

TinyThinker: Distilling Reasoning through Coarse-to-Fine Knowledge Internalization with Self-Reflection

Shengmin Piao
Yonsei University
Seoul, South Korea
shengminp@yonsei.ac.kr

Sanghyun Park[†]
Yonsei University
Seoul, South Korea
sanghyun@yonsei.ac.kr

Abstract

Large Language Models exhibit impressive reasoning capabilities across diverse tasks, motivating efforts to distill these capabilities into smaller models through generated reasoning data. However, direct training on such synthesized reasoning data may lead to superficial imitation of reasoning process, rather than fostering a genuine integration of reasoning capabilities with underlying knowledge. To address this, we propose TinyThinker, a framework introducing two novel approaches. First, we introduce a three-stage process that incrementally guides the student model through the reasoning process, progressively refining knowledge from coarse to fine granularity. Second, we develop a two-phase training framework comprising an initial reasoning acquisition phase followed by a self-reflection phase utilizing self-generated data. Experiments on commonsense reasoning benchmarks demonstrate that TinyThinker achieves superior performance compared to baselines. Ablation studies further validate the effectiveness of each component in our framework. We expect that TinyThinker can be extended to other knowledge-intensive reasoning tasks, offering an alternative strategy for developing effective reasoning capabilities in smaller language models. Codes are available at <https://github.com/shengminp/TinyThinker>.

1 Introduction

Large language models (LLMs) demonstrate impressive reasoning capabilities across a wide range of tasks (Kaplan et al., 2020; Wei et al., 2022a), largely facilitated by *in-context learning* (Brown et al., 2020) and Chain-of-Thought (CoT) prompting (Wei et al., 2022b). These methods enable LLMs to generalize from examples of intermediate reasoning steps during inference, enabling them to effectively address complex multi-step reasoning

[†] Corresponding author.

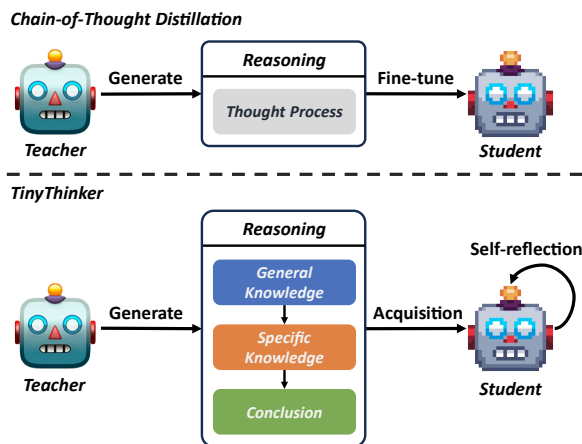


Figure 1: Comparison between TinyThinker and standard Chain-of-Thought Distillation. **Top:** Fine-tuning the student directly on teacher-generated reasoning data. **Bottom:** TinyThinker acquires reasoning capabilities through a three-stage process, further refined via self-reflection.

tasks, particularly in commonsense and arithmetic reasoning domains (Achiam et al., 2023).

A notable limitation of CoT prompting is its reliance on models with over 100 billion parameters, making it computationally expensive and less feasible for real-world deployment (Wei et al., 2022b). To address these limitations, prior works (Hsieh et al., 2023; Ho et al., 2023; Magister et al., 2023) have explored leveraging LLMs to generate reasoning data for knowledge distillation, thereby strengthening the reasoning capabilities of smaller language models (LMs)¹. Building on these advances, recent studies generally follow two main steps: (1) applying CoT prompting or its variants (Kojima et al., 2022) to generate reasoning data from a teacher model, and (2) designing specialized fine-tuning strategies to transfer reasoning ca-

¹In this study, "large" refers to models with over 100 billion parameters, whereas "small" denotes those with fewer than 1 billion parameters.

pabilities to a student model² (Shridhar et al., 2023; Jiang et al., 2023; Kang et al., 2024).

Although these methods have yielded promising results, it remains uncertain whether the student model genuinely develops reasoning capabilities. A recent study attributes the success of CoT prompting to its decomposition of compositional functions during in-context learning, enabling the model to focus on relevant data at each step and learn single-step composition functions in context (Li et al., 2024c). This interpretation aligns with the perspective that LLMs perform multi-step reasoning internally (Hou et al., 2023). However, when the student model lacks such intrinsic capabilities, direct training on synthesized reasoning data raises a critical concern: the student model may imitate the style of the synthesized data rather than its factuality (Gudibande et al., 2024). Given that reasoning typically involves multiple inference steps—each requiring explicit or implicit application of knowledge (Yu et al., 2023)—there is a risk that the student model might mimic these step-by-step reasoning processes superficially, without truly internalizing the underlying knowledge.

To enable flexible incorporation of internal knowledge for effective reasoning, we propose *TinyThinker*, a novel framework that structures the reasoning process into three stages, progressively refining knowledge from coarse to fine granularity (Figure 1). To facilitate the acquisition and refinement of this structured reasoning process, we introduce a two-phase training framework, consisting of a *reasoning acquisition* phase followed by a *self-reflection* phase.

This study focuses on commonsense reasoning, which inherently requires the integration of extensive knowledge with effective reasoning capabilities. Given that most commonsense reasoning datasets are structured in a multiple-choice question (MCQ) format (Talmor et al., 2019; Mihaylov et al., 2018; Geva et al., 2021), we designed a three-stage process inspired by human problem-solving strategies. In the reasoning acquisition phase, the model first *recalls* general knowledge relevant to the question and options. It then conducts an *analysis* of specific knowledge for each option, using the recalled general knowledge as context. Finally, the model integrates the acquired knowledge to *summarize* and identify the correct answer.

²In accordance with knowledge distillation terminology, LLMs are referred to as teacher models, and smaller LMs as student models.

After acquiring new reasoning capabilities, independent reflection is essential for consolidating these capabilities. In the self-reflection phase, the student revisits previous training data and generates new reasoning data through the same three-stage process. These self-generated data, combined with iterative Direct Preference Optimization (DPO) (Rafailov et al., 2024), further refine the reasoning capabilities, which fosters a deeper internalization of the underlying knowledge.

