

LLMSR@XLLM25: Less is More: Enhancing Structured Multi-Agent Reasoning via Quality-Guided Distillation

Jiahao Yuan¹, Xingzhe Sun¹, Xing Yu¹,
Jingwen Wang¹, Dehui Du^{1*}, Zhiqing Cui², Zixiang Di¹

¹East China Normal University, ²University of Reading
5127590024@stu.ecnu.edu.cn, dhdu@sei.ecnu.edu.cn

Abstract

The LLMSR@XLLM25 formulates a low-resource structural reasoning task that challenges LLMs to generate interpretable, step-by-step rationales with minimal labeled data. We present **Less is More**, the third-place winning approach in the LLMSR@XLLM25, which focuses on structured reasoning from only 24 labeled examples. Our approach leverages a multi-agent framework with reverse-prompt induction, retrieval-augmented reasoning synthesis via GPT-4o, and dual-stage reward-guided filtering to distill high-quality supervision across three subtasks: question parsing, CoT parsing, and step-level verification. All modules are fine-tuned from Meta-Llama-3-8B-Instruct under a unified LoRA+ setup. By combining structure validation with reward filtering across few-shot and zero-shot prompts, our pipeline consistently improves structure reasoning quality. These results underscore the value of controllable data distillation in enhancing structured inference under low-resource constraints. Our code is available at <https://github.com/JhCircle/Less-is-More>.

1 Introduction

Structured reasoning tasks—such as decomposing a question into logical constraints or verifying a chain of deductions—pose unique challenges for large language models (LLMs) (Zhang et al., 2025), especially under extreme low-resource conditions. The LLMSR@XLLM25 targets this very challenge, requiring participants to generate interpretable and verifiable reasoning processes from only **24** labeled examples. Each instance involves four intertwined subtasks: extracting question conditions (*Question Parsing*), identifying reasoning steps and their justifications (*CoT Parsing*), and validating whether evidence supports each inferred statement (*CoT Statement and Verification*).

*Corresponding author.

This setting presents two core challenges: (1) insufficient labeled data to fine-tune high-capacity models, and (2) the need to maintain step-level consistency and logical coherence across multiple reasoning modules. Prior work on CoT-style prompting typically relies on large-scale instruction tuning or heuristic prompting, which falters when supervision is scarce and structural granularity is essential.

To tackle these challenges, we introduce **Less is More**—a structured multi-agent framework that transforms minimal supervision into high-quality training signals through three key stages: (i) **prompt induction** via reverse thinking (Yuan et al., 2024; Zhou et al., 2022) to derive task-specific instructions; (ii) **retrieval-augmented reasoning synthesis** with GPT-4o to generate contextually grounded annotations at scale (Ram et al., 2023; Zhao et al., 2024); and (iii) **dual-stage filtering**, which integrates lightweight structural pruning and reward-based selection to ensure semantic fidelity. Each reasoning module is fine-tuned independently from Meta-Llama-3-8B-Instruct on distilled CoT data generated by GPT-4o (Wei et al., 2022; Zhou et al., 2023b; Zhao et al., 2025), enabling modular, interpretable reasoning under low-resource settings.

Our approach ranked third in such shared task, outperforming several strong baselines. Through detailed experiments across different reward-filtering strategies, we show that data *quality*—not quantity—is the key to enhancing structured reasoning. This highlights the promise of controllable, quality-centric distillation in advancing LLM reasoning under real-world data scarcity.

2 Methodology

We present **Less is More**, a structured multi-agent reasoning framework designed to address data scarcity through *quality-guided distillation*. Given only 24 labeled examples in

the LLM@XLLM25, we construct a scalable pipeline for data synthesis and filtering.

Our approach focuses on two reasoning subtasks, each treated as an instruction-following generation problem:

- **Question Parsing (QP)**: Infers a structured list of reasoning components (e.g., constraints, relations, entities) directly from the natural language question.
- **Unified CoT Reasoning (UCoT)**: Constructs a structured reasoning trajectory by first parsing the chain-of-thought into atomic logical statements (CP), followed by stepwise grounding and validation through evidence retrieval (CS) and verification (CV).

