine-tuning of LLMs. Our method operates through four stages: (1) *Seed Question Filter* that selects questions with high information richness, complexity, and clarity; (2) an *Agentic Question Rephrase* step that employs a multi-agent system to generate diverse, logically consistent paraphrases; (3) an *Answer Augment* step where rewrite answers using chain-of-thought reasoning to enhance numerical and logical correctness, without reliance on human-provided labels; and (4) a final *Question and Answer Evaluation* that retains only the most superior pairs. Extensive experiments demonstrate that, fine-tuning 3B-8B parameter LLMs on **AgenticMath** generated datasets (comprising only 30-60K math samples) achieves competitive or superior performance on diverse in-domain and out-of-domain mathematical reasoning benchmarks compared to baselines trained on much more data (e.g., 400K or 2.3M samples). Our work demonstrates that targeted, high-quality data generation is a more efficient path to improving mathematical reasoning in LLMs than large-scale, low-quality alternatives.

1 INTRODUCTION

Large language models (LLMs) (Brown et al., 2020; Achiam et al., 2023; Chowdhery et al., 2023; Touvron et al., 2023) have achieved strong results across many domains, showing impressive general reasoning and knowledge transfer. However, when applied to mathematical reasoning, open models (Touvron et al., 2023; Bai et al., 2023; Bi et al., 2024) still perform far below human levels, struggling with the precision and consistency required. Mathematical problems demand long chains of logic that combine symbolic manipulation, cross-domain knowledge, and step-by-step numerical accuracy (Ahn et al., 2024; Long et al., 2024). These requirements make math reasoning more complex and error-prone than typical natural language tasks.

Limitations in Existing Math Reasoning Methods. To improve the mathematical proficiency of LLMs, research has mainly followed two paths. The first uses prompt engineering (Fu et al., 2022), such as Chain-of-Thought (Wei et al., 2022) and Self-Consistency (Wang et al., 2022), which guide models to produce reasoning chains at test time. This method is simple and training-free but its gains are limited by model capacity and often unstable across problem types. The second path relies on powerful base models to synthesize large numbers of question–solution pairs for supervised fine-tuning (SFT) (Yu et al., 2023; Li et al., 2024a; Yue et al., 2023; Gou et al., 2023). This reduces annotation costs and boosts benchmark scores, yet performance is capped by the quality of the synthetic data. When generated problems lack clarity, rigor, or diversity, the resulting models remain far below the performance attainable with expert-annotated corpora. The core challenge is not just producing solutions but enforcing strict quality control during problem synthesis, since the problem statement shapes both the reasoning process and the useful information in the dataset.

Limitations in Multi-Agent Data Generation for Mathematics. Recent work has introduced LLM-based multi-agent frameworks to improve synthetic corpora (Abdullin et al., 2024; Chen et al.,

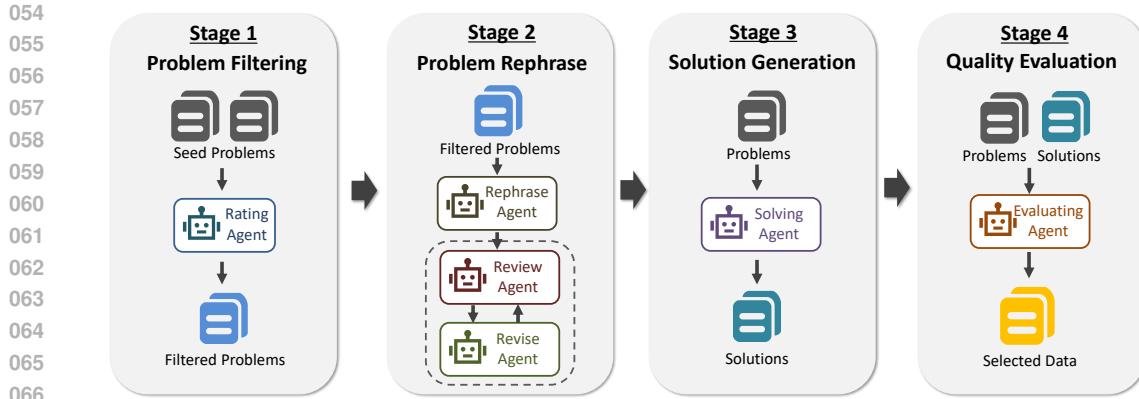


Figure 1: The Overview of AgenticMath Pipeline.

2024a; Mitra et al., 2024a; Ge et al., 2024; Chen et al., 2024b; Ye et al., 2024). Most of these methods target general-purpose instruction data, where task formulation is relatively shallow. In mathematics, the quality of the problem itself is decisive: precise formulation, logical coherence, and sufficient variability not only ensure solvability but also drive the generation of rigorous solutions. Without careful problem design, even advanced solution-generation strategies cannot compensate for poorly posed questions, keeping the dataset far from its upper bound. Existing multi-agent methods rarely enforce such domain-specific constraints, and prior attempts in mathematical data generation (Mitra et al., 2024b; Motwani et al., 2024) still lack systematic control at the problem construction stage.

How AgenticMath Tackles the Challenges. To address these challenges, we propose *AgenticMath*, an automated multi-agent framework that enforces quality control at every stage of mathematical data generation. The framework leverages LLMs for generation, evaluation, and coordinated decision-making. It proceeds in four stages: (1) *Seed filtering* extracts high-value problems from human-authored corpora; (2) *Problem synthesis* engages cooperative agents to rephrase and diversify questions under explicit quality-control criteria; (3) *Solution generation* employs a solver agent to produce complete reasoning chains with rigor and correctness; and (4) *Quality evaluation* aggregates multi-dimensional scores to assess each problem–solution pair. By retaining only top-scoring samples, AgenticMath resolves the data quality bottleneck and follows the “*Less Is More*” principle. The result is a **data-efficient, high-quality dataset** that directly addresses the challenges of clarity, rigor, and diversity in mathematical reasoning tasks.

Empirical Results and Contributions. We evaluate AgenticMath on six mathematical reasoning benchmarks, including in-domain tasks (GSM8K (Cobbe et al., 2021), MATH (Hendrycks et al., 2021)) and out-of-domain settings (CollegeMath (Tang et al., 2024), DeepMind Mathematics (Saxton et al., 2019), OlympiadBench (He et al., 2024), TheoremQA (Chen et al., 2023)). AgenticMath matches or surpasses previous methods that rely on hundreds of thousands or even millions of samples (e.g., 400K or 2.3M), while using far fewer data. With **only 30K–60K** samples, performance improves by over 10 points on average, showing clear data efficiency and strong generalization to out-of-domain tasks. These results establish *AgenticMath* as an efficient and competitive approach to advancing mathematical reasoning. The main contributions of this work are as follows:

- **Agentic Math Data Generation:** We propose *AgenticMath*, an effective multi-agent framework for synthesizing, evaluating, and refining mathematical problems and solutions, offering a systematic and scalable paradigm for building high-quality reasoning corpora.
- **High-Quality Math Data:** We release *AgenticMathQA*, a curated dataset in 30K, 60K, and 90K versions. Unlike prior approaches that rely on scale, our dataset emphasizes clarity, correctness, and diversity, providing higher-quality supervision for mathematical reasoning.
- **Comprehensive Empirical Validation and Insights:** Extensive experiments show that with **only 5%–15%** of the data size, *AgenticMath* matches or even surpasses methods trained on 400K–2M samples. This result demonstrates that data quality, rather than dataset scaling alone, is the main factor behind improvements in mathematical reasoning.

108

2 RELATED WORK

110 **LLM for Math Reasoning.** Large language models (Brown et al., 2020; Achiam et al., 2023;
 111 Touvron et al., 2023; ?; Chowdhery et al., 2023; Bi et al., 2024; Team et al., 2023; 2024; Grattafiori
 112 et al., 2024) show strong general ability and are increasingly applied to mathematical problem solv-
 113 ing (Cobbe et al., 2021; Hendrycks et al., 2021; Zhang et al., 2024a; Xia et al., 2025). Prompt-based
 114 approaches (Wei et al., 2022; Wang et al., 2022; Fu et al., 2022) extend reasoning paths but yield lim-
 115 ited improvements. Recent work thus emphasizes synthesizing math reasoning data for supervised
 116 fine-tuning (Yu et al., 2023; Luo et al., 2023; Tang et al., 2024; Li et al., 2024a; Zhang et al., 2024b;
 117 Liu et al., 2025a; Tong et al., 2024). WizardMath (Luo et al., 2023) adds evolution directives and
 118 reinforcement learning; MathFusion (Pei et al., 2025) fuses problems for relational reasoning. Other
 119 methods integrate code tools (Yue et al., 2023; Wang et al., 2023; Hosseini et al., 2014; Toshniwal
 120 et al., 2024; Li et al., 2024b; Lu et al., 2024). In this work, we advance mathematical reasoning by
 121 improving both the question formulation and answer quality in synthetic data.

122 **Multi-Agent for Data Generation** Multi-agent systems (Hong et al., 2023; Wu et al., 2023; Li
 123 et al., 2023; aut, 2023) show strong ability and are increasingly applied to data synthesis. Abdullin
 124 et al. (2024) proposed a multi-intelligence framework for dialog generation, while MAGDi (Chen
 125 et al., 2024a) used graph-based interactions and MALLM-GAN (Ling et al., 2024) employed gener-
 126 ator–discriminator agents for tabular data. AgentInstruct (Mitra et al., 2024a) and Orca-Math (Mitra
 127 et al., 2024b) iteratively refined instructions, whereas role-driven synthesis was explored by Ge et al.
 128 (2024) and VCR (Liu et al., 2025b). MALT (Motwani et al., 2024) introduced generator, verifier,
 129 and refiner agents for math problems. Despite these advances, ensuring high-quality data for math-
 130 ematical reasoning remains challenging. To address this, we introduce seed filtering and quality
 131 evaluate agents to guarantee reliable math reasoning data.

132

3 AGENTICMATH: MULTI-AGENT DESIGN FOR MATH REASONING

133 This section details the proposed AgenticMath (see Figure 1), which is designed to generate high-
 134 quality mathematical problems and reasoning solutions based on the GSM8K (Cobbe et al., 2021)
 135 and MATH (Hendrycks et al., 2021) seed datasets. The framework consists of four stages: seed
 136 problem filtering, agentic problem generation, reasoning-solution generation, and synthetic-data
 137 evaluation. Using AgenticMath, we construct a high-quality math dataset to fine-tune LLMs and
 138 enhance their mathematical reasoning ability. All agent prompts are provided in Appendix A.8.

139

3.1 PROBLEM DEFINITION

140 Given a seed dataset $\mathcal{D}_{\text{seed}} = \{q_i\}_{i=1}^N$, where each q_i denotes an original mathematical problem from
 141 MATH (Hendrycks et al., 2021) and GSM8K (Cobbe et al., 2021), we employ large language models
 142 (LLMs) to construct a new dataset of problem–solution pairs, eliminating ground-truth labels and
 143 reducing costly human annotations. Formally, the transformation can be summarized as $\mathcal{D}_{\text{seed}} \Rightarrow$
 144 $\mathcal{D}_{\text{final}}$, where the resulting dataset is denoted as $\mathcal{D}_{\text{final}} = \{(q_i, r_i)\}_{i=1}^{N'}$. The problem component q_i
 145 consists of both original problems from $\mathcal{D}_{\text{seed}}$ and newly synthesized problems, while the solution
 146 component r_i is entirely generated by the LLM. This dataset $\mathcal{D}_{\text{final}}$ is subsequently used as training
 147 data for supervised fine-tuning (SFT). The SFT objective is to maximize likelihood of the target
 148 response given the prompt query. Specifically, the loss function is defined as:

$$149 \mathcal{L}(\theta) = -\frac{1}{N'} \sum_{i=1}^{N'} \log P(r_i | q_i; \theta), \quad (1)$$

150 where θ denotes model parameters, q_i the input problem, and r_i the generated solution.