Experimental results on three commonsense reasoning datasets (CommonsenseQA, OpenBookQA, and StrategyQA) demonstrate that *TinyThinker* significantly enhances the reasoning capabilities of the student model. Furthermore, the scalability of our approach is validated by consistent performance improvements as the model size increases. Additionally, ablation studies on both the three-stage reasoning process and the self-reflection phase further underscore the effectiveness of the coarse-to-fine reasoning approach, along with the model’s iterative self-reflection capabilities through the generation of internal knowledge.

2 Related Work

Knowledge distillation transfers knowledge from a larger teacher model to a smaller student model (Gou et al., 2021), enabling the student to leverage the features learned by the teacher (Hinton, 2015). As LLMs advance in generating high-quality data, this process has been extended to transfer reasoning capabilities via teacher-generated data. Recent research in this area can be categorized into five main approaches: supervised fine-tuning, decomposer-solver framework, feedback-based framework, retriever-augmented framework, and self-improvement framework.

Supervised fine-tuning This method trains the student model exclusively on teacher-generated reasoning data, typically employing a language modeling objective. It has become a standard approach for reasoning distillation (Ho et al., 2023; Magister et al., 2023; Fu et al., 2023; Wang et al., 2023a; Li et al., 2023). Alternatively, reasoning generation is combined with direct answer prediction within a multi-task learning framework, which has been shown to yield further performance improvements (Hsieh et al., 2023; Chen et al., 2024; Li et al., 2024a). However, these methods rely solely on teacher-generated data and lack an explicit focus on the integration of reasoning with underlying

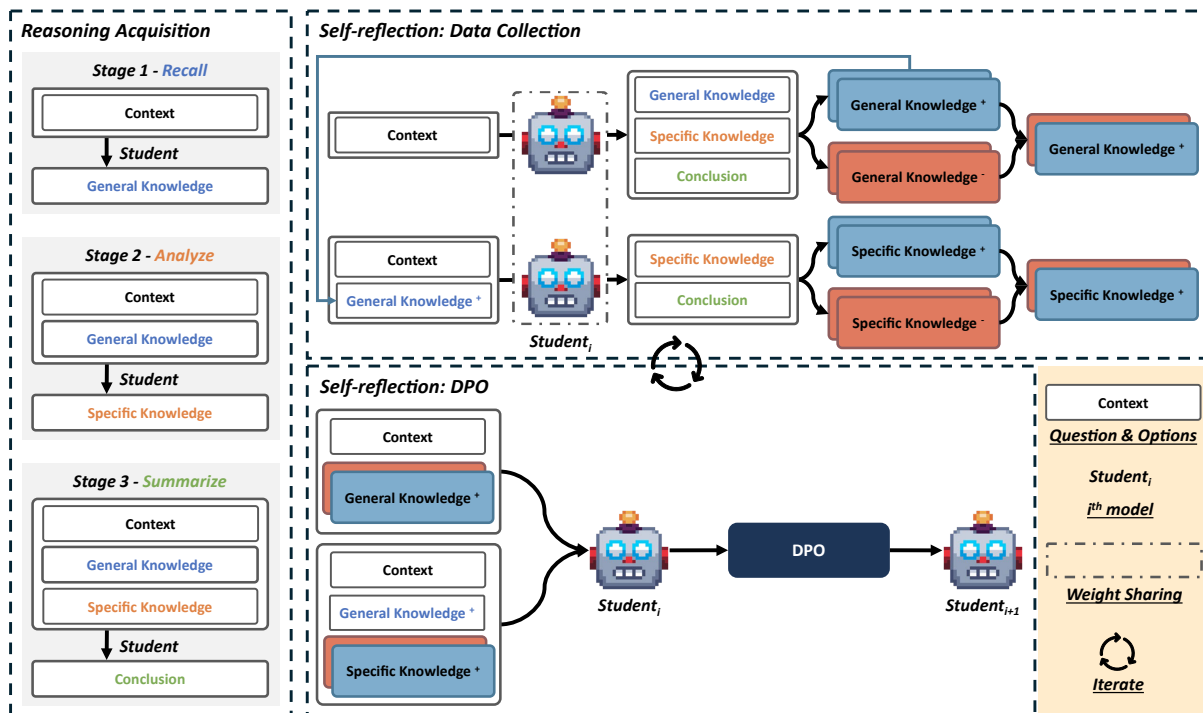


Figure 2: Detailed process of TinyThinker. **Reasoning Acquisition:** The student model follows a recall-analyze-summarize process, refining reasoning from coarse to fine granularity. **Self-reflection:** The model iteratively collects data and applies DPO. Pairwise data is first collected during the recall stage, and the preferred data from this stage informs the collection of pairwise data in the analyze stage. Once sufficient data is gathered, DPO is applied to refine the student’s reasoning capabilities, facilitating progression to the next iteration of self-reflection.

knowledge.

Decomposer-solver framework This framework trains two student models: one responsible for decomposing complex questions into simpler sub-questions, and another tasked with solving these sub-questions. The approach can be implemented in a single-step process, where both decomposition and solution occur in one iteration (Shridhar et al., 2023), or as a turn-based approach, where decomposition and solving are repeated iteratively until a complete solution is reached (Han et al., 2023). While this method effectively simplifies complex reasoning questions, it primarily focuses on semantic decomposition.

Feedback-based framework In this framework, the student model is initially trained to generate reasoning processes. When the student produces an incorrect answer, the associated reasoning process is sent back to the teacher, which provides feedback on the errors. This feedback is incorporated into subsequent rounds of fine-tuning, creating an iterative cycle of corrective guidance that gradually improves the student’s reasoning capabilities (Jiang et al., 2023; Wang et al., 2023c). While this

approach enables the student to address its deficiencies through targeted feedback, the reliance on the teacher limits the student to independently infer connections between reasoning and underlying knowledge.