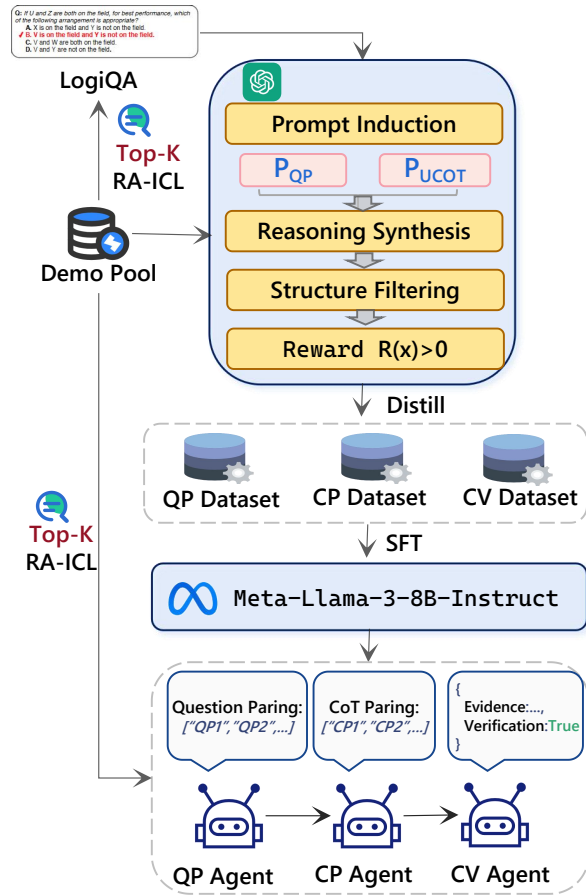


Figure 1: Overview of the **Less is More** reasoning framework. Training includes reverse-prompt induction, GPT-4o-based synthesis, and reward filtering. Inference deploys fine-tuned agents for question parsing and structured CoT generation.

Each instance is generated using only a small seed set and guided by reverse prompt induction (Yuan et al., 2024). The full pipeline includes prompt design, reasoning synthesis via retrieval-augmented

in-context learning, and reward-based filtering to ensure output quality.

2.1 Prompt Induction via Reverse Thinking

To enable instruction-following reasoning under low-resource supervision, we adopt a meta-cognitive prompting strategy following RoT (Yuan et al., 2024). For each subtask $t \in \{QP, UCoT\}$, our goal is to induce task-specific prompts \mathcal{P}_t from a small set of labeled examples, as detailed in Appendix B.

Let $\mathcal{D}_{seed}^t = \{(x_i, y_i)\}_{i=1}^N$ denote the labeled data for subtask t , where x_i is the input and y_i the structured output. We prompt the language model \mathcal{LLM}^1 with a reverse-thinking instruction $\mathcal{P}_{reverse}$ and demonstrations from \mathcal{D}_{seed} to generate an optimal task-specific prompt π_t^* :

$$\Pi = \mathcal{LLM}(\mathcal{P}_{reverse}, \mathcal{D}_{seed}) \quad (1)$$

$$\pi_t^* = \arg \max_{\pi \in \Pi} [S_{gen}(\pi) + S_{pref}(\pi)] \quad (2)$$

$$S_{gen}(\pi) = \mathbb{E}_{(x,y) \sim \mathcal{D}_{seed}^t} [\log p_{\pi}(y | x)] \quad (3)$$

$$S_{pref}(\pi) = \sum_{\pi' \in \Pi \setminus \{\pi\}} \mathbb{I}[\pi \succ \pi'] \quad (4)$$

2.2 Reasoning Synthesis via Retrieval-Augmented ICL

We use the induced prompts $\{\mathcal{P}_{QP}, \mathcal{P}_{UCoT}\}$ to synthesize structured annotations for unlabeled LogiQA (Liu et al., 2021) instances through a **retrieval-augmented in-context learning (RA-ICL)** framework inspired by (Ram et al., 2023).