151

3.2 STAGE 1: SEED PROBLEM FILTERING

152 Using seed problems as references to synthesize more problems can effectively enhance the model’s
 153 mathematical capabilities. However, current methods ignore that low-quality, low-difficulty prob-
 154 lems in the seed dataset may not be worthy of serving as references. Hence, we propose a training-
 155 free and label-free filtering method to identify high-value seed problems.

162
 163 **LLM-based Scoring.** Each candidate
 164 problem from the seed dataset $\mathcal{D}_{\text{seed}} =$
 165 $\{q_i^{\text{seed}}\}_{i=1}^N$ is scored by a large language
 166 model (LLM) along three dimensions:
 167 *Complexity* (c_i), *Information Value* (v_i),
 168 and *Clarity & Precision* (p_i), with
 169 ratings in the range $\{0, 1, \dots, 5\}$. The Evaluator,
 170 based on GPT-4o-mini (Achiam et al., 2023),
 171 processes a scoring prompt to generate the score list $\mathbf{s}_i = [c_i, v_i, p_i]$.
 172 The overall score \bar{s}_i is defined as the
 173 arithmetic mean of these three dimensions.
 174 As a result of this evaluation, we
 175 obtain the scored dataset $\mathcal{D}_{\text{scored}} =$
 176 $\{(q_i^{\text{seed}}, \mathbf{s}_i, \bar{s}_i)\}_{i=1}^N$, where each problem q_i
 177 is associated with both its dimension-wise scores \mathbf{s}_i and aggregated score \bar{s}_i .
 178

Problem Filtering Dimensions and Prompt Example

179 **Complexity:** Does it integrate multiple mathematical domains (e.g., algebra + geometry) or demand critical thinking?
 180

181 **Information Value:** Does it contain useful knowledge or reasoning opportunities? Can it help learners discover concepts, strategies, or patterns?
 182

183 **Clarity & Precision:** Is the question unambiguous, logically consistent, and free of errors?
 184 Poorly framed questions score lower.
 185

186 **Score Curation.** To mitigate potential rating errors introduced by LLM-based evaluation, we apply a score curation procedure inspired by DS2 (Pang et al., 2024) and the clusterability-based method of (Zhu et al., 2021). Starting from the scored dataset $\mathcal{D}_{\text{scored}}$, we construct a Score Transition Matrix T to capture consistency patterns among neighboring problems in the embedding space. By leveraging k -nearest neighbor agreement, problems whose ratings deviate from those of their local neighborhood are adjusted toward more reliable estimates. This process yields the curated dataset $\mathcal{D}_{\text{curated}} = \{(q_i^{\text{seed}}, \tilde{s}_i)\}_{i=1}^N$, where each problem q_i^{seed} is paired with its corrected overall score \tilde{s}_i , representing a refined estimate of problem quality.
 187
 188
 189
 190
 191
 192
 193

194 **Filtering Rule.** In the final step, we impose a quality threshold of $\tau = 3$ on the curated scores.
 195 The resulting dataset $\mathcal{D}_{\text{filter}}$ is derived from $\mathcal{D}_{\text{curated}}$ by retaining only those problems whose corrected
 196 overall score \tilde{s}_i meets or exceeds this threshold. This filtering process excludes problems that are
 197 overly simplistic, ambiguous, or uninformative, ensuring that the retained problems are well-formed
 198 and valuable. The overall pipeline for seed problem filtering can be summarized as: $\mathcal{D}_{\text{seed}} \Rightarrow$
 199 $\mathcal{D}_{\text{scored}} \Rightarrow \mathcal{D}_{\text{curated}} \Rightarrow \mathcal{D}_{\text{filter}}$.
 200

201 3.3 STAGE 2: AGENTIC PROBLEM GENERATION

202 Although closed-source models can generate complex new problems by following
 203 instructions, hallucinations still appear, leading to low-quality or poorly
 204 phrased outputs. In multi-agent settings, self-reflection provides a way to correct
 205 such errors. Building on this idea, we design a framework for problem synthesis
 206 that ensures quality through three roles: a rephrase agent, a review agent, and a
 207 revise agent.
 208
 209
 210
 211
 212
 213

214 **Problem Rephrase Agent.** From the filtered dataset $\mathcal{D}_{\text{filter}} = \{q_i^{\text{seed}}\}_{i=1}^M$, each
 215

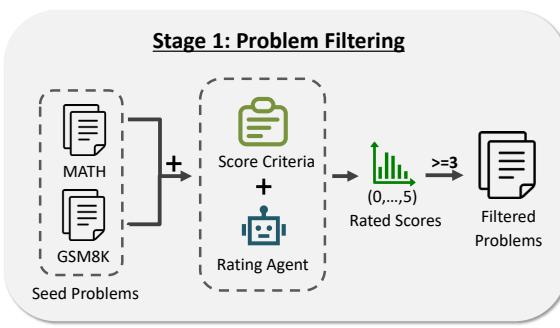


Figure 2: Workflow of Stage 1 showing the filtering process that removes low-quality seed problems and retains high-value ones for subsequent synthesis.

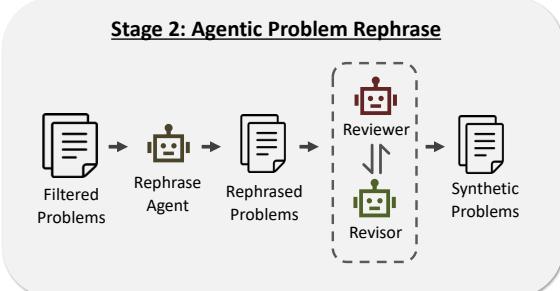


Figure 3: Workflow of Stage 2 showing multi-agent problem generation with Rephrase, Review, and Revise agents, along with iterative refinement to ensure clarity, coherence, and mathematical validity.

216 problem is expanded into six paraphrased variants by the Problem Rephrase Agent. The new col-
 217 lection is denoted as $\mathcal{D}_{\text{rephrase}} = \{q_i^{\text{rep}}\}_{i=1}^{M'}$, where each q_i^{rep} corresponds to a rephrased version of
 218 its seed problem. Rephrasing is guided by task-specific instructions to GPT-4o-mini, designed to
 219 preserve the mathematical intent while introducing greater difficulty, lexical richness, and syntactic
 220 diversity.
 221

222 **Problem Review Agent.** The rephrased dataset $\mathcal{D}_{\text{rephrase}} = \{q_i^{\text{rep}}\}_{i=1}^{M'}$ is passed to the Problem Re-
 223 view Agent for evaluation. Each candidate problem is checked against its original version following
 224 a review instruction. The assessment spans three dimensions: *Clarity & Grammar*, *Logical Co-
 225 herence & Completeness*, and *Mathematical Validity & Solvability*. For every candidate, the agent
 226 assigns a score in the range [1, 5] and, if needed, provides textual feedback for improvement. The
 227 outcome is the reviewed dataset $\mathcal{D}_{\text{review}} = \{(q_i^{\text{rep}}, \bar{s}_i^{\text{rev}}, a_i^{\text{rev}})\}_{i=1}^{M'}$, where each rephrased problem is
 228 paired with its score and an optional suggestion.
 229

Problem Review Dimensions and Prompt Example

231 **Clarity & Grammar:** The question must be grammatically correct, precisely phrased, and easy
 232 to understand. It should avoid ambiguity in wording or phrasing.
 233

234 **Logical Coherence & Completeness:** All elements of the problem (e.g., given information, con-
 235 straints, relationships, objectives) must be logically interconnected and sufficient. The problem
 236 should present a clear, sequential path for reasoning, without missing information required for
 237 the specified solution approach.
 238

239 **Mathematical Validity & Solvability:** The problem must be fundamentally a mathematics prob-
 240 lem, with all its premises and conditions being mutually consistent and mathematically sound...
 241

242 **Problem Revise Agent.** Based on the reviewed dataset $\mathcal{D}_{\text{review}}$, the Problem Revise Agent targets
 243 rephrased problems with scores below the threshold $\tau_{\text{rev}} = 4.5$. For each low-scoring case, the
 244 problem q_i^{rep} is revised according to reviewer feedback a_i^{rev} . This step corrects issues such as unclear
 245 phrasing, weak logical flow, or invalid mathematical conditions. The result is the revised dataset
 246 $\mathcal{D}_{\text{revise}} = \{q_i^{\text{rev}}\}_{i=1}^{M''}$, which retains only problems that reach the required quality level.
 247

248 **Problem Review–Revise Interaction.** To further strengthen problem quality, an iterative loop
 249 between the Review and Revise agents is applied. Starting from $\mathcal{D}_{\text{review}}$, all problems scoring below
 250 τ_{rev} enter this refinement process. In each round, the Review agent re-evaluates a candidate, assigns
 251 a new score, and may suggest specific improvements. The Revise Agent incorporates this feedback
 252 to produce an updated version. The loop runs for at most three iterations, with early stopping once
 253 the threshold is met. Afterward, only problems with final scores above 4.5 are kept, while the
 254 rest are discarded. The outcome is the refined dataset $\mathcal{D}_{\text{refined}} = \{q_i^{\text{ref}}\}_{i=1}^K$, containing high-quality
 255 rephrasings that meet the required standard.
 256

3.4 STAGE 3: SOLUTION GENERATION

257 **Solution Solver Agent.** To provide
 258 high-quality reasoning traces for training,
 259 we employ a one-shot Chain-of-Thought
 260 (CoT) prompting scheme that elicits
 261 multi-step reasoning solution paths. The
 262 Solver Agent works on two distinct
 263 datasets: the original seed problems
 264 $\mathcal{D}_{\text{seed}} = \{q_i^{\text{seed}}\}_{i=1}^N$ and the refined
 265 rephrased problems $\mathcal{D}_{\text{ref}} = \{q_j^{\text{ref}}\}_{j=1}^K$. For
 266 each problem, GPT-4o-mini is prompted
 267 with a single CoT exemplar to generate
 268 a detailed, step-by-step solution. This
 269 process produces two corresponding
 270 solution-augmented datasets:
 271

$$\mathcal{D}_{\text{seed}}^{\text{sol}} = \{(q_i^{\text{seed}}, a_i^{\text{sol}})\}_{i=1}^N, \quad \mathcal{D}_{\text{ref}}^{\text{sol}} = \{(q_j^{\text{ref}}, a_j^{\text{sol}})\}_{j=1}^K,$$

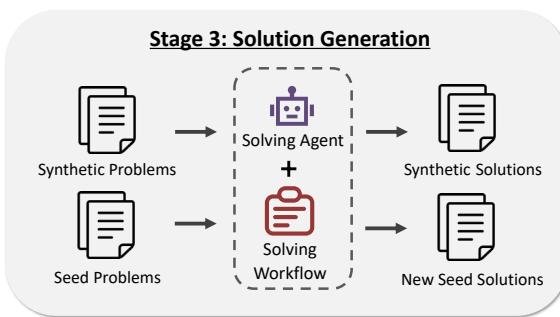


Figure 4: Workflow of Stage 3 showing solution generation by the Solver Agent with step-by-step reasoning.