Retriever-augmented framework In contrast to the feedback-based framework, which focuses on refining reasoning through direct guidance, the retriever-augmented framework improves reasoning by integrating external information. Specifically, a retriever is employed to search for relevant information from external sources, such as Wikipedia, and incorporated into the reasoning process to improve performance (Kang et al., 2024; Li et al., 2024b). Although this framework strengthens the student’s reasoning by supplementing it with external information, the retrieved information is not integrated into the student’s internal knowledge.

Self-improvement framework Unlike the feedback-based and retriever-augmented frameworks, which rely on external assistance from either the teacher model or knowledge bases, the self-improvement framework refines the

Stage	Data Format
Data Curation	Prompt: Question Options: (A)... (B)... (C)... (D)... \n Key Information: <>\n Explanations: A is correct. Because <> B is incorrect. Because <> C is incorrect. Because <> D is incorrect. Because <> (Note: "<>" denotes knowledge to be generated by the teacher model.)
Recall	Input: Question Options: (A)... (B)... (C)... (D)... \n Recall: Label: General Knowledge
Analyze	Input: Question Options: (A)... (B)... (C)... (D)... \n Recall: General Knowledge. \n Analyze: For option A, Label: Specific Knowledge (Note: Apply the same format for option B, C, D)
Summarize	Input: Question Options: (A)... (B)... (C)... (D)... \n Recall: General Knowledge. \n Analyze: For option A, Specific Knowledge. For option B, Specific Knowledge. For option C, Specific Knowledge. For option D, Specific Knowledge. \n Summarize: Label: Summary

Table 1: Data format across various stages of TinyThinker: data curation, recall, analyze, and summarize stages.

reasoning capabilities through self-generated data (Liu et al., 2023). In this approach, the student initially learns basic reasoning capabilities from the teacher. It then generates both correct and incorrect relevant knowledge, which is used as paired data for reinforcement learning. While this approach underscores the complementarity between knowledge and reasoning, its primary focus is on enabling the student to generate and utilize its internal knowledge effectively, rather than on improving the reasoning process itself.

Our approach extends the self-improvement framework with two key innovations: (1) a novel three-stage reasoning process that equips the student model with foundational reasoning capabilities by internalizing underlying knowledge, and (2) a self-reflection method that iteratively applies the DPO approach to further refine the relationship between reasoning capabilities and underlying knowledge.

3 Methodology

To enable the student to internalize the underlying knowledge for effective reasoning, we propose a two-phase training approach (Figure 2): reasoning acquisition (§3.1) and self-reflection (§3.2). During the reasoning acquisition phase, the student model progresses through three stages—recall, analyze, and summarize—which progressively refine its reasoning from general to specific knowledge. In the subsequent self-reflection phase, the student generates multiple reasoning processes based on its learned capabilities to further refine the recall and analyze stages with DPO.

3.1 Reasoning Acquisition

Data Curation The initial dataset consists of a question q and options $O = \{o_1, o_2, \dots, o_n\}$. To generate data for each stage, we utilized a few-shot prompting method with 7 to 8 examples, adapted from prior work (Wei et al., 2022b; Wang et al., 2023b; Li et al., 2024a) and tailored to our task (Table 1). Following Magister et al. (2023), we explicitly indicate the correctness of each option, encouraging the teacher model to generate accurate knowledge. Full prompt details are provided in Appendix C.

Once sufficient data had been collected, they were reorganized for each stage as shown in Table 1. Specifically, stage-specific cue phrases such as "Recall:", "Analyze: For Option", and "Summarize:" were introduced to guide the student model, parameterized by θ , in generating general knowledge k_{general} , option-specific knowledge k_{specific} , and a final summary k_{summary} .

Recall Stage The student performs a cursory scan of the question and options, retrieving general background knowledge. This broad knowledge serves as a foundation for subsequent reasoning processes. The training objective for this stage is expressed as:

$$\mathcal{L}_{\text{recall}}(\theta) = -\log P(k_{\text{general}} \mid q, O; \theta) \quad (1)$$

Analyze Stage After recalling general knowledge, the student carefully evaluates each option $o_i \in O$, generating specific knowledge relevant to each option. Although the model may infer the correctness of certain options, it can be misled by confounding alternatives, necessitating a detailed analysis to reach a reliable summary. The training

objective for this stage is formalized as:

$$\mathcal{L}_{\text{analyze}}(\theta) = -\log P(k_{\text{specific}} | q, O, k_{\text{general}}; \theta) \quad (2)$$

Summarize Stage At this stage, all previously generated knowledge is fed back into the model, enabling it to derive a final summary and select the correct answer. The training objective for this stage is defined as:

$$\begin{aligned} \mathcal{L}_{\text{summary}}(\theta) \\ = -\log P(k_{\text{summary}} | q, O, k_{\text{general}}, k_{\text{specific}}; \theta) \end{aligned} \quad (3)$$

3.2 Self-Reflection

Iterative DPO Direct Preference Optimization (DPO) aligns models with human preferences using pairwise comparisons $\mathcal{D} = \{x_i, y_i^w, y_i^l\}_{i=1}^N$, where x_i is the input to the model π_θ , and y_i^w and y_i^l are outputs generated by the model. Human evaluators identify one output as "preferred" (y^w) and the other as "dispreferred" (y^l).

In this study, we adapt DPO to refine the student's reasoning capabilities by aligning it with its internal knowledge. We employ an iterative DPO approach (Yuan et al., 2024; Pang et al., 2024) to facilitate the self-reflection process. Specifically, this approach trains a series of models $\{\pi_1, \dots, \pi_T\}$, where each model at iteration t is trained on new pairwise data generated by the model from the preceding iteration ($t - 1$). At each iteration, the model π_θ is initialized with the parameters of the $(t - 1)^{\text{th}}$ model.