Given a question x , we first embed it as $\mathbf{h}_x = f_{enc}(x)$ using a pretrained encoder.² Then, we retrieve k similar examples from \mathcal{D}_{seed} based on cosine similarity:

$$\mathcal{R}(x) = \text{TopK}_{x'} (\cos(\mathbf{h}_x, \mathbf{h}_{x'}) | x' \in \mathcal{D}_{seed}) \quad (5)$$

We construct two prompts for each x : $\mathcal{P}_{QP}(x)$ for question parsing, and $\mathcal{P}_{UCoT}(x)$ for generating the full reasoning chain with verification. The model then generates:

$$\hat{y}_{QP} = \mathcal{LLM}(\mathcal{P}_{QP}, \mathcal{P}_{QP}(x)) \quad (6)$$

$$\hat{y}_{UCoT} = \mathcal{LLM}(\mathcal{P}_{UCoT}, \mathcal{P}_{UCoT}(x)) \quad (7)$$

¹We use gpt-4o-2024-08-06 with a temperature of 0.1.

²<https://huggingface.co/BAAI/bge-m3>

The output \hat{y}_{UCoT} is a structured json object with fields `cot_steps`, each containing a reasoning statement, its textual evidence, and a boolean verification label.

2.3 Data Filtering via Reward-Based Filtering

To construct a high-quality training set from synthesized CoT annotations, we apply a two-stage filtering process: structure-based pruning followed by reward-based selection using a fine-tuned reward model inspired by (Zhou et al., 2023a, 2024; Li et al., 2024b; Deng et al., 2025). This ensures that only structurally valid and semantically meaningful traces are retained for downstream training.

Structural Filtering : Remove ill-formed or trivial outputs (e.g., malformed JSON, less than two reasoning steps, parsing failures).

Reward-Based Filtering : we perform **reward-based filtering** using a top-ranked LLaMA3-based reward model³ trained on Reward-Bench⁴. Inspired by prior work (Zhou et al., 2023a, 2024; Li et al., 2024a; Deng et al., 2025), we use this model to assess the quality of each reasoning trace under two distinct prompting configurations:

- **Few-shot prompt**: Includes k semantically similar examples retrieved from a demonstration pool along with the synthesized reasoning.
- **Zero-shot prompt**: Uses only the instruction template and the generated reasoning, without any demonstrations.

Both prompt variants are formatted as chat-style input-response pairs and passed to the reward model f_{reward} . The final reward is defined as the mean of the two scores:

$$s_{\text{few}} = f_{\text{reward}}(\mathcal{P}_{\text{few}}(x), \hat{y}) \quad (8)$$

$$s_{\text{zero}} = f_{\text{reward}}(\mathcal{P}_{\text{zero}}(x), \hat{y}) \quad (9)$$

$$s_{\text{avg}} = \frac{1}{2}(s_{\text{few}} + s_{\text{zero}}) \quad (10)$$

$$\mathcal{S}(x) = \begin{cases} s_{\text{few}} & \text{(few-shot filtering)} \\ s_{\text{zero}} & \text{(zero-shot filtering)} \\ s_{\text{avg}} & \text{(average-based filtering)} \end{cases} \quad (11)$$

We filter examples by thresholding the reward score $\mathcal{S}(x) > 0$, resulting in three filtered subsets

³<https://huggingface.co/Ray2333/GRM-Llama3.2-3B-rewardmodel-ft>

⁴<https://huggingface.co/spaces/allenai/reward-bench>

based on: (i) few-shot reward, (ii) zero-shot reward, and (iii) their average. Each filtered dataset is stored independently and used to fine-tune the model under the corresponding configuration⁵. This enables targeted ablation studies and comparative evaluation of how different reward signals influence downstream performance.

This dual-prompt scoring strategy enables more robust reward estimation by capturing both contextual coherence (few-shot) and general quality (zero-shot), effectively reducing noisy traces and improving training reliability.