270 where every problem from $\mathcal{D}^{\text{seed}}$ and \mathcal{D}^{ref} is paired with a synthetic solution a^{sol} that explicitly
 271 demonstrates the intermediate reasoning steps.
 272

273 3.5 STAGE 4: SYNTHETIC DATA EVALUATION

274 In this stage, the scoring and curation
 275 framework from Stage 1 is extended
 276 to problem–solution pairs. The evalua-
 277 tion targets the synthetic set $\mathcal{D}_{\text{ref}}^{\text{sol}} =$
 278 $\{(q_j^{\text{ref}}, a_j^{\text{sol}})\}_{j=1}^K$. Each pair is judged along
 279 three dimensions—clarity of the problem,
 280 correctness of the solution, and comple-
 281 teness of reasoning. Scores are further
 282 stabilized using the Score Transition
 283 Matrix and refined through k -NN consistency
 284 checks. To build a high-quality and
 285 diverse subset, a ranking-based selection is
 286 applied instead of a fixed threshold. Pairs
 287 are first sorted by quality and grouped into
 288 discrete score levels. We first group all samples by quality score (from 5 down to 0) and select
 289 groups in descending order. When the remaining quota falls within a group larger than needed, we
 290 rank samples inside that group using the long-tail diversity score. This strategy ensures that we
 291 always take the highest-quality data available while promoting diversity when selecting from over-
 292 sized groups. This yields the curated dataset $\mathcal{D}_{\text{selected}}^{\text{sol}} = \{(q_j^{\text{ref}}, a_j^{\text{sol}})\}_{j=1}^L$. The final training dataset
 293 combines this curated set with the seed-based solutions: $\mathcal{D}_{\text{final}} = \mathcal{D}_{\text{selected}}^{\text{sol}} \cup \mathcal{D}_{\text{seed}}^{\text{sol}}$, ensuring both
 294 rigor and diversity for downstream fine-tuning.
 295

296 4 EXPERIMENTS

297 4.1 EXPERIMENTAL SETUP

300 **Data Synthesis:** We employed GPT-4o-mini (2024-07-18) (Achiam et al., 2023), following (Pei
 301 et al., 2025), for all agents across the four stages, including evaluation scoring, problem synthesis,
 302 and solution synthesis. Seed problems were sourced from the MATH (Hendrycks et al., 2021) and
 303 GSM8K (Cobbe et al., 2021) datasets. For the 30K setting, the final dataset consists of 15K seed
 304 problems and 15K AgenticMath-synthesized problems. In Stage 1, we filtered seed problems with
 305 scores above 3. In Stage 2, each seed problem was expanded into six rephrased variants, with a
 306 review–revise loop requiring scores above 4.5 and running up to three iterations, keeping only those
 307 exceeding the threshold. In Stage 4, we applied ranking-based selection with a target of 15K high-
 308 quality problem–solution pairs. During all data generation steps, we used a temperature of 0.7 and
 309 a maximum token length of 4096.

310 **Training:** We perform standard instruction tuning on the proposed AgenticMathQA. Following
 311 DART-Math (Tong et al., 2024) and Mathfusion (Pei et al., 2025), experiments cover both math-
 312 specialized and general base models. For the math-specialized category, we use DeepSeekMath-
 313 7B (sha, 2024); for general models, we fine-tune Qwen2.5-3B (Team, 2024), Mistral-7B (Jiang et al.,
 314 2023), and Llama3-8B (Grattafiori et al., 2024). The 30K dataset is built from 15K seed problems
 315 (sourced from GSM8K and MATH) with corresponding solutions, together with 15K AgenticMath-
 316 synthesized problem–solution pairs. Scaling to larger sizes is achieved by augmenting each 30K
 317 problem with multiple solutions: duplicating once yields 60K (30K \times 2), and duplicating twice yields
 318 90K (30K \times 3). More training details are provided in Appendix A.2.

319 **Evaluation:** Following DART-Math (Tong et al., 2024) and MathFusion (Pei et al., 2025), we
 320 evaluate on six benchmarks spanning both in-domain and out-of-domain (OOD) settings. The in-
 321 domain benchmarks are GSM8K (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021), while
 322 the OOD benchmarks include CollegeMath (Tang et al., 2024), DeepMind-Mathematics (Saxton
 323 et al., 2019), OlympiadBench-Math (He et al., 2024), and TheoremQA (Chen et al., 2023). Further
 324 benchmark details are provided in the Appendix A.3.

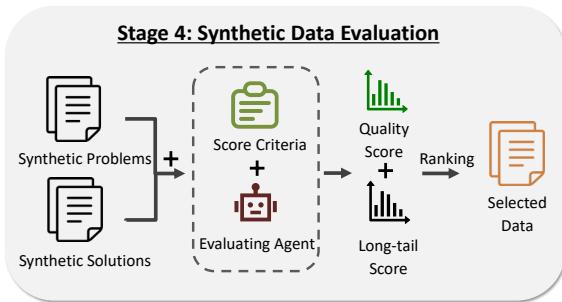


Figure 5: Workflow of Stage 4 showing evaluation of problem–solution pairs for quality and diversity.

324	325	Model	# Samples	In-Domain		Out-of-Domain			AVG
				MATH	GSM8K	College	DM	Olympiad	
Qwen2.5-3B (3-8B General Base Model)									
327	Qwen2.5-3B-RefAug [†]	30K	40.9	69.7	32.4	42.5	10.7	11.4	34.6
328	Qwen2.5-3B-MathFusion (<i>Sequential</i>) [†]	30K	39.9	72.1	28.9	50.0	23.0	14.6	38.1
329	AgenticMath-Qwen2.5-3B	30K	62.0	83.4	46.8	72.8	25.6	31.4	53.7
330	Qwen2.5-3B-MetaMath [†]	60K	43.4	79.8	34.5	46.3	11.3	19.0	39.1
331	Qwen2.5-3B-MMIQC [†]	60K	47.3	78.2	35.6	51.2	14.7	17.1	40.7
332	Qwen2.5-3B-DART-Math [†]	60K	53.9	84.3	42.3	59.2	18.4	26.4	47.4
333	Qwen2.5-3B-MathFusion [†]	60K	40.5	72.7	29.1	52.4	25.5	15.3	39.3
334	AgenticMath-Qwen2.5-3B	60K	62.4	83.6	46.3	74.3	27.3	29.3	53.9
DeepSeekMath (7B Math-Specialized Base Model)									
335	DeepSeekMath-7B-RefAug	30K	32.1	71.2	26.0	38.4	10.1	14.4	32.0
336	DeepSeekMath-7B-MathFusion (<i>Sequential</i>)	30K	49.9	76.6	38.8	64.6	21.6	22.8	45.7
337	AgenticMath-DSMath-7B	30K	52.4	80.1	42.6	66.8	18.2	26.9	47.8
338	DeepSeekMath-7B-MetaMath	60K	40.0	79.0	33.2	45.9	9.5	18.9	37.8
339	DeepSeekMath-7B-MMIQC	60K	26.3	60.6	19.2	41.5	10.4	6.8	27.5
340	DeepSeekMath-7B-RefAug	60K	33.1	71.6	26.2	35.4	10.5	14.0	31.8
341	DeepSeekMath-7B-DART-Math	60K	51.4	82.9	39.1	62.8	21.0	27.4	47.4
342	DeepSeekMath-7B-MathFusion	60K	53.4	77.9	39.8	65.8	23.3	24.6	47.5
343	AgenticMath-DSMath-7B	60K	55.0	80.1	43.6	69.9	20.0	27.0	49.3
Mistral-7B (3-8B General Base Model)									
344	Mistral-7B-RefAug	30K	15.1	61.1	10.4	15.4	3.1	11.0	19.4
345	Mistral-7B-MathFusion (<i>Sequential</i>)	30K	32.7	73.9	18.9	29.3	9.3	15.5	29.9
346	AgenticMath-Mistral-7B	30K	35.3	79.5	27.0	41.9	11.9	19.3	35.8
347	Mistral-7B-MetaMath	60K	22.7	70.8	14.1	27.2	5.0	12.2	25.3
348	Mistral-7B-MMIQC	60K	17.3	61.4	11.1	13.5	5.0	5.9	19.0
349	Mistral-7B-RefAug	60K	17.4	63.1	12.5	18.1	3.9	11.1	21.0
350	Mistral-7B-DART-Math	60K	34.1	77.2	23.4	36.0	8.7	18.2	32.9
351	Mistral-7B-MathFusion	60K	41.6	79.8	24.3	39.2	13.6	18.1	36.1
352	AgenticMath-Mistral-7B	60K	39.5	82.3	28.7	47.1	12.4	20.5	38.4
Llama3-8B (3-8B General Base Model)									
353	Llama3-8B-RefAug	30K	20.8	67.3	15.7	25.9	4.7	13.6	24.7
354	Llama3-8B-MathFusion (<i>Sequential</i>)	30K	38.8	77.9	25.1	42.0	12.6	17.0	35.6
355	AgenticMath-Llama3-8B	30K	36.8	78.4	29.6	40.3	11.4	20.4	36.2
356	Llama3-8B-MetaMath	60K	28.7	78.5	19.7	31.3	5.3	16.1	29.9
357	Llama3-8B-MMIQC	60K	24.4	69.7	13.4	30.9	5.2	10.6	25.7
358	Llama3-8B-RefAug	60K	20.3	68.6	15.5	29.1	5.5	13.0	25.3
359	Llama3-8B-DART-Math	60K	39.6	82.2	27.9	36.9	12.9	22.9	37.6
360	Llama3-8B-MathFusion	60K	46.5	79.2	27.9	43.4	17.2	20.0	39.0
361	AgenticMath-Llama3-8B	60K	40.4	80.1	31.6	46.7	14.1	22.6	39.3

Table 1: Evaluation results across in-domain and out-of-domain math benchmarks with 30K–60K samples. Most baseline results are reported from (Pei et al., 2025), while entries marked with [†] denote results reproduced by us. **Bold** numbers indicate the best performance within the same type of sample size and base model. Rows highlighted in blue correspond to our AgenticMath results.

Baseline: We compare AgenticMath with state-of-the-art methods in two settings. For large-scale training, we include MetaMath (Yu et al., 2023), RFT (Yuan et al., 2023), DART-Math (Tong et al., 2024), MathScale (Tang et al., 2024), DeepSeekMath-7B-Instruct (sha, 2024), RefAug (Zhang et al., 2024b), MMIQC (Liu et al., 2025a), and WizardMath (Luo et al., 2023), all using 400K–2.3M samples. For small-scale, we evaluate 30K and 60K subsets: RefAug and MathFusion provide native 30K versions, while other baselines are randomly down-sampled to 60K from the large-scale dataset above.