To further stabilize training, an additional negative log-likelihood (NLL) loss term is applied to the preferred data, following Pang et al. (2024) and Dubey et al. (2024). The overall training objective for DPO combined with NLL is defined as:

$$\begin{aligned} \mathcal{L}_{\text{DPO+NLL}}(\pi_t; \pi_\theta) \\ = \mathcal{L}_{\text{DPO}}(y_i^w, y_i^l | x_i) + \alpha \mathcal{L}_{\text{NLL}}(y_i^w | x_i) \\ = -\log \sigma \left(\beta \log \frac{\pi_t(y_i^w | x_i)}{\pi_\theta(y_i^w | x_i)} - \beta \log \frac{\pi_t(y_i^l | x_i)}{\pi_\theta(y_i^l | x_i)} \right) \\ - \alpha \log \pi_t(y_i^w | x_i) \end{aligned} \quad (4)$$

Data Collection Unlike standard DPO, which leverages continuous rewards to determine preferred outputs, we adopt a binary approach that classifies the generated outputs as either correct or incorrect based on their corresponding summaries.

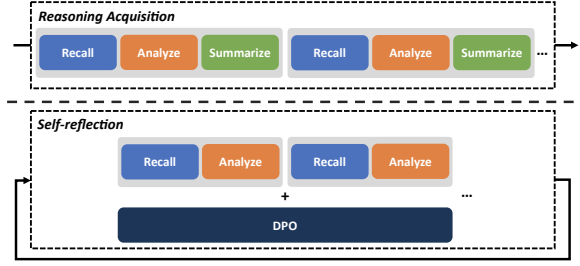


Figure 3: Overall process of the training strategy. **Top:** During the reasoning acquisition phase, the recall-analyze-summarize process is repeated iteratively. **Bottom:** During the self-reflection phase, the recall-analyze process is iterated with DPO.

This assessment directly reflects the accuracy of the model's underlying knowledge: a correct summary indicates the presence of accurate knowledge, which warrants prioritization for further refinement. In contrast, an incorrect summary reveals deficiencies in relevant knowledge and signals the need for improvement.

In our approach, knowledge is generated primarily during the recall and analyze stages. To strengthen each stage individually, we collect data by following the three-stage process twice, with different generation strategies for each pass. In the first data collection, temperature sampling is used in the recall stage to generate diverse knowledge, while greedy search is applied in the analyze and summarize stages to minimize influence on the final summary. Since the analyze stage depends on the outputs from the recall stage, the second data collection starts directly at the analyze stage, utilizing the "preferred" general knowledge gathered from the first collection. Similarly, temperature sampling is applied to the analyze stage, and greedy search is employed for the summarize stage.

3.3 Training Strategy

To effectively apply the three-stage reasoning process to novel questions, we repeat this process multiple times per epoch during the reasoning acquisition phase (Figure 3). Each iteration consists of a predetermined number of training steps per stage, with the model advancing sequentially through the stages. Since the analyze stage requires generating option-specific knowledge, the corresponding dataset is proportionally larger, requiring more training steps compared to the recall and summarize stages.

After learning the three-stage reasoning process, the student transitions into the self-reflection phase

to further refine its reasoning capabilities. This phase retains the iterative structure but differs in two key aspects: (1) the summarize stage is excluded since only the recall and analyze stages require further refinement, and (2) the analyze stage now focuses on pairwise comparisons for each option individually, leading to an equal number of training steps for both recall and analyze stages.

4 Experimental Setup

4.1 Dataset

We evaluate the student model on three common-sense reasoning benchmarks: CommonsenseQA (CSQA) (Talmor et al., 2019), OpenBookQA (OBQA) (Mihaylov et al., 2018), and StrategyQA (Geva et al., 2021). These benchmarks require substantial knowledge and reasoning capabilities to achieve competitive performance. A detailed description of each dataset is provided in Appendix B.

4.2 Baselines

We compare the performance of TinyThinker with several existing reasoning distillation approaches that use the same datasets and student model architectures:

Fine-tune-CoT (Ho et al., 2023) applies standard supervised fine-tuning to train the student model to generate CoT-style reasoning. Building on this, **DSS** (Hsieh et al., 2023) and **MT-CoT** (Li et al., 2024a) incorporate multi-task learning to train the student model to generate reasoning while simultaneously predicting the correct answer. **MI Distillation** (Chen et al., 2024) further improves multi-task learning by maximizing the mutual information between reasoning generation and answer prediction. The retrieval-augmented method **KARD** trains a retriever to assist the reasoning generation process directly, while **D&R Distillation** trains two separate student models: one to decompose questions into subquestions and another to answer these subquestions using an external knowledge base. Finally, **Crystal** (Liu et al., 2023) generates substantial amounts of relevant knowledge and incorporates it into the reasoning process simultaneously, then applies reinforcement learning to encourage the student model to generate more relevant knowledge.

4.3 Implementation Details

Teacher Model We employ GPT-4o (gpt-4o-2024-05-13) as the teacher model, accessed via the

OpenAI API³. The teacher model generates reasoning data for each question in the dataset following the three-stage reasoning process as outlined in our prompt design.

Student Model We utilize the T5 models—Small (60M), Base (220M), and Large (770M) (Raffel et al., 2020)—as the backbone for the student model. All models are trained using 4 A100 GPUs via the Huggingface library (Wolf, 2019), with pre-trained weights from publicly available sources⁴. Detailed hyperparameter settings are provided in Appendix A.

5 Results and Analysis

5.1 Overall Performance

We applied TinyThinker to train all student models across the datasets, with accuracy (%) as the primary evaluation metric. The performance results, summarized in Table 2, are organized by student model size and compared against baseline methods that trained the same student models using different approaches on the same datasets.

TinyThinker consistently achieves the best performance on both the OBQA and StrategyQA datasets, surpassing the best-performing baseline by margins of 1.5% to 2% on OBQA and 3% to 7% on StrategyQA across all model sizes. This demonstrates the advantage of internalizing underlying knowledge for effective reasoning.