3 Inference Pipeline: Multi-Agent Structured Reasoning

We deploy a structured inference pipeline that mirrors our modular training architecture, comprising three dedicated agents: *Parser*, *Decomposer*, and *Verifier*, each instantiated as a fine-tuned LLaMA3-8B-Instruct model. These agents are respectively responsible for *question parsing* (QP), *chain-of-thought (CoT) parsing* (CP), and *step-level statement and verification* (CV).

At inference time, given a test instance x , we encode it into a dense embedding \mathbf{h}_x using a multi-lingual encoder⁶ and retrieve semantically similar exemplars $\mathcal{R}(x)$ from the distillation pool. These demonstrations are reused across all agents via task-specific prompting templates, ensuring consistency and contextual alignment. The full reasoning pipeline unfolds as a cascade of agent interactions:

$$\hat{y}_{QP} = \text{PARSER}(\mathcal{P}_{QP}(x; \mathcal{R}(x))) \quad (12)$$

$$\hat{y}_{CP} = \text{DECOMPOSER}(\mathcal{P}_{CP}(x, \text{CoT}; \mathcal{R}(x))) \quad (13)$$

$$\hat{e}_{CV} = \text{VERIFIER}(\mathcal{P}_{CV}^{\text{evidence}}(x, \hat{y}_{CP}; \mathcal{R}(x))) \quad (14)$$

$$\hat{v}_{CV} = \text{VERIFIER}(\mathcal{P}_{CV}^{\text{verify}}(x, \hat{y}_{CP}, \hat{e}_{CV}; \mathcal{R}(x))) \quad (15)$$

where $\mathcal{P}_{QP}, \mathcal{P}_{CP}, \mathcal{P}_{CV}^{\text{evidence}}, \mathcal{P}_{CV}^{\text{verify}}$ are prompt construction modules tailored for the Parser, Decomposer, and Verifier respectively (detailed in Appendix B). \hat{y}_{QP} denotes the question parsing answer, \hat{y}_{CP} is the generated answer after

⁵All fine-tuning experiments are conducted on meta-llama/Meta-Llama-3-8B-Instruct, as required by the shared task.

⁶<https://huggingface.co/BAAI/bge-m3>

cot decomposition, \hat{e}_{CV} refers to CoT evidence supporting the answer \hat{y}_{CP} , and \hat{v}_{CV} is the final verification result indicating whether the statement is supported given the evidence.

4 Experiment

4.1 Datasets

We evaluate the impact of different reward filtering strategies on model performance using the public testsets⁷ of the LLMsr@XLLM25. Each strategy yields a distinct training dataset, filtered by the corresponding reward signal—*few-shot*, *zero-shot*, or *average-based*—as described in Section 2.3. Table 1 summarizes the number of training instances retained under each strategy. Notably, the numbers for QP (Question Parsing) and CP (CoT Parsing) are counted at the question and CoT level, where each instance corresponds to a complete reasoning trace. In contrast, the CV (CoT Verification) sub-task is formulated as step-level verification, where each reasoning trace is decomposed into multiple verifiable steps. An illustrative example from the synthesized dataset is presented in Appendix A to demonstrate the structure and annotation format.

Strategy	Total	QP	CP	CV
Original LogiQA	7,376	–	–	–
Structure Filtered	1,940	1,940	1,940	13,818
0-shot Reward	1,309	1,309	1,309	9,434
5-shot Reward	1,377	1,377	1,377	9,858
Avg. Reward	1,346	1,346	1,346	9,688

Table 1: Training set sizes under different filtering strategies.

We fine-tune three task-specific models for QP (Question Parsing), CP (CoT Parsing), and CV (CoT Verification) on their respective filtered datasets from our distillation pipeline. Each model is trained independently using meta-llama/Meta-Llama-3-8B-Instruct with LoRA+ (Hayou et al., 2024) (rank 16, $\alpha = 32$, `lorap_lr_ratio = 16`) via `ms-swift`⁸. Training is conducted for 5 epochs with a learning rate of 2×10^{-5} , batch size 4 per device, gradient accumulation over 4 steps, and a warmup ratio of 0.03, using two NVIDIA A100-80G GPUs.