4.2 MAIN RESULTS

AgenticMath Achieves SOTA Performance at 30K–60K Data Scale. Table 1 shows that AgenticMath consistently outperforms all baselines at both 30K and 60K scales. Across every base model (Qwen2.5-3B, DeepSeekMath-7B, Mistral-7B, and Llama3-8B), our method achieves the highest average score and sets new state-of-the-art performance under small-scale training. For example, with 30K samples, AgenticMath-Qwen2.5-3B reaches 53.7 average accuracy, surpassing MathFusion by over 15 points. At 60K, AgenticMath continues to improve and outperforms all other baseline methods trained with the same number of samples. These results demonstrate that rigorous multi-agent synthesis and quality control provide significantly better data efficiency than prior methods.

378	379	380	381	382	383	384	385	386	Model	# Samples	In-Domain		Out-of-Domain			AVG
									MATH	GSM8K	College	DM	Olympiad	Theorem		
									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-7B-MathFusion-7B	195K	58.2	79.5	40.3	69.1	25.5	27.0	49.9								
AgenticMath-DSMath-7B	30K	52.4	80.1	42.6	66.8	18.2	26.9	47.8								
AgenticMath-DSMath-7B	60K	55.0	80.1	43.6	69.9	20.0	27.0	49.3								
Mistral-7B (3-8B 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-7B-MathScale	2.0M	35.2	74.8	21.8	—	—	—	—								
Mistral-7B-MMIQC	2.3M	37.4	75.4	28.5	38.0	9.4	16.2	34.2								
AgenticMath-Mistral-7B	30K	35.3	79.5	27.0	41.9	11.9	19.3	35.8								
AgenticMath-Mistral-7B	60K	39.5	82.3	28.7	47.1	12.4	20.5	38.4								
Llama3-8B (3-8B 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-MMIQC	2.3M	39.5	77.6	29.5	41.0	9.6	16.2	35.6								
Llama3-8B-DART-Math	590K	46.6	81.1	28.8	48.0	14.5	19.4	39.7								
AgenticMath-Llama3-8B	30K	36.8	78.4	29.6	40.3	11.4	20.4	36.2								
AgenticMath-Llama3-8B	60K	40.4	80.1	31.6	46.7	14.1	22.6	39.3								
AgenticMath-Llama3-8B	90K	42.8	81.4	33.0	48.6	13.9	21.8	40.3								

Table 2: Results on math benchmarks comparing AgenticMath (30K/60K/90K) with large-scale baselines trained on 400K–2.3M data. All baseline results are reported from their respective papers. **Bold** numbers indicate the best performance within the same type of base model. Rows highlighted in blue correspond to our AgenticMath results.

AgenticMath Matches or Surpasses Larger-Scale Baselines with Much Less Data. Table 2 shows that AgenticMath, even with only 30K–90K samples, matches or surpasses baselines trained on hundreds of thousands or even millions of samples. For example, AgenticMath-DSMath-7B (60K) achieves an average score of 49.3, close to DeepSeekMath-7B-MathFusion (195K, 49.9) and higher than DeepSeekMath-7B-RFT (590K, 48.3). On general models, AgenticMath-Mistral-7B (60K) reaches 38.4, comparable to Mistral-7B-DART-Math (590K, 38.8) and outperforming Mistral-7B-RFT (590K, 34.3). Most notably, AgenticMath-Llama3-8B (90K) achieves 40.3, surpassing Llama3-8B-DART-Math (590K, 39.7) and all other large-scale Llama3 baselines. These results confirm that AgenticMath delivers competitive or superior performance with much fewer samples, highlighting its strong data efficiency.

4.3 UNDERSTANDING AGENTICMATH: ABLATIONS AND INSIGHTS

All Modules Contribute to Performance Gains.

To more clearly quantify the contribution of each module in the AgenticMath pipeline, we perform ablation studies on Mistral-7B using a fixed training set of 15K synthesized samples. As shown in Table 3, the pipeline exhibits consistent, stage-by-stage improvements. The initial Problem Rephrase step provides a strong baseline (31.4 AVG). Introducing Seed Filtering further improves performance by removing low-quality or noisy seed questions, yielding a +0.6 gain. The subsequent Problem Review–Revise loop brings the largest additional improvement (+1.0), demonstrating the importance of iterative refinement for producing clearer and more logically coherent problems. Finally, the Synthetic Data Evaluation stage contributes a modest but steady boost (+0.2), as most low-quality samples have already been filtered out in earlier stages. Overall, these results show that the full pipeline is cumulatively beneficial, and that each stage plays a meaningful role in improving data quality and downstream model performance.

Table 3: Ablation study on the contribution of different pipeline stages.

Method	Samples	AVG
Problem Rephrase	15K	31.4
+ Seed Filtering	15K	32.0 $\uparrow 0.6$
+ Problem Review–Revise	15K	33.0 $\uparrow 1.0$
+ Synthetic Data Evaluation	15K	33.2 $\uparrow 0.2$

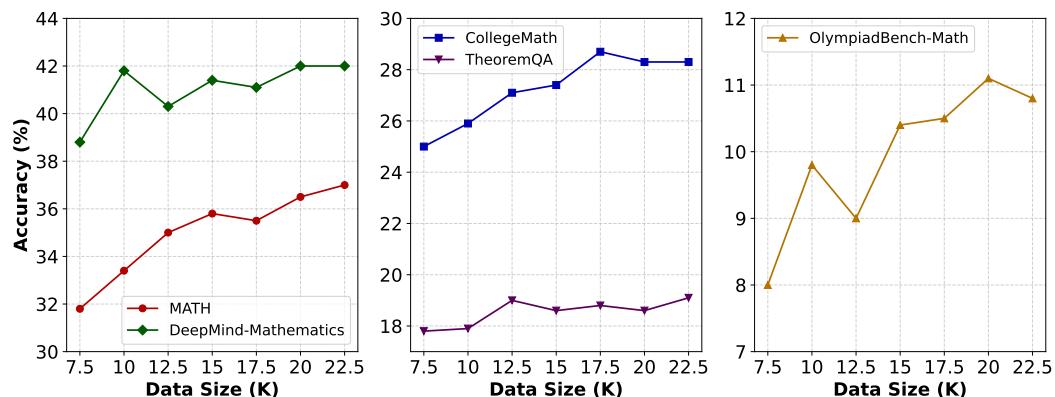


Figure 6: Llama3-8B performance across benchmarks as training size increases.

Problem Quality Directly Boosts Performance. We further investigate the impact of different filtering thresholds in Stage 1 using Mistral-7B as the base model. Table 4 shows that the use of higher thresholds for the filtering of seed problems leads to better results, confirming that the quality of the selected problems directly impacts the performance of reasoning. Although our main experiments adopt a threshold of score 3, increasing it to score 4 yields further gains. This indicates that *AgenticMath* benefits from stricter quality control and still offers further optimization space for even stronger performance.

Larger Training Samples Yield Stronger Reasoning Performance. We analyze how varying the amount of training data affects model performance. Starting from a base setting with 7.5K MATH samples, we gradually add synthetic data in increments of 2.5K, up to a total of 22.5K samples. As shown in Figure 6, Llama3-8B shows consistent accuracy gains on different benchmarks as the dataset grows, confirming a strong positive correlation between training size and reasoning ability. This upward trend demonstrates that increasing data with our multi-agent framework steadily strengthens performance.

Illustrative Cases of Enhanced Problem Quality Appendix A.7 provides several illustrative cases refined by the Reviewer and Revise Agents, showing how our method improves clarity and correctness of mathematical problems.

5 CONCLUSION

In this work, we introduced *AgenticMath*, a multi-agent framework for high-quality synthetic data generation of mathematical problems and solutions. By coordinating agents for filtering, rephrasing, revision, solution generation, and joint evaluation, *AgenticMath* provides a systematic and scalable approach to generating high-quality math reasoning data. The resulting dataset, *AgenticMathQA*, is released in curated 30K, 60K, and 90K versions, emphasizing clarity, correctness, and diversity rather than data scale. Extensive experiments across multiple open-source base models show that with only 5%–15% of the data size scale, *AgenticMath* matches or surpasses methods trained on 400K–2.3M samples, achieving SOTA performances by referring to baselines with the same data scale. These results highlight that data quality—supported by rigorous multi-agent design—plays a more decisive role than dataset size in advancing mathematical reasoning in large language models.

Table 4: Performance with different thresholds for seed problem filtering.

Threshold	Samples	AVG
Score = 2	30K	33.4
Score = 3	30K	34.9
Score = 4	30K	35.0

486
487**ETHICS STATEMENT**

488

This work does not involve human subjects or sensitive personal data. All datasets used are publicly available (MATH and GSM8K), and our synthetic data generation follows the ICLR Code of Ethics by avoiding the release of harmful or biased content. The proposed framework, AgenticMath, is intended purely for advancing research in mathematical reasoning and does not introduce applications that could cause societal harm. We release our curated datasets and code in compliance with licensing terms to ensure transparency, reproducibility, and fair access for the research community.

489

490

REPRODUCIBILITY STATEMENT

491

492

We have taken multiple steps to ensure reproducibility. Details of the multi-agent pipeline design, training and evaluation are provided in the main text. Additional training details, hyperparameters are included in the Appendix. All benchmarks used (MATH, GSM8K, CollegeMath, DeepMind-Mathematics, OlympiadBench, and TheoremQA) are publicly available. The full AgenticMathQA dataset and implementation code will be released upon publication to enable independent verification of our results.

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 REFERENCES
541

542 Autogpt: An open-source autonomous ai agent. <https://github.com/Significant-Gravitas/AutoGPT>, 2023.

544 Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

546 Yelaman Abdullin, Diego Molla-Aliod, Bahadorreza Ofoghi, John Yearwood, and Qingyang Li. Synthetic dialogue dataset generation using llm agents. *arXiv preprint arXiv:2401.17461*, 2024.

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

553 Janice Ahn, Rishu Verma, Renze Lou, Di Liu, Rui Zhang, and Wenpeng Yin. Large language models for mathematical reasoning: Progresses and challenges. *arXiv preprint arXiv:2402.00157*, 2024.

555 Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023.

558 Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiushi Du, Zhe Fu, et al. Deepseek llm: Scaling open-source language models with longtermism. *arXiv preprint arXiv:2401.02954*, 2024.

561 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

565 Justin Chih-Yao Chen, Swarnadeep Saha, Elias Stengel-Eskin, and Mohit Bansal. Magdi: Structured distillation of multi-agent interaction graphs improves reasoning in smaller language models. *arXiv preprint arXiv:2402.01620*, 2024a.

568 Mingyang Chen, Haoze Sun, Tianpeng Li, Fan Yang, Hao Liang, Keer Lu, Bin Cui, Wentao Zhang, Zenan Zhou, and Weipeng Chen. Facilitating multi-turn function calling for llms via compositional instruction tuning. *arXiv preprint arXiv:2410.12952*, 2024b.

572 Wenhui 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.

575 Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113, 2023.

580 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.

583 Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. Complexity-based prompting for multi-step reasoning. *arXiv preprint arXiv:2210.00720*, 2022.

586 Tao Ge, Xin Chan, Xiaoyang Wang, Dian Yu, Haitao Mi, and Dong Yu. Scaling synthetic data creation with 1,000,000,000 personas. *arXiv preprint arXiv:2406.20094*, 2024.

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

592 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.