Although TinyThinker does not outperform MT-CoT on the CSQA dataset, it outperforms MT-CoT on both the OBQA and StrategyQA datasets. In contrast, TinyThinker maintains a consistent performance gap of approximately 5% compared to DSS and MI Distillation on the CSQA dataset. Notably, all three methods—MT-CoT, DSS, and MI Distillation—utilize multi-task learning to develop reasoning capabilities. However, directly predicting answers in this approach may lead the student to rely on spurious correlations between question and options, potentially exploiting reasoning shortcut to arrive at the answer (Wang et al., 2023a, 2022).

5.2 Performance Across Model Sizes

We evaluated TinyThinker’s performance across different model sizes on the CSQA and StrategyQA

³<https://platform.openai.com/docs/model/gpt-4o>

⁴<https://huggingface.co/google-t5>

Model	Methods	CSQA	OBQA	StrategyQA ^a
<i>T5-Small (60M)</i>	Fine-tune-CoT (Ho et al., 2023)	29.48	-	56.04
	DSS (Hsieh et al., 2023)	43.24	-	-
	MI Distillation (Chen et al., 2024)	43.90	-	-
	MT-CoT (Li et al., 2024a)	49.17	51.72	-
	D&R Distillation (Li et al., 2024b)	-	-	55.00
	TinyThinker	46.36	53.60	60.26
<i>T5-Base (220M)</i>	Fine-tune-CoT (Ho et al., 2023)	45.37	-	59.68
	DSS (Hsieh et al., 2023)	63.29	-	-
	KARD (Kang et al., 2024)	-	59.33	56.57
	MI Distillation (Chen et al., 2024)	63.88	-	-
	MT-CoT (Li et al., 2024a)	64.50	60.68	61.05
	D&R Distillation (Li et al., 2024b)	-	-	59.00
TinyThinker	59.79	62.40	66.38	
<i>T5-Large (770M)</i>	Fine-tune-CoT (Ho et al., 2023)	54.22	-	62.15
	DSS (Hsieh et al., 2023)	70.43	-	-
	KARD (Kang et al., 2024)	-	66.40	66.04
	Crystal (Liu et al., 2023)	70.52	64.20	-
	MT-CoT (Li et al., 2024a)	74.37	64.60	-
	D&R Distillation (Li et al., 2024b)	-	-	63.30
TinyThinker	65.44	68.80	69.00	

^a Due to varying data splits in StrategyQA across papers, the results are for reference only.

Table 2: Accuracy (%) of the student model across baselines. **Bold** values indicate the best performance.

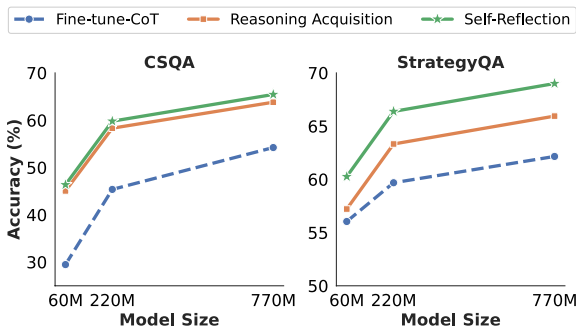


Figure 4: Accuracy (%) on CSQA and StrategyQA datasets across different model sizes.

datasets, comparing it with the Fine-tune-CoT baseline. As illustrated in Figure 4, TinyThinker consistently outperforms Fine-tune-CoT across all model sizes, indicating the advantages of the proposed three-stage reasoning process over the standard fine-tuning approach. Both the reasoning acquisition and self-reflection phases improve with increasing model size, demonstrating the scalability of our approach. Furthermore, the self-reflection phase consistently enhances the performance of the reasoning acquisition phase, confirming that refinement through self-generated data strengthens the student’s reasoning capabilities.

5.3 Ablation Study

The effect of three-stage process TinyThinker employs a structured reasoning process to progressively refine the student’s knowledge. To evaluate the necessity of this stage-wise refinement, we conducted the following experiments: In **Summarize**, the student directly derives the summary from the question and options; in **Recall-Summarize**, the model first generates general knowledge during the recall stage, which is then used alongside the question and options to infer summary; in **Analyze-Summarize**, the model first generates specific knowledge during the analyze stage and combines it with the question and options to reach a summary.

As indicated in Figure 5, performance generally improves with an increasing amount of available knowledge, particularly on the CSQA dataset, which demonstrates the efficacy of the reasoning process in progressively enhancing knowledge from coarse to fine granularity. Conversely, for the OBQA and StrategyQA datasets, the optimal performance for the T5-small and T5-base models was achieved using only general knowledge. This can be attributed to the model’s parameter size, which limits its capacity to manage the increasing

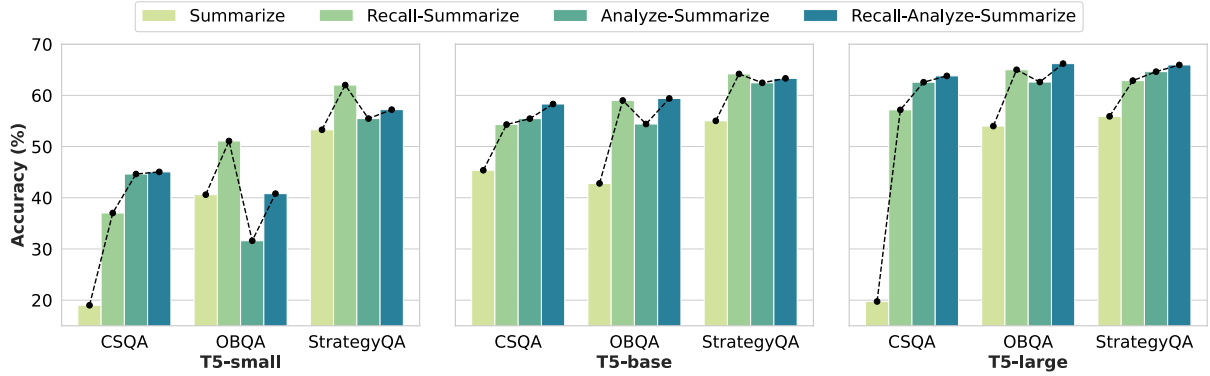


Figure 5: Ablation study on the effects of each stage in the recall-analyze-summarize reasoning process.