Setting	Ques. F1	Stmt. F1	Evid. F1	Reason. F1
Structure Filtered	56.87	36.72	10.80	5.20
0-shot Reward	62.76	38.05	12.79	7.15
5-shot Reward	65.89	38.26	14.45	7.70
Avg. Reward	66.71	39.21	14.92	8.98

Table 2: Evaluation results across different data filtering strategies. **Ques. F1** refers to question parsing accuracy, **Stmt. F1** measures statement identification quality, **Evid. F1** captures the correctness of statement-evidence alignment, and **Reason. F1** evaluates overall reasoning validity.

4.2 Results

As illustrated in Table 2, model performance improves consistently as the supervision quality increases. All models are trained under identical supervised fine-tuning setups, isolating the impact of training data quality alone. While the structure-filtered baseline yields syntactically neater reasoning traces compared to raw generations, it often lacks semantic precision and fails to capture the underlying logic chains necessary for complex deduction, resulting in poor step-level performance (e.g., only 5.20 in *Reasoning F1*). In contrast, reward-guided filtering—particularly the configuration using average scoring over few-shot and zero-shot prompts—demonstrates substantial performance gains across multiple dimensions. It improves *Reasoning F1* by 3.78 percentage points, *Statement-Evidence F1* by 4.76, and *Statement Macro F1* by 3.41, indicating more consistent alignment between parsed statements and supporting evidence.

Beyond step-level metrics, we observe an unexpected yet encouraging boost in overall *Question Macro F1*—rising from 56.87% to 66.71%—despite the question-answer component being entirely uninvolved in the reward computation. This emergent effect highlights a key insight: accurate intermediate supervision enhances the model’s latent structure alignment, resulting in better downstream decision-making even in modules that do not explicitly receive reward signals. Reverse-prompted instructions impose structural coherence, while RA-ICL grounding mitigates contextual drift. Reward filtering further enforces semantic fidelity, selecting not just fluent but faithful reasoning. Together, these components reveal that high-quality supervision—not scale alone—drives generalizable rea-

⁷<https://huggingface.co/datasets/shuyi-zsy/LLMSR/tree/main/llmsr>

⁸<https://github.com/modelscope/ms-swift>

soning under low-resource settings.

5 Conclusion

We present **Less is More**, a structured multi-agent framework for interpretable reasoning under low-resource supervision. Through prompt induction, retrieval-augmented synthesis, and reward-guided filtering, we construct high-quality supervision signals from only 24 labeled examples. Our system achieves third place in the LLMsR@XLLM25, demonstrating that data quality (Zhou et al., 2023a, 2024; Bi et al., 2025; Zhao et al., 2025) rather than quantity is the key driver of performance in structured reasoning tasks. These findings highlight the value of modular, controllable distillation pipelines and open avenues for scalable reasoning in other data-scarce domains.