594 Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Thai, Junhao Shen, Jinyi Hu, Xu Han,
 595 Yujie Huang, Yuxiang Zhang, et al. Olympiadbench: A challenging benchmark for promoting
 596 agi with olympiad-level bilingual multimodal scientific problems. In *Proceedings of the 62nd*
 597 *Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp.
 598 3828–3850, 2024.

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

602 Sirui Hong, Xiawu Zheng, Jonathan Chen, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang,
 603 Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, et al. Metagpt: Meta programming for multi-
 604 agent collaborative framework. *arXiv preprint arXiv:2308.00352*, 3(4):6, 2023.

605 Mohammad Javad Hosseini, Hannaneh Hajishirzi, Oren Etzioni, and Nate Kushman. Learning to
 606 solve arithmetic word problems with verb categorization. In *Proceedings of the 2014 conference*
 607 *on empirical methods in natural language processing (EMNLP)*, pp. 523–533, 2014.

608 Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chap-
 609 lot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier,
 610 Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas
 611 Wang, Timothée Lacroix, and William El Sayed. Mistral 7b. *arXiv preprint arXiv:2310.06825*,
 612 2023. URL <https://arxiv.org/abs/2310.06825>.

613 Chen Li, Weiqi Wang, Jingcheng Hu, Yixuan Wei, Nanning Zheng, Han Hu, Zheng Zhang, and
 614 Houwen Peng. Common 7b language models already possess strong math capabilities. *arXiv*
 615 *preprint arXiv:2403.04706*, 2024a.

616 Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbulin, and Bernard Ghanem. Camel: Com-
 617 municative agents for “mind” exploration of large language model society. *Advances in Neural*
 618 *Information Processing Systems*, 36:51991–52008, 2023.

619 Jia Li, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Huang, Kashif
 620 Rasul, Longhui Yu, Albert Q Jiang, Ziju Shen, et al. Numinamath: The largest public dataset in
 621 ai4maths with 860k pairs of competition math problems and solutions. *Hugging Face repository*,
 622 13:9, 2024b.

623 Yaobin Ling, Xiaoqian Jiang, and Yejin Kim. Mallm-gan: Multi-agent large language model as
 624 generative adversarial network for synthesizing tabular data. *arXiv preprint arXiv:2406.10521*,
 625 2024.

626 Haoxiong Liu, Yifan Zhang, Yifan Luo, and Andrew C Yao. Augmenting math word problems via
 627 iterative question composing. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
 628 volume 39, pp. 24605–24613, 2025a.

629 Sannyyua Liu, Jintian Feng, Xiaoxuan Shen, Shengyingjie Liu, Qian Wan, and Jianwen Sun. Vcr:
 630 A “cone of experience” driven synthetic data generation framework for mathematical reasoning.
 631 In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 24650–24658,
 632 2025b.

633 Lin Long, Rui Wang, Ruixuan Xiao, Junbo Zhao, Xiao Ding, Gang Chen, and Haobo Wang.
 634 On llms-driven synthetic data generation, curation, and evaluation: A survey. *arXiv preprint*
 635 *arXiv:2406.15126*, 2024.

636 Zimu Lu, Aojun Zhou, Houxing Ren, Ke Wang, Weikang Shi, Junting Pan, Mingjie Zhan, and
 637 Hongsheng Li. Mathgenie: Generating synthetic data with question back-translation for enhanc-
 638 ing mathematical reasoning of llms. *arXiv preprint arXiv:2402.16352*, 2024.

639 Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng, Qing-
 640 wei Lin, Shifeng Chen, and Dongmei Zhang. Wizardmath: Empowering mathematical reasoning
 641 for large language models via reinforced evol-instruct. *arXiv preprint arXiv:2308.09583*, 2023.

648 Arindam Mitra, Luciano Del Corro, Guoqing Zheng, Shweta Mahajan, Dany Rouhana, Andres Co-
 649 das, Yadong Lu, Wei-ge Chen, Olga Vrousgos, Corby Rosset, et al. Agentinstruct: Toward gen-
 650 erative teaching with agentic flows. *arXiv preprint arXiv:2407.03502*, 2024a.

651 Arindam Mitra, Hamed Khanpour, Corby Rosset, and Ahmed Awadallah. Orca-math: Unlocking
 652 the potential of slms in grade school math. *arXiv preprint arXiv:2402.14830*, 2024b.

653 Sumeet Ramesh Motwani, Chandler Smith, Rocktim Jyoti Das, Rafael Rafailov, Ivan Laptev,
 654 Philip HS Torr, Fabio Pizzati, Ronald Clark, and Christian Schroeder de Witt. Malt: Improv-
 655 ing reasoning with multi-agent llm training. *arXiv preprint arXiv:2412.01928*, 2024.

656 Jinlong Pang, Jiaheng Wei, Ankit Parag Shah, Zhaowei Zhu, Yaxuan Wang, Chen Qian, Yang Liu,
 657 Yujia Bao, and Wei Wei. Improving data efficiency via curating llm-driven rating systems. *arXiv
 658 preprint arXiv:2410.10877*, 2024.

659 Qizhi Pei, Lijun Wu, Zhuoshi Pan, Yu Li, Honglin Lin, Chenlin Ming, Xin Gao, Conghui He,
 660 and Rui Yan. Mathfusion: Enhancing mathematical problem-solving of llm through instruction
 661 fusion. *arXiv preprint arXiv:2503.16212*, 2025.

662 David Saxton, Edward Grefenstette, Felix Hill, and Pushmeet Kohli. Analysing mathematical rea-
 663 soning abilities of neural models. In *International Conference on Learning Representations*,
 664 2019.

665 Zhengyang Tang, Xingxing Zhang, Benyou Wang, and Furu Wei. Mathscale: Scaling instruction
 666 tuning for mathematical reasoning. *arXiv preprint arXiv:2403.02884*, 2024.

667 Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut,
 668 Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly
 669 capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.

670 Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya
 671 Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. Gemma: Open
 672 models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*, 2024.

673 Qwen Team. Qwen2.5: A party of foundation models, September 2024. URL <https://qwenlm.github.io/blog/qwen2.5/>.

674 Yuxuan Tong, Xiwen Zhang, Rui Wang, Ruidong Wu, and Junxian He. Dart-math: Difficulty-aware
 675 rejection tuning for mathematical problem-solving. *Advances in Neural Information Processing
 676 Systems*, 37:7821–7846, 2024.

677 Shubham Toshniwal, Ivan Moshkov, Sean Narenthiran, Daria Gitman, Fei Jia, and Igor Gitman.
 678 Openmathinstruct-1: A 1.8 million math instruction tuning dataset. *Advances in Neural Infor-
 679 mation Processing Systems*, 37:34737–34774, 2024.

680 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
 681 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and
 682 efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.

683 Ke Wang, Houxing Ren, Aojun Zhou, Zimu Lu, Sichun Luo, Weikang Shi, Renrui Zhang, Linqi
 684 Song, Mingjie Zhan, and Hongsheng Li. Mathcoder: Seamless code integration in llms for en-
 685 hanced mathematical reasoning. *arXiv preprint arXiv:2310.03731*, 2023.

686 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdh-
 687 ery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models.
 688 *arXiv preprint arXiv:2203.11171*, 2022.

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

692 Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun
 693 Zhang, Shaokun Zhang, Jiale Liu, et al. Autogen: Enabling next-gen llm applications via multi-
 694 agent conversation. *arXiv preprint arXiv:2308.08155*, 2023.

702 Shijie Xia, Xuefeng Li, Yixin Liu, Tongshuang Wu, and Pengfei Liu. Evaluating mathematical
 703 reasoning beyond accuracy. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
 704 volume 39, pp. 27723–27730, 2025.

705
 706 Hai Ye, Mingbao Lin, Hwee Tou Ng, and Shuicheng Yan. Multi-agent sampling: Scaling inference
 707 compute for data synthesis with tree search-based agentic collaboration. *arXiv preprint*
 708 *arXiv:2412.17061*, 2024.

709 Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T Kwok, Zhen-
 710 guo Li, Adrian Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical questions
 711 for large language models. *arXiv preprint arXiv:2309.12284*, 2023.

712
 713 Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting Dong, Keming Lu, Chuanqi Tan, Chang Zhou,
 714 and Jingren Zhou. Scaling relationship on learning mathematical reasoning with large language
 715 models. *arXiv preprint arXiv:2308.01825*, 2023.

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

719 Jiaxin Zhang, Zhongzhi Li, Mingliang Zhang, Fei Yin, Chenglin Liu, and Yashar Moshfeghi. Geoval:
 720 benchmark for evaluating llms and multi-modal models on geometry problem-solving. *arXiv*
 721 *preprint arXiv:2402.10104*, 2024a.

722
 723 Zhihan Zhang, Tao Ge, Zhenwen Liang, Wenhao Yu, Dian Yu, Mengzhao Jia, Dong Yu, and Meng
 724 Jiang. Learn beyond the answer: Training language models with reflection for mathematical
 725 reasoning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language
 726 Processing*, pp. 14720–14738, 2024b.

727 Zhaowei Zhu, Yiwen Song, and Yang Liu. Clusterability as an alternative to anchor points when
 728 learning with noisy labels. In *International Conference on Machine Learning*, pp. 12912–12923.
 729 PMLR, 2021.

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 **A APPENDIX**
757758 **A.1 THE USE OF LARGE LANGUAGE MODELS (LLMs)**
759760 In this work, LLMs are used to assist with text revision and grammar refinement, ensuring concise
761 and fluent writing. LLMs further support formatting adjustments for figures and tables, improving
762 readability and consistency across the paper. LLMs are also applied to refine mathematical notation,
763 adjust formula symbols, and standardize technical expressions, helping maintain clarity and precision
764 throughout the manuscript. LLMs serve only as auxiliary tools, with all final decisions and
765 edits made by the authors.
766767 **A.2 TRAINING AND INFERENCE DETAILS**
768769 All models—including our baseline reproductions—are fine-tuned for 3 epochs using a global
770 batch size of 96 on 6×NVIDIA A800 GPUs. We adopt a peak learning rate of 1e-6 (5e-6 for
771 DeepSeekMath-7B), combined with a linear warm-up over the first 3% of steps and cosine decay
772 thereafter. The maximum sequence length is fixed at 4096 tokens.
773774 During inference, we fix the sampling temperature to 0 to ensure deterministic outputs, and set the
775 maximum generation length (max tokens) to 2048 for all models. We use a fixed random seed of 0
776 for reproducibility and set the number of inference trials to 1 for every evaluation. For our primary
777 models, we adopt a standard Chain-of-Thought (CoT) prompting scheme. Specifically:
778779

- **Training prompt:** Question: {problem} Answer:
- **Evaluation prompt:** Question: {problem} Answer: Let's think step
780 by step.