Model	Recall	Analyze	CSQA	OBQA	StrategyQA
<i>T5-Small (60M)</i>	✗	✗	45.05	40.80	57.21
	✓	✗	45.54 (+0.50)	45.20 (+4.40)	58.08 (+0.87)
	✗	✓	46.02 (+0.97)	50.00 (+9.20)	59.39 (+2.18)
	✓	✓	46.36 (+1.31)	53.60 (+12.80)	60.26 (+3.05)
<i>T5-Base (220M)</i>	✗	✗	58.31	59.40	63.32
	✓	✗	58.72 (+0.41)	60.19 (+0.79)	64.63 (+1.31)
	✗	✓	59.54 (+1.31)	61.80 (+2.40)	65.07 (+1.75)
	✓	✓	59.79 (+1.48)	62.40 (+3.00)	66.38 (+3.06)
<i>T5-Large (770M)</i>	✗	✗	63.80	66.20	65.94
	✓	✗	64.13 (+0.33)	66.80 (+0.60)	65.94 (+0.00)
	✗	✓	65.19 (+1.39)	67.40 (+1.20)	68.12 (+2.18)
	✓	✓	65.44 (+1.64)	68.80 (+2.60)	69.00 (+3.06)

Table 3: Ablation study on the effect of applying DPO at recall and analyze stages. ✓ and ✗ indicate whether DPO is applied. All ✗ symbols for both stages denote performance after the reasoning acquisition phase, whereas individual ✓ symbols correspond to the recall-DPO and analyze-DPO setting, respectively. **Bold** values indicate the best performance, and values in parentheses represent the performance change after applying DPO in each setting.

complexity of later stages, particularly during the analyze and summarize stages. While generating general knowledge requires a broad understanding of the question and options, generating specific knowledge requires a more detailed assessment of each option, significantly increasing the computational complexity of the later stages.

Consequently, as model size increases, the performance gap between Recall-Summarize and Analyze-Summarize gradually narrows, leading to improved performance on Recall-Analyze-Summarize. This trend is observable across both the OBQA and StrategyQA datasets.

The effect of self-reflection After training the student model to reason through the three-stage process, we applied the DPO algorithm to both the recall and analyze stages to further consolidate its reasoning capabilities. This approach has two key objectives: reinforcing correct knowledge and refining incorrect knowledge. In this experiment,

we investigate the performance gains of applying DPO individually to the recall and analyze stages, denoted as **recall-DPO** and **analyze-DPO**, respectively.

As shown in Table 3, applying DPO yielded performance improvements at both stages, with the most substantial gains observed when DPO was applied concurrently to both the recall and analyze stages. Notably, the improvement in the analyze stage was more pronounced than in the recall stage, supporting our earlier observation that the analyze stage is inherently more challenging to learn. This suggests that further performance gains could be achieved by refining learning strategies or increasing model size to better handle the complexities of the analyze stage.

6 Conclusion

In this study, we introduced TinyThinker to enhance reasoning capabilities through effective knowledge internalization. In contrast to prior

methods, we developed a structured three-stage reasoning process that progressively refines knowledge from coarse to fine granularity. This process is complemented by a two-phase training approach, consisting of reasoning acquisition and self-refinement phases. Experiments on common-sense reasoning benchmarks demonstrate that TinyThinker achieves superior performance compared to existing baselines. Additionally, the ablation study further confirms the contributions of each component, highlighting the overall effectiveness of our approach. We expect that TinyThinker provides a flexible framework extendable to other knowledge-intensive reasoning tasks, offering a promising strategy for developing effective reasoning capabilities in smaller LMs.

7 Limitation

Quality of curated data Although LLMs are capable of generating semantically coherent sentences, their inherent issue of hallucination sometimes leads to content that lacks factual accuracy and safety (Ji et al., 2023). This lack of factual accuracy also affects reasoning-related content generation, particularly through the "error cascade" problem, where a factual error in an intermediate reasoning step propagates inaccuracies through subsequent steps (Chu et al., 2024). Therefore, ensuring the factual quality of the data generated by the teacher model remains a challenge in this study.

Efficient generation strategy While the proposed three-stage reasoning process is effective for multiple-choice tasks, there is room for improving efficiency, especially in the analyze stage. Currently, the model generates specific knowledge independently for each option, leading to a time-intensive process. A more efficient approach would generate specific knowledge for all options in parallel, thereby accelerating the analyze stage.

Acknowledgments

This research was supported by the National Research Foundation (NRF) funded by the Korean government (MSIT) (No. RS-2023-00229822). We sincerely thank Shengen Piao, Jieun Lee, and Huijun Jin for their constructive suggestions and discussions, which significantly contributed to improving this work.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Xin Chen, Hanxian Huang, Yanjun Gao, Yi Wang, Jishen Zhao, and Ke Ding. 2024. [Learning to maximize mutual information for chain-of-thought distillation](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 6857–6868, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Zheng Chu, Jingchang Chen, Qianglong Chen, Weijiang Yu, Tao He, Haotian Wang, Weihua Peng, Ming Liu, Bing Qin, and Ting Liu. 2024. [Navigate through enigmatic labyrinth a survey of chain of thought reasoning: Advances, frontiers and future](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1173–1203, Bangkok, Thailand. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Yao Fu, Hao Peng, Litu Ou, Ashish Sabharwal, and Tushar Khot. 2023. [Specializing smaller language models towards multi-step reasoning](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 10421–10430. PMLR.
- Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. [Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies](#). *Transactions of the Association for Computational Linguistics*, 9:346–361.
- Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. 2021. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129(6):1789–1819.