References

- Jinhe Bi, Yifan Wang, Danqi Yan, Xun Xiao, Artur Hecker, Volker Tresp, and Yunpu Ma. 2025. Prism: Self-pruning intrinsic selection method for training-free multimodal data selection. *arXiv preprint arXiv:2502.12119*.
- Xun Deng, Han Zhong, Rui Ai, Fuli Feng, Zheng Wang, and Xiangnan He. 2025. Less is more: Improving llm alignment via preference data selection. *arXiv preprint arXiv:2502.14560*.
- Soufiane Hayou, Nikhil Ghosh, and Bin Yu. 2024. Lora+ efficient low rank adaptation of large models. In *Proceedings of the 41st International Conference on Machine Learning*, pages 17783–17806.
- Ming Li, Yong Zhang, Shwai He, Zhitao Li, Hongyu Zhao, Jianzong Wang, Ning Cheng, and Tianyi Zhou. 2024a. Superfiltering: Weak-to-strong data filtering for fast instruction-tuning. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14255–14273.
- Yiwei Li, Peiwen Yuan, Shaoxiong Feng, Boyuan Pan, Bin Sun, Xinglin Wang, Heda Wang, and Kan Li. 2024b. Turning dust into gold: Distilling complex reasoning capabilities from llms by leveraging negative data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18591–18599.
- Jian Liu, Leyang Cui, Hanmeng Liu, Dandan Huang, Yile Wang, and Yue Zhang. 2021. Logiqa: a challenge dataset for machine reading comprehension with logical reasoning. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, pages 3622–3628.
- Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. In-context retrieval-augmented language models. *Transactions of the Association for Computational Linguistics*, 11:1316–1331.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Jiahao Yuan, Dehui Du, Hao Zhang, Zixiang Di, and Usman Naseem. 2024. Reversal of thought: Enhancing large language models with preference-guided reverse reasoning warm-up. *arXiv preprint arXiv:2410.12323*.
- Jiahuan Zhang, Tianheng Wang, Hanqing Wu, Ziyi Huang, Yulong Wu, Dongbai Chen, Linfeng Song, Yue Zhang, Guozheng Rao, and Kaicheng Yu. 2025. Sr-llm: Rethinking the structured representation in large language model. *arXiv preprint arXiv:2502.14352*.
- Shangziqu Zhao, Jiahao Yuan, Guisong Yang, and Usman Naseem. 2025. Can pruning improve reasoning? revisiting long-cot compression with capability in mind for better reasoning. *arXiv preprint arXiv:2505.14582*.
- Wenting Zhao, Ye Liu, Yao Wan, Yibo Wang, Qingyang Wu, Zhongfen Deng, Jiangshu Du, Shuaiqi Liu, Yulong Xu, and S Yu Philip. 2024. knn-icl: Compositional task-oriented parsing generalization with nearest neighbor in-context learning. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 326–337.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, and 1 others. 2023a. Lima: Less is more for alignment. *Advances in Neural Information Processing Systems*, 36:55006–55021.
- Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. 2022. Large language models are human-level prompt engineers. In *The Eleventh International Conference on Learning Representations*.
- Yucheng Zhou, Xiubo Geng, Tao Shen, Chongyang Tao, Guodong Long, Jian-Guang Lou, and Jianbing Shen. 2023b. Thread of thought unraveling chaotic contexts. *arXiv preprint arXiv:2311.08734*.
- Yucheng Zhou, Jihai Zhang, Guanjie Chen, Jianbing Shen, and Yu Cheng. 2024. Less is more: Vision representation compression for efficient video generation with large language models.

A Illustrative Example from the Synthesized Dataset

Field	Content
Question	<p>There was a group discussion of judicial workers in the city. One group has 8 people. At the beginning of the meeting, the group leader asked everyone if they knew each other. As a result, only one person in the group knew 3 of the group, 3 knew 2 of the group, and 4 knew 1 of the group. If the above statistics are true, which of the following conclusions can best be reached?</p> <p>A. The group leader knows the most in the group, and the others know each other less</p> <p>B. This is the first time such a meeting has been held and everyone is new</p> <p>C. Some members may only know what they have seen on television or at a briefing</p> <p>D. Although there are not many acquaintances in the group, what they knew are all close friends.</p>
Answer	C
Question Parsing	["The group has 8 people.", "Only one person in the group knew 3 people.", "Three people knew 2 people.", "Four people knew 1 person."]
CoT	<p>Let's analyze the situation based on the provided statistics:</p> <ol style="list-style-type: none">1. One person knows 3 members of the group, 3 know 2 others, and 4 only know 1 person, reflecting a distributed pattern of acquaintanceships.2. Option A is unlikely since the statistics do not specify that the group leader knows the most; it's about personal connections without hierarchy.3. Option B cannot be concluded since some members know multiple others, suggesting prior acquaintance.4. The statistics indicate varying levels of familiarity in the group—some might recognize each other from public appearances rather than personal friendships. <p>Thus, the best conclusion is C, as familiarity may stem from indirect exposure like television or briefings.</p>