781782 This prompt design follows common practice in mathematical reasoning tasks and encourages the
783 model to generate explicit intermediate reasoning steps. For Mistral 7B and Llama 3 8B, we instead
784 use the Alpaca instruction-following template during inference:
785
786787 Below is an instruction that describes a task. Write
788 a response that appropriately completes the request:
789 ### Instruction:
790 {problem}
791 ### Response:792 We adopt this template because our preliminary experiments showed that Alpaca-style instructions
793 consistently yield better reasoning quality on these two architectures compared with the CoT-style
794 prompt. This observation is also aligned with the findings reported in MathFusion, where Alpaca-
795 style prompting was similarly found to be more effective.
796797 **A.3 EVALUATION BENCHMARKS**
798799 We provide detailed descriptions of the six benchmarks used in our evaluation:
800801 **In-Domain:** (i) GSM8K (Cobbe et al., 2021), consisting of grade-school arithmetic word problems
802 that are relatively simple. (ii) MATH (Hendrycks et al., 2021), a large-scale dataset of competition-
803 level problems that are significantly more challenging.
804805 **Out-of-Domain (OOD):** (i) CollegeMath (Tang et al., 2024), with 2,818 college-level problems
806 drawn from nine textbooks across seven domains (e.g., linear algebra, differential equations), de-
807 signed to test generalization to complex mathematics. (ii) DeepMind-Mathematics (Saxton et al.,
808 2019), a collection of 1,000 problems covering a national school curriculum (up to age 16), assess-
809 ing basic reasoning across varied types. (iii) OlympiadBench-Math (He et al., 2024), providing 675
810 Olympiad-level problems (English text-only subset) targeting the most challenging reasoning tasks.
811 (iv) TheoremQA (Chen et al., 2023), consisting of 800 problems that require applying mathematical
812 theorems across mathematics, physics, and engineering, testing theoretical reasoning in STEM.
813

810
811
812
813
814
815
816
817
818
819
Table 5: 10K-sample comparison using the same teacher model.

Dataset	# Samples	MATH	GSM8K	College	DM	Olympiad	Theorem	AVG
MetaMath	10K	24.9	70.4	19.0	28.5	5.1	13.8	26.9
DARTMath	10K	29.3	66.4	21.3	37.4	8.7	16.7	29.9
ScaleQuest	10K	24.9	66.4	17.5	27.3	7.7	14.1	26.3
AgenticMath (ours)	10K	29.6	73.8	24.6	38.2	7.4	16.3	31.6

820
821
822
823
824
825
826
827
828
829
A.4 FAIRNESS CONCERNs REGARDING TEACHER MODELS830
831
832
833
For further strengthen the fairness discussion, we conduct an additional 10K-sample controlled comparison across MetaMath, DARTMath, ScaleQuest, and our AgenticMath, all trained under the same SFT configuration (Mistral-7B) and, where applicable, using the same teacher model (GPT-4o-mini, 2024-07-18) for solution generation. As shown in Table 5, AgenticMath achieves the highest average performance among all methods. This controlled experiment confirms that the observed improvements do not arise from using a stronger teacher model—since all datasets share the same teacher—but rather from the design of our synthesis pipeline itself.834
835
836
837
838
839
A.4.1 SENSITIVITY ANALYSIS ON THE REVISE THRESHOLD830
831
832
833
To further examine the impact of the revise threshold, we conduct a sensitivity study with $\tau_{\text{rev}} \in \{3.5, 4.0, 4.5\}$ using Llama3-8B under a fixed 30K SFT setting. The results are reported in Table 6.840
841
842
843
844
Table 6: Sensitivity analysis of the revise threshold τ_{rev} using Llama3-8B.

τ_{rev}	Samples	MATH	GSM8K	College	DM	Olympiad	Theorem	AVG
4.5	30K	36.8	78.4	29.6	40.3	11.4	20.4	36.2
4.0	30K	36.6	77.5	28.2	43.1	11.5	20.0	36.2
3.5	30K	37.8	77.4	27.4	41.0	10.3	20.0	35.7

840
841
842
843
844
As shown in Table 6, the settings $\tau_{\text{rev}} = 4.0$ and 4.5 produce highly consistent results, demonstrating that the review–revise mechanism remains stable across reasonable threshold choices. In contrast, $\tau_{\text{rev}} = 3.5$ yields lower overall performance, which is expected since a looser threshold admits more low-quality candidates into subsequent stages, ultimately weakening the final dataset quality.845
846
847
848
849
850
851
A.4.2 ANALYSIS OF THE THREE REVIEW–REVISE ITERATIONS.849
850
851
To further address the reviewer’s question regarding the choice of three review–revise iterations, we evaluate the quality of the problems that pass each round using GPT-4o-mini. Specifically, we compute the average *complexity*, *information value*, and *clarity* scores for all accepted problems after each iteration.853
854
855
856
857
858
859
Table 7: Quality metrics across three review–revise iterations.

Metric	Round 1	Round 2	Round 3
Complexity	3.86	3.93	3.92
Information Value	3.96	4.03	4.02
Clarity	4.35	4.44	4.45
Avg	4.06	4.13	4.13

860
861
862
863
As shown in Table 7, the first two review–revise iterations produce consistent improvements across all metrics. By the third iteration, however, the gains largely stabilize, indicating diminishing returns. This analysis supports our design choice of using three iterations: it captures most of the quality improvements while avoiding unnecessary computation beyond the point of saturation.

864

865

866 **A.5 ADDITIONAL ANALYSIS OF SYNTHETIC DATA QUALITY AND CHARACTERISTICS**

867

868 To provide further quantitative evidence of the quality and semantic composition of the synthesized
 869 data, we conduct a post-hoc analysis using GPT-4o-mini as an external evaluator. Specifically,
 870 we (i) assign quality scores to the refined problems and (ii) classify each problem into standard
 871 mathematical topics.

872

873

874 **A.5.1 QUALITY SCORE DISTRIBUTION**

875

876

877

Table 8 reports the quality distribution assigned by GPT-4o-mini over the 18,679 refined problems.
 The majority of synthesized questions receive a score of ≥ 4 , indicating strong clarity, coherent
 reasoning, and non-trivial complexity across the dataset.

878

879

Table 8: Quality score distribution of 18,679 refined problems, evaluated by GPT-4o-mini.

880

Quality Score	Count	Percentage
1	307	1.64%
2	1175	6.29%
3	5053	27.05%
4	12142	65.00%
5	2	0.01%

881

882

883

884

885 **A.5.2 TOPIC DISTRIBUTION OF THE FINAL 15K DATASET**

886

GPT-4o-mini is further used to classify each problem in the final 15,000-sample dataset into standard
 mathematical domains. The resulting topic distribution is shown in Table 9. The distribution demon-
 strates broad semantic coverage across major mathematical disciplines, with strong representation
 in combinatorics, geometry, algebra, and number theory.

887

888

889

Table 9: Topic distribution of the final 15K synthetic dataset (classified by GPT-4o-mini).

890

891

892

893

894

895

Topic	Count	Percentage
Counting & Probability	3705	24.70%
Geometry	3326	22.17%
Algebra	2446	16.31%
Number Theory	1475	9.83%
Calculus	1470	9.80%
Precalculus	1456	9.71%
Intermediate Algebra	907	6.05%
Prealgebra	123	0.82%
Linear Algebra	71	0.47%
Others	21	0.14%

896

897

898

899

900

901

902

903

904

905

906

907

911 **A.6 DETAILED STATISTICS OF THE DATA GENERATION PIPELINE**

912

913

914

915

916

917

To provide a clearer understanding of the robustness and transparency of our data generation
 pipeline, we report detailed statistics for all major stages, including seed scoring, rephrase expansion,
 the multi-round review–revise refinement process, and the final data evaluation. These analyses
 illustrate how the pipeline progressively improves *problem diversity, clarity, and complexity*, which
 are key for enhancing downstream mathematical reasoning ability.

918 A.6.1 SEED DATA SCORING
919920 We first evaluate all raw seed questions using our label-free scoring mechanism. A total of 7,001
921 questions satisfy the filtering threshold (score ≥ 3). The full distribution is shown in Table 10. This
922 stage ensures that only seed questions with sufficient structural soundness and baseline complexity
923 are used for further synthesis.924 Table 10: Score distribution for GSM8K and MATH seed datasets.
925926
927
928
929

Dataset	Score=0	Score=1	Score=2	Score=3	Score=4	Score=5	Score ≥ 3
GSM8K	823	3522	1849	1188	91	0	1279
MATH	69	825	884	2053	3415	254	5722

930
931
932
933 A.6.2 REPHRASE EXPANSION934 To enhance problem *complexity* and *diversity* while preserving core semantics, each filtered seed
935 question is rephrased six times. All 42,006 synthesized candidates proceed to the review–revise
936 process. This expands the problem pool as follows:
937938 Table 11: Rephrase expansion of filtered seed questions.
939940
941
942

Stage	Count Calculation	Total
Rephrase Expansion	$(1279 + 5722) \times 6$	42,006

943
944
945 A.6.3 REVIEW–REVISE LOOP
946947 Our three-round review–revise process progressively improves *clarity*, *logical correctness*, and
948 *mathematical validity*. Across rounds, vague or low-quality questions are removed, while clearer
949 and more coherent problems are retained. Table 12 summarizes the filtering behavior across rounds.
950951 Table 12: Statistics of the three-round review–revise refinement.
952953
954
955
956
957

Round	Total Inputs	Passed	Pass Rate
1	42,006	7,438	17.71%
2	34,568	6,526	18.88%
3	28,042	4,718	16.83%
All Rounds	–	18,682	–

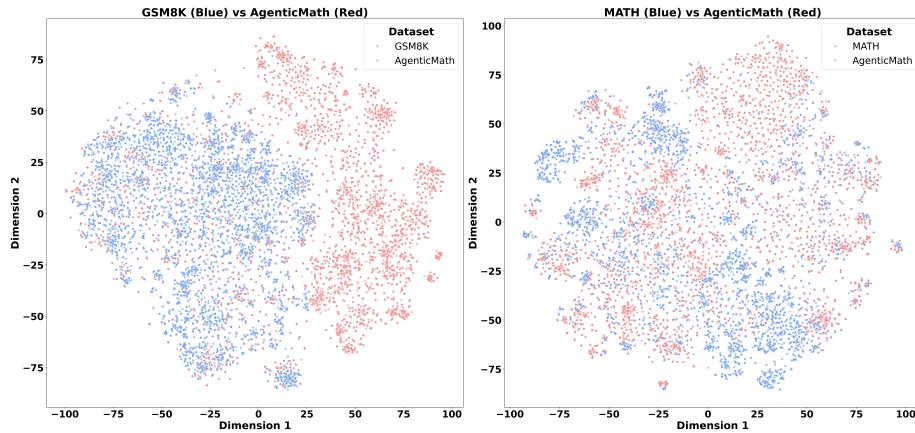


Figure 7: t-SNE Semantic Distribution.

A.6.4 SYNTHETIC DATA QUALITY DISTRIBUTION

After refinement, 18,679 high-quality synthetic problems remain. The fact that 65% of the questions are assigned a score of 4, with another 27% scoring 3, demonstrates that the majority of synthesized problems exhibit strong clarity, coherent reasoning, and meaningful complexity. Their quality distribution (evaluated by GPT-4o-mini) is shown in Table 13.

Table 13: Quality score distribution of 18,679 refined synthetic problems.

	Score	Count	Percentage
1000	1	307	1.64%
1001	2	1175	6.29%
1002	3	5053	27.05%
1003	4	12,142	65.00%
1004	5	2	0.01%

A.6.5 FINAL DATASET CONSTRUCTION

To construct the final 15K dataset used in our experiments, we jointly rank all refined samples using a combination of the quality score and the long-tail diversity score. This ranking procedure prioritizes both *overall quality* and *distributional diversity*. The top 15,000 problems from this ranked list form the final synthetic dataset.