- Arnav Gudibande, Eric Wallace, Charlie Victor Snell, Xinyang Geng, Hao Liu, Pieter Abbeel, Sergey Levine, and Dawn Song. 2024. [The false promise of imitating proprietary language models](#). In *The Twelfth International Conference on Learning Representations*.
- Chengcheng Han, Xiaowei Du, Che Zhang, Yixin Lian, Xiang Li, Ming Gao, and Baoyuan Wang. 2023. [DiCoT meets PPO: Decomposing and exploring reasoning paths in smaller language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 8055–8068, Singapore. Association for Computational Linguistics.
- Geoffrey Hinton. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Namgyu Ho, Laura Schmid, and Se-Young Yun. 2023. [Large language models are reasoning teachers](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14852–14882, Toronto, Canada. Association for Computational Linguistics.
- Yifan Hou, Jiaoda Li, Yu Fei, Alessandro Stolfo, Wangchunshu Zhou, Guangtao Zeng, Antoine Bosselut, and Mrinmaya Sachan. 2023. [Towards a mechanistic interpretation of multi-step reasoning capabilities of language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4902–4919, Singapore. Association for Computational Linguistics.
- Cheng-Yu Hsieh, Chun-Liang Li, Chih-kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alex Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. 2023. [Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8003–8017, Toronto, Canada. Association for Computational Linguistics.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38.
- Yuxin Jiang, Chunkit Chan, Mingyang Chen, and Wei Wang. 2023. [Lion: Adversarial distillation of proprietary large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3134–3154, Singapore. Association for Computational Linguistics.
- Minki Kang, Seanie Lee, Jinheon Baek, Kenji Kawaguchi, and Sung Ju Hwang. 2024. Knowledge-augmented reasoning distillation for small language models in knowledge-intensive tasks. *Advances in Neural Information Processing Systems*, 36.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.
- Liunian Harold Li, Jack Hessel, Youngjae Yu, Xiang Ren, Kai-Wei Chang, and Yejin Choi. 2023. [Symbolic chain-of-thought distillation: Small models can also “think” step-by-step](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2665–2679, Toronto, Canada. Association for Computational Linguistics.
- Shiyang Li, Jianshu Chen, yelong shen, Zhiyu Chen, Xinlu Zhang, Zekun Li, Hong Wang, Jing Qian, Baolin Peng, Yi Mao, Wenhui Chen, and Xifeng Yan. 2024a. [Explanations from large language models make small reasoners better](#). In *2nd Workshop on Sustainable AI*.
- Xiang Li, Shizhu He, Fangyu Lei, JunYang JunYang, Tianhuang Su, Kang Liu, and Jun Zhao. 2024b. [Teaching small language models to reason for knowledge-intensive multi-hop question answering](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 7804–7816, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Yingcong Li, Kartik Sreenivasan, Angeliki Gianou, Dimitris Papailiopoulos, and Samet Oymak. 2024c. Dissecting chain-of-thought: Compositionality through in-context filtering and learning. *Advances in Neural Information Processing Systems*, 36.
- Jiacheng Liu, Ramakanth Pasunuru, Hannaneh Hajishirzi, Yejin Choi, and Asli Celikyilmaz. 2023. [Crystal: Introspective reasoners reinforced with self-feedback](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 11557–11572, Singapore. Association for Computational Linguistics.
- Lucie Charlotte Magister, Jonathan Mallinson, Jakub Adamek, Eric Malmi, and Aliaksei Severyn. 2023. [Teaching small language models to reason](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1773–1781, Toronto, Canada. Association for Computational Linguistics.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. [Can a suit of armor conduct electricity? a new dataset for open book question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2381–2391, Brussels, Belgium. Association for Computational Linguistics.

- Richard Yuanzhe Pang, Weizhe Yuan, Kyunghyun Cho, He He, Sainbayar Sukhbaatar, and Jason Weston. 2024. Iterative reasoning preference optimization. *arXiv preprint arXiv:2404.19733*.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- Kumar Shridhar, Alessandro Stolfo, and Mrinmaya Sachan. 2023. [Distilling reasoning capabilities into smaller language models](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 7059–7073, Toronto, Canada. Association for Computational Linguistics.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A question answering challenge targeting commonsense knowledge](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Peifeng Wang, Aaron Chan, Filip Ilievski, Muhao Chen, and Xiang Ren. 2022. Pinto: Faithful language reasoning using prompt-generated rationales. *arXiv preprint arXiv:2211.01562*.
- Peifeng Wang, Zhengyang Wang, Zheng Li, Yifan Gao, Bing Yin, and Xiang Ren. 2023a. [SCOTT: Self-consistent chain-of-thought distillation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5546–5558, Toronto, Canada. Association for Computational Linguistics.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023b. [Self-consistency improves chain of thought reasoning in language models](#). In *The Eleventh International Conference on Learning Representations*.
- Zhaoyang Wang, Shaohan Huang, Yuxuan Liu, Jiahai Wang, Minghui Song, Zihan Zhang, Haizhen Huang, Furu Wei, Weiwei Deng, Feng Sun, and Qi Zhang. 2023c. [Democratizing reasoning ability: Tailored learning from large language model](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1948–1966, Singapore. Association for Computational Linguistics.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022a. [Emergent abilities of large language models](#). *Transactions on Machine Learning Research*. Survey Certification.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022b. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- T Wolf. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Fei Yu, Hongbo Zhang, Prayag Tiwari, and Benyou Wang. 2023. Natural language reasoning, a survey. *ACM Computing Surveys*.
- Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Xian Li, Sainbayar Sukhbaatar, Jing Xu, and Jason E Weston. 2024. [Self-rewarding language models](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 57905–57923. PMLR.

A Hyperparameters

The hyperparameters utilized in this study are summarized in Table 4.

B Datasets

CSQA CSQA is a multiple-choice question-answering dataset with five answer options per question, which requires diverse commonsense knowledge. Since the test set is not publicly available, we report performance on the validation set, following Li et al. (2024a); Wang et al. (2023a); Hsieh et al. (2023). Additionally, we randomly sampled 1,221 questions from the training set to form a development set, resulting in a final split of 8,520/1,221/1,221 questions for the training/validation/test sets.

OBQA OBQA is a four-choice question-answering dataset designed to evaluate the ability to apply broad common knowledge, particularly for elementary-level science questions. The dataset contains 4,957/500/500 questions for the training/validation/test sets.