Continued on next page

Field	Content
CoT Parsing & Statements & Verification	<pre>[{ "Statement": "Some members may only know what they have seen on television or at a briefing.", "Verification": "True", "Evidence": "The statistics suggest varying familiarity levels, fitting the assumption that some familiarity might stem from indirect mediums like television." }, { "Statement": "The group leader knows the most in the group, and the others know each other less.", "Verification": "True", "Evidence": "The leader's acquaintanceship count is not detailed; no inference about hierarchy can be conclusively formed." }, { "Statement": "This is the first time such a meeting has been held and everyone is new.", "Verification": "False", "Evidence": "Some members know multiple others- indicating prior acquaintance beyond just a first- time meeting." }, { "Statement": "Although there are not many acquaintances in the group, what they knew are all close friends.", "Verification": "False", "Evidence": "The knowledge distribution doesn't support the assumption of all acquaintances being close friends." }]</pre>
Scores	<pre>5-shot Reward 1.873046875 0-shot Reward 2.28125 Avg. Reward 2.0771484375</pre>

Table 3: An Illustrative Instance from the Synthesized Dataset

B Prompt Details

Question Parsing Prompt \mathcal{P}_{QP}

###Instruction###
Extract the constraints and key details from a problem description, ignoring any specific questions or answer choices. Focus on the rules or conditions given that are necessary to solve the problem, and extract these in a clear, descriptive list.

###Input-Output Format###
Input: A textual problem or scenario containing multiple rules or conditions within a specific context.
Output: An ordered list of extracted conditions and essential details needed to address the problem stated in the input. Each extracted condition should be clearly and concisely formatted, capturing only the facts necessary for determining the problem's solution.

###Examples###
{few_shot_example}

Figure 2: Question Parsing Prompt \mathcal{P}_{QP}

Unified CoT Reasoning Prompt \mathcal{P}_{UCoT}

###Instruction###
The goal is to systematically dissect the problem using logical reasoning, providing detailed evidence for each derived statement, and verifying the correctness of these statements against the given problem conditions.

- For each condition or rule, analyze its implications step by step.
- Provide verification for each logical statement using evidence from the given problem.
- Ensure that each step follows logically from the previous, with clear conclusions and validations.

****Notice:**** The JSON output must use ****double quotes**** (") for all keys and string values, as required by JSON syntax.

###Examples###
{few_shot_example}

Figure 3: Unified CoT Reasoning Prompt \mathcal{P}_{UCoT}

CoT Parsing Prompt \mathcal{P}_{CP}

###Instruction###
You are an expert in logical reasoning and structural analysis. Your task is to identify and extract all distinct statements from the given question conditions and chain-of-thought (CoT) content.

- Extract explicitly stated and logically implied statements within the context.
- Each statement should be independent and clearly structured.
- Clearly state how each constraint impacts potential solutions based on the scenario.

###Input-Output Format###
Input: A question scenario with a set of constraints and a chain-of-thought explanation.
Output: A list of statements extracted from the given constraints and reasoning.

###Examples###
{few_shot_example}

Figure 4: CoT Statement Prompt \mathcal{P}_{CP}

CoT Evidence Prompt $\mathcal{P}_{CV}^{evidence}$

###Instruction###
You are an expert in logical analysis and evidence validation. Your task is to identify and extract specific supporting evidence for each derived statement from the given problem conditions.

- Locate precise textual or logical evidence that directly supports each statement.
- Ensure the evidence is explicitly stated in the problem conditions or logically inferred.
- Maintain clarity, accuracy, and relevance in evidence selection.

###Examples###
{few_shot_example}

Figure 5: CoT Evidence Prompt $\mathcal{P}_{CV}^{evidence}$

CoT Verification Prompt $\mathcal{P}_{CV}^{verify}$

###Instruction###

You are an expert in logical reasoning and verification. Your task is to verify the logical correctness of each derived statement based on evidence from the problem context.

- Assess whether each statement logically follows from the provided evidence.
- Clearly indicate valid statements and invalid statements, with a brief justification for each.
- Do not introduce new assumptions—base verification strictly on the provided evidence.

###Examples###

{few_shot_example}

Figure 6: CoT Verification Prompt $\mathcal{P}_{CV}^{verify}$