1026 A.7 EXAMPLES OF PROBLEM REVISION
10271028 A.7.1 EXAMPLES OF PROBLEM REVISIONS FROM GSM8K
10291030 Example of problem refinement by the Reviewer and Revise Agents.
10311032 **Original Seed Problem:** Ivory got four more riddles than Josh did. Taso got twice as many
1033 riddles as Ivory did. If Josh has 8 riddles, how many riddles does Taso have?
10341035 **Rephrased Problem:** Josh is participating in a school riddle competition. He has 8 riddles.
1036 Ivory, his friend, has four more riddles than Josh. Taso, another competitor, has twice as many
1037 riddles as Ivory. If Josh decides to give away 3 of his riddles and that affects the total number of
1038 riddles in the competition, how many riddles does Taso have after accounting for Josh's decision?
10391040 **Review Score:** 3,4,4.
10411042 **Review Agent Suggestions:**
10431044 (1) Clarify the influence of Josh giving away his riddles on Taso's total. For instance, specify that
1045 Taso's riddle count is independent and remains unchanged despite Josh's actions. This removes
1046 ambiguity regarding the competition's overall structure.
10471048 (2) Rephrase the question for better clarity. Instead of saying "and that affects the total number
1049 of riddles in the competition," use "Josh's decision to give away his riddles does not change
1050 how many riddles Taso has." This direct statement enhances comprehension and ensures readers
1051 understand that Taso's initial count remains constant.
10521053 (3) Consider providing context as to whether the riddles given away affect only Josh or extend
1054 to others. A statement like "regardless of how many riddles Josh gives away, Taso's number of
1055 riddles is based on the original count" would maintain logical clarity.
10561057 **Revised Problem:** Josh is participating in a school riddle competition. He has 8 riddles. Ivory,
1058 his friend, has four more riddles than Josh. Taso, another competitor, has twice as many riddles
1059 as Ivory. Josh's decision to give away 3 of his riddles does not change how many riddles Taso
1060 has. Regardless of how many riddles Josh gives away, Taso's number of riddles is based on the
1061 original count. How many riddles does Taso have?
1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1080

Example of problem refinement by the Reviewer and Revise Agents.

1081

1082

1083

1084

1085

1086

1087

1088

1089

1090

1091

1092

1093

1094

1095

1096

1097

1098

1099

Problem: Stella and Twinkle are filling a truck with a capacity of 6000 stone blocks at the rate of 250 blocks per hour per person. They work for four hours without any breaks, after which they are joined by 6 other people who work at the same rate. Calculate the total time taken to fill the truck. Additionally, if the blocks are in the shape of cubes with a side length of 2 units, determine the area of the base of the truck that would hold these blocks when filled to capacity. Provide the total hours taken and the area of the base in square units.

Rephrased Problem: Stella and Twinkle are filling a truck with a capacity of 6000 cubic blocks at the rate of 250 blocks per hour per person. They work for four hours without any breaks. After this, they are joined by 6 additional people who work at the same rate. Calculate the total time taken to fill the truck. To begin, calculate the total number of blocks filled by both Stella and Twinkle during their initial 4 hours. Next, determine how long it will take for all workers (Stella, Twinkle, and the 6 additional people) to finish filling the truck. Additionally, the blocks are in the shape of cubes with a side length of 2 units. Calculate the area of the rectangular base of the truck needed to accommodate the cubes when filled to its full capacity, expressed in square units. Provide the total hours taken and the area of the base in square units.

Review Score: 4,4,4.

Review Agent Suggestions:

(1) Rephrase the initial sentence to: "Stella and Twinkle are filling a truck with a capacity of 6000 cubic blocks, working at a rate of 250 blocks per hour each for four hours without breaks." This enhances clarity by simplifying the phrase structure.

(2) Provide a clearer description of the mathematical operations to carry out after working four hours. For instance, "After 4 hours of work, determine the total blocks filled, and then calculate the time required for the remaining blocks with the help of the additional workers." This improves logical coherence by clarifying the task transition.

(3) Clearly request the output format by stating, "Finally, report both the total hours taken to fill the truck and the area of the rectangular base in square units," ensuring the conclusion is definitive and straightforward for the reader to apprehend.

Revise Problem: Stella and Twinkle are filling a truck with a capacity of 6000 cubic blocks, working at a rate of 250 blocks per hour each for four hours without breaks. After 4 hours of work, determine the total blocks filled by both Stella and Twinkle, and then calculate the time required for the remaining blocks with the help of 6 additional workers who work at the same rate. Finally, report both the total hours taken to fill the truck and the area of the rectangular base in square units needed to accommodate the blocks when filled to its full capacity, given that each block is a cube with a side length of 2 units.

1100

1101

1102

1103

1104

1105

1106

1107

1108

1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

1122

1123

1124

1125

1126

1127

1128

1129

1130

1131

1132

1133

1134

Example of problem refinement by the Reviewer and Revise Agents.

1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

Problem: Tom's rabbit can run at a speed of 25 miles per hour, while his cat runs at 20 miles per hour. If the cat gets a 15-minute head start, during which it covers a distance on a flat path, determine the time (in hours) it will take the rabbit to catch up with the cat assuming that for every mile, the rabbit encounters terrain that slows it down by 1 minute for every 5 miles traveled. Additionally, calculate the total distance (in square miles) covered by both animals until the rabbit catches up, assuming their paths can be represented as lines on a coordinate plane.

Rephrased Problem: Tom's rabbit can run at a speed of 25 miles per hour, while his cat runs at a speed of 20 miles per hour. The cat gets a 15-minute head start, covering a distance on a flat path during this time. Determine the time (in hours) it will take for the rabbit to catch up to the cat, defined as the moment when both animals have traveled the same linear distance from the starting point. Consider that for every mile the rabbit travels, it encounters terrain that slows its speed by an additional minute for every 5 miles traveled. Finally, calculate the total distance traveled (in miles) by both animals until the rabbit catches up with the cat.

Review Score: 3,4,3.

Review Agent Suggestions:

(1) Streamline the question for clarity. For example, state: "Tom's rabbit runs at 25 mph and his cat at 20 mph. The cat has a 15-minute head start. When will the rabbit catch up to the cat, considering that every mile the rabbit travels adds an extra minute to its speed?" This revision reduces complexity and provides clearer instructions.

(2) Clarify the impact of the terrain on the rabbit's speed. Instead of stating that its speed will decrease, specify how to adjust the time taken due to the terrain. For instance, rephrase the condition regarding the rabbit's speed to enhance understanding, such as, "For every 5 miles the rabbit runs, add 1 extra minute to its travel time."

(3) Remove unnecessary details and ensure the question leads to a unique solution. Consider separating the catch-up time and total distance calculations into distinct questions or minimizing the effects of terrain to make the problem more straightforward. This way, students can focus on the primary aspects of the problem without layers of complexity that introduce potential for confusion.

Revised Problem: Tom's rabbit runs at 25 mph and his cat at 20 mph. The cat has a 15-minute head start. For every 5 miles the rabbit travels, add an extra minute to its travel time. How long (in hours) will it take for the rabbit to catch up to the cat? Also, calculate the total distance (in miles) traveled by both animals when the rabbit catches up.

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

1188
1189

A.7.2 EXAMPLES OF PROBLEM REVISIONS FROM MATH

1190
1191
1192
1193

Example of problem refinement by the Reviewer and Revise Agents.

Problem: I have 5 marbles numbered 1 through 5 in a bag. Suppose I take out two different marbles at random. What is the expected value of the sum of the numbers on the marbles?**Rephrased Problem:** You have a bag containing 5 marbles numbered 1 through 5. You randomly draw two different marbles. If marble 1 is drawn first, then the drawn marbles are put back in the bag, and you draw again. If the first marble drawn is not 1, you keep it and draw the second marble without replacement. What is the expected value of the sum of the numbers on the marbles drawn, given that if you draw an odd-numbered marble first, you will always draw marble 2 next, while if you draw an even-numbered marble first, you will draw the next higher number available if possible? Calculate the sum for each scenario and give the overall expected value.**Review Score:** 3,4,3.**Review Agent Suggestions:**

(1) Simplify the language used in the problem statement for better clarity. For instance, instead of saying “if marble 1 is drawn first, then the drawn marbles are put back in the bag,” you could say “if marble 1 is drawn first, put the marbles back and draw again.” This succinctly conveys the same information without ambiguity.

(2) Clearly define what happens in all scenarios after each draw. It would be beneficial to write separate scenarios or cases explicitly (e.g., drawing 1, drawing an odd number, drawing an even number), ensuring that all potential outcomes are evident to the reader to aid in understanding and solving the problem.

Revised Problem: You have a bag with 5 marbles numbered 1 through 5. You will randomly draw two different marbles according to the following rules: 1. If you draw marble 1 first, put it back in the bag and draw again. 2. If the first marble drawn is not marble 1, you keep that marble and draw a second marble without putting the first back. Additionally: - If the first marble drawn is odd (3 or 5), you will draw marble 2 next. - If the first marble drawn is even (2 or 4), you will draw the next higher number available (for example, if you draw 2 first, you will draw 3 next; if you draw 4, you will draw 5 next). Calculate the expected value of the sum of the numbers on the marbles drawn, considering each scenario.1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

1242

Example of problem refinement by the Reviewer and Revise Agents.

1243

Problem: Find the number of ordered pairs of positive integers (a,b) such that $a+b=1000$ and neither a nor b has a zero digit.

1244

Rephrased Problem: In a community project, two teams need to collectively gather exactly 1000 units of resources. Let (a) represent the units collected by Team A and (b) represent those collected by Team B, such that $(a + b = 1000)$. If neither team can collect resources in amounts represented by a number containing the digit zero, how many ordered combinations of (a, b) can the teams realistically plan for their collection, considering all above constraints?

1245

Review Score: 4,4,3.

1246

Review Agent Suggestions:

1247

(1) Clarify the constraint regarding the digit zero by explicitly stating that both (a) and (b) can only consist of the digits 1-9. This can help eliminate any potential confusion about what is meant by not containing the digit zero.

1248

(2) Define what is meant by "ordered combinations". It could be interpreted in different ways—does it mean different pairings of (a) and (b) that total 1000, or is it simply asking for the number of distinct pairs? Providing examples or additional context could further enhance comprehension.

1249

(3) Instead of stating "the teams realistically plan for their collection," consider specifying that this refers to the mathematical combinations that satisfy all conditions. This would reinforce the focus of the question on mathematics rather than a broader planning context.

1250

Revised Problem: In a community project, two teams need to collectively gather exactly 1000 units of resources. Let (a) represent the units collected by Team A and (b) represent those collected by Team B, such that $(a + b = 1000)$. Both (a) and (b) must be composed exclusively of the digits 1 through 9, meaning neither can include the digit zero. How many distinct ordered pairs $((a, b))$ satisfy this condition, ensuring all constraints are met?

1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

1266

1267

1268

1269

1270

1271

1272

1273

1274

1275

1276

1277

1278

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295

1296
1297

A.8 AGENT PROMPT

1298
1299

Example 1: Problem Rating Prompt

1300
1301
1302

As a mathematics quality evaluator, your task is to rigorously assess whether a given mathematical question is high-quality and valuable as a reference for synthesizing new questions. Use the following criteria:

1303
1304
1305

1. Complexity (1–10): How many logical steps or conceptual layers are required? Does it integrate multiple mathematical domains (e.g., algebra + geometry) or demand critical thinking?
2. Information Value: Does it contain useful knowledge or reasoning opportunities? Can it help learners discover concepts, strategies, or patterns?
3. Clarity & Precision (1–10): Is the question unambiguous, logically consistent, and free of errors? Poorly framed questions score lower.

1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317

** Scoring Guidelines **:

- Please rate the sample on a scale from 1 to 10 for each criterion, and return an overall rating on a scale from 1 to 10, where a higher score indicates higher level of quality.
- Ensure that the ratings are not overly concentrated around a specific score. If multiple samples have similar qualities, consider spreading the scores more evenly to reflect subtle differences.
- Penalize heavily for ambiguity, errors, or oversimplification.

Please carefully evaluate the following data sample and return the integral evaluation scores using the JSON format:

```
{
  "Complexity": <number, 1–10>,
  "Information Value": <number, 1–10>,
  "Clarity": <number, 1–10>,
  "Overall rating": <number, 1–10>
}
```

1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

1350
 1351 Example 2: Problem Rephrase Prompt
 1352 Act as an expert mathematics educator specializing in problem complexity escalation.
 1353 Systematically transform the given problem while preserving its core concepts, using the
 1354 following framework:
 1355
 1356 ****Stage 1: Problem Deconstruction****
 1357 - Domain Identification: [Algebra/Geometry/Calculus/etc.]
 1358 - Core Competencies: [List specific theorems/formulas/methods]
 1359 - Baseline Difficulty: [Level 1–5 using Krathwohl’s Cognitive Rigor Index]
 1360
 1361 ****Stage 2: Escalation Protocol****
 1362 Select ≥ 3 complexity dimensions from:
 1363 1. Multi-stage Transformation: Designs a single, cohesive mathematical problem where the
 1364 complete solution inherently demands multiple, sequentially dependent calculations. The
 1365 output of one implicit intermediate step must serve as the essential and sole input for the
 1366 next, creating a longer chain of necessary computational derivation for the solver to reach
 1367 the definite final answer.
 1368 2. Cross-domain Integration: Create hybrid problems combining ≥ 2 mathematical disci-
 1369 plines
 1370 3. Real-world Parameterization: Embed contextual constraints with multivariate relation-
 1371 ships
 1372 4. Conditional Branching: Introduce layered constraints requiring decision-tree analysis
 1373 5. Inverse Problem Design: Reverse-engineer given solutions to reconstruct premises
 1374 6. Uncertainty Integration: Incorporate measurement errors/probabilistic factors
 1375 7. Optimization Extension: Convert closed solutions into multi-objective optimization
 1376 challenges
 1377
 1378 ****Stage 3: Revise question****
 1379 - Must be a definitive mathematical problem: The question must require mathematical
 1380 reasoning, calculation, or logical deduction.
 1381 - Must have a unique and specific mathematical answer: The problem should lead to a sin-
 1382 gle, verifiable numerical or analytical solution, avoiding open-ended questions, subjective
 1383 evaluations, or non-mathematical tasks.
 1384
 1385 Please reply strictly in the following format:
 1386 Stage 1
 1387 #Problem Deconstruction#:
 1388 Stage 2
 1389 #Escalation Protocol#:
 1390 Stage 3
 1391 #Finally Rewritten question#:
 1392
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403

1404
1405

Example 3: Problem Review Prompt

1406

As a mathematics quality checker, your task is to rigorously assess whether a given mathematical question is high-quality and provide rewrite suggestions:

1407

1. Clarity & Grammar (1–5): The question must be grammatically correct, precisely phrased, and easy to understand. It should avoid ambiguity in wording or phrasing.

1408

2. Logical Coherence & Completeness (1–5): All elements of the problem (e.g., given information, constraints, relationships, objectives) must be logically interconnected and sufficient. The problem should present a clear, sequential path for reasoning, without missing information required for the specified solution approach.

1409

3. Mathematical Validity & Solvability (1–5): The problem must be fundamentally a mathematics problem, with all its premises and conditions being **mutually consistent** and **mathematically sound**. It must lead to a **unique, solvable numerical or analytical answer** that adheres to all mathematical rules and specified ranges (e.g., probabilities summing to 1, valid geometric properties, real number solutions). If any condition leads to a mathematical contradiction or an impossible/undefined solution (e.g., total probability > 1 after adjustments, an equation with no valid solution within given constraints), this criterion rates very low, and the exact mathematical inconsistency must be pinpointed. Avoid open-ended or non-mathematical questions.

1410

**** Scoring Guidelines **:**

1411

- Please rate the sample on a scale from 1 to 5 for each criterion, and return an overall rating on a scale from 1 to 5, where a higher score indicates higher level of quality.

1412

Rephrased question: {rephrased_question}

1413

****Output Requirements****

1414

Respond in the following plain-text format ****only**** (do not include JSON or any additional commentary):

1415

###thought###

1416

<Analytical reasoning addressing each criterion sequentially, especially for rephrased_question >

1417

###rating_score###

1418

<Clarity & Grammar score >, <Logical Consistency score >, <Mathematical Relevance & Solvability score >

1419

###suggestions###

1420

###Specific improvement 1###

1421

<Specific improvement 1 >

1422

###Specific improvement 2###

1423

<Specific improvement 2 >

1424

...more improvements if needed...

1425

Noice:

1426

- "rating_score" represents evaluate score of Rephrased question.

1427

- when generate "suggestions", please give more details and reasons for each improvement.

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

1458

1459

1460

1461

1462

1463

1464

1465

1466

1467

1468

1469

1470

1471

1472

1473

1474

1475

1476

1477

1478

1479

1480

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

1491

1492

1493

1494

1495

1496

1497

1498

1499

1500

1501

1502

1503

1504

1505

1506

1507

1508

1509

1510

1511

Example 4: Problem Revise Prompt

As an expert in mathematical question improvement, please optimize the question according to the following suggestions:

{suggestions}

Optimization requirements:

1. Clarity & Grammar (1–5): The question must be grammatically correct, precisely phrased, and easy to understand. It should avoid ambiguity in wording or phrasing.
2. Logical Coherence & Completeness (1–5): All elements of the problem (e.g., given information, constraints, relationships, objectives) must be logically interconnected and sufficient. The problem should present a clear, sequential path for reasoning, without missing information required for the specified solution approach.
3. Mathematical Validity & Solvability (1–5): The problem must be fundamentally a mathematics problem, with all its premises and conditions being *mutually consistent* and *mathematically sound*. It must lead to a *unique, solvable numerical or analytical answer* that adheres to all mathematical rules and specified ranges (e.g., probabilities summing to 1, valid geometric properties, real number solutions). If any condition leads to a mathematical contradiction or an impossible/undefined solution (e.g., total probability exceeds 1 after adjustments, an equation with no valid solution within given constraints), this criterion rates very low, and the exact mathematical inconsistency must be pinpointed. Avoid open-ended or non-mathematical questions.

original question: {rephrased_question}

** Output Requirements **

Respond in the following plain-text format **only** (do not include JSON or any additional commentary):

###revised_question###
 <improved full question>
 ###revision_notes###
 <Specific revision note>

1512
1513

Example 5: Solution Generation Prompt (GSM8K)

1514
1515

As a mathematics problem solving expert, analyze and answer the following question.

1516

Workflow:

1517

1. Analyze and Deconstruct:

1518

- First, systematically break down the problem into its core components.

1519

- Explicitly list all given data, variables, constraints, and the final objective of the problem.

1520

2. Clarify Ambiguities:

1521

- Before starting calculations, if any part of the problem statement is ambiguous, you must state your interpretation and the reasoning behind it.

1522

3. Step-by-Step Derivation and Process Demonstration:

1523

- For each component of the problem, provide a detailed step-by-step derivation.

1524

- You must show all intermediate calculation steps, formulas used, and logical judgments. Do not skip or summarize critical calculation processes.

1525

- For any step involving complex calculations, multi-case analysis, or iterative enumeration (e.g., filtering combinations that meet a condition, solving systems of equations, analyzing multiple scenarios), you must clearly list all cases or combinations considered.

1526

4. Synthesis and Final Calculation:

1527

- Integrate the results from all preceding steps to perform the final calculation.

1528

- Clearly show the final calculation that leads to the final answer.

1529

Respond in the following plain-text format ****only**** (do not include JSON or any additional commentary):

###thought### <step-by-step reasoning process> ###answer### <final answer>

1530

Output Notice:

1531

- Replace <step-by-step reasoning process> with your detailed derivation.

1532

- Replace <final answer> with the concise final answer (e.g., a number or fraction), without units or extra words.

1533

Output Example 1:

1534

Question: A cleaning company produces two sanitizer sprays. One spray kills 50% of germs, and another spray kills 25% of germs. However, 5% of the germs they kill are the same ones. What percentage of germs would be left after using both sanitizer sprays together?

1535

Output(must match the specified format exactly):

1536

###thought### To correctly calculate the percentage of germs left, we must use the Principle of Inclusion-Exclusion to find the total percentage of unique germs killed

1537

###answer### 30

1538

Question: {question}

1539

Output:

1540

1541

1542

1543

1544

1545

1546

1547

1548

1549

1550

1551

1552

1553

1554

1555

1556

1557

1558

1559

1560

1561

1562

1563

1564

1565

1566
1567

Example 6: Solution Generation Prompt (MATH)

1568
1569

As a mathematics problem solving expert, analyze and answer the following question.

1570

Workflow:

1571

1. Analyze and Deconstruct:

1572

- First, systematically break down the problem into its core components.
- Explicitly list all given data, variables, constraints, and the final objective of the problem.

1573

2. Clarify Ambiguities:

1574

- Before starting calculations, if any part of the problem statement is ambiguous, you must state your interpretation and the reasoning behind it.

1575

3. Step-by-Step Derivation and Process Demonstration:

1576

- For each component of the problem, provide a detailed step-by-step derivation.

1577

- You must show all intermediate calculation steps, formulas used, and logical judgments. Do not skip or summarize critical calculation processes.

1578

- For any step involving complex calculations, multi-case analysis, or iterative enumeration (e.g., filtering combinations that meet a condition, solving systems of equations, analyzing multiple scenarios), you must clearly list all cases or combinations considered.

1579

4. Synthesis and Final Calculation:

1580

- Integrate the results from all preceding steps to perform the final calculation.

1581

- Clearly show the final calculation that leads to the final answer.

1582

Respond in the following plain-text format ****only**** (do not include JSON or any additional commentary):

###thought### <step-by-step reasoning process> ###answer### <final answer>

1583

Output Notice:

1584

- Replace **<step-by-step reasoning process>** with your detailed derivation.

1585

- Replace **<final answer>** with the concise final answer (e.g., a number or fraction), without units or extra words.

1586

Output Example 1:

1587

Question: A box contains 5 white balls and 6 black balls. Two balls are drawn out of the box at random. What is the probability that they both are white?

1588

Output(must match the specified format exactly):

1589

###thought### To solve for the probability of drawing two white balls from a box containing 5 white and 6 black balls, we'll use.....

1590

###answer### $\frac{2}{11}$

1591

Question: {question}

1592

Output:

1593

1594

1595

1596

1597

1598

1599

1600

1601

1602

1603

1604

1605

1606

1607

1608

1609

1610

1611

1612

1613

1614

1615

1616

1617

1618

1619