StrategyQA StrategyQA is a binary (yes/no) question-answering dataset that requires implicit reasoning across diverse topics. The training set

contains 2,290 questions, while the test set includes 490 questions. As the test set is not publicly available, we split the training set into 80% for training, 10% for validation, and 10% for testing, following the procedure outlined in [Magister et al. \(2023\)](#) to ensure reproducibility.

To enhance the diversity of teacher-generated data during data curation, we utilize temperature sampling to generate 4-8 distinct reasoning processes per question to ensure diversity in the generated data. As demonstrated by SCoTD ([Li et al., 2023](#)), diverse reasoning data is crucial for effective reasoning distillation. After generation, we apply filtering procedures, such as de-duplication, to ensure data quality. The statistics of the curated datasets are summarized in [Table 5](#).

C Full Prompts

The prompt design follows a consistent instruction template across all datasets, as illustrated in [Table 6](#). This instruction provides foundational guidance for the teacher model, facilitating effective data generation. Subsequently, few-shot examples for each dataset are listed in [Table 7](#), [Table 8](#), and [Table 9](#).

Hyperparameter	Value	Note
Data Curation		
max_tokens	256	Maximum tokens generated per input.
n	[4, 8]	Number of generations for each input message.
temperature	0.8	Temperature value for sampling.
Reasoning Acquisition		
epochs	10	Total number of training epochs.
batch_size	64	Batch size for training.
interval	100	Default number of steps for recall, analyze, summarize stage.
lr	5×10^{-4}	Learning rate of AdamW optimizer.
Self-Reflection		
iterations	5	Number of iterations for DPO training.
n	10	Generation number for pairwise data collection.
temperature	0.7	Temperature value for sampling.
epochs	10	Epochs for each DPO iteration.
batch_size	64	Batch size for DPO training.
lr	5×10^{-6}	Learning rate of AdamW optimizer.
β	0.5	Weight of beta value.
α	0.5	Weight of NLL lose.

Table 4: Hyperparameter settings used in the experiments.

Dataset	# Recall	# Analyze	# Summarize	# Validation	# Test
CSQA	66,739	334,331	66,869	1,221	1,221
OBQA	19,444	77,985	19,497	500	500
StrategyQA	7,089	14,290	7,145	229	229

Table 5: Dataset statistics following data curation.

INSTRUCTION
<p>Reference examples delimited with ### as a guide to outline the reasoning steps towards the correct answer in a step-by-step manner. First, concisely summarize the key information or core issue from the question statement. Then, systematically evaluate each option according to the following format:</p> <p>Key Information: [Summarize the key information from the question.] Explanations: [Option] is [correct/incorrect]. Because [Explain the option’s features and relevance to the key information.]</p> <p>Ensure explanations are clear, concise, and diverse by using varied logic, examples, and perspectives. Be creative to prevent repetition and make sure each step leads to the summary. Highlighting the uniqueness of each example to support diverse reasoning.</p>

Table 6: The instruction of data curation.

PROMPT FOR CSQA
<p>1. The fox walked from the city into the forest, what was it looking for? Options: (A) pretty flowers (B) hen house (C) natural habitat (D) storybook Key Information: The motivations of a fox moving from an urban to a forest environment, focusing</p>

on natural instincts or necessities.

Explanations: A is incorrect. Because Foxes seeking pretty flowers does not align with their instinctual needs for survival, such as food or shelter.

B is incorrect. Because a hen house, usually near human habitats, does not match the forest setting the fox is moving towards.

C is correct. Because the fox's migration to the forest implies a search for its natural habitat, indicating a pursuit of basic needs and instincts.

D is incorrect. Because a storybook, as a human creation, doesn't meet any of a fox's natural survival instincts in a forest setting.

2. Sammy wanted to go to where the people were. Where might he go?

Options: (A) populated areas (B) race track (C) desert (D) apartment (E) roadblock

Key Information: Sammy seeks location that is likely to be frequented by or filled with individuals.

Explanations: A is correct. Because populated areas naturally have many people, aligning with Sammy's goal.

B is incorrect. Because a race track may be crowded during events but doesn't consistently draw people, not fully aligning with Sammy's aim.

C is incorrect. Because deserts are sparsely populated and do not meet Sammy's desire to be around people.

D is incorrect. Because while an apartment indicates residence, it doesn't broadly represent a gathering place for many people.

E is incorrect. Because roadblocks might temporarily gather people but aren't places people intentionally seek for socializing.

3. What do people use to absorb extra ink from a fountain pen?

Options: (A) shirt pocket (B) calligrapher's hand (C) inkwell (D) desk drawer (E) blotter

Key Information: People use a specific tool or material to absorb extra ink from a fountain pen.

Explanations: A is incorrect. Because a shirt pocket may catch ink accidentally but is not used intentionally for absorbing extra ink from a fountain pen.

B is incorrect. Because a calligrapher's hand might contact with ink but is not used purposefully to absorb extra ink.

C is incorrect. Because an inkwell is for storing ink, not absorbing excess ink from a pen.

D is incorrect. Because a desk drawer is for storage, not specifically for absorbing excess ink.

E is correct. Because a blotter is designed to absorb excess ink, preventing smudges and keeping the writing area clean.

4. Before getting a divorce, what did the wife feel who was doing all the work?

Options: (A) harder (B) anguish (C) bitterness (D) tears (E) sadness

Key Information: The emotional state of a wife who perceived imbalance of responsibilities before deciding on divorce, suggesting feelings stemming from stress, imbalance, or discontent.

Explanations: A is incorrect. Because 'harder' refers to the level of effort, not an emotional state, and doesn't match the emotional context of imbalance leading to divorce.

B is incorrect. Because 'anguish' indicates severe distress, it may not fully capture the specific feelings of imbalance and discontent in this scenario.

C is correct. Because 'bitterness' accurately reflects feelings of anger and disappointment from unmet expectations and perceived imbalance, aligning with the scenario.

D is incorrect. Because 'tears' indicate a physical manifestation of emotions but do not specify the emotional state related to the feeling of being overburdened.

E is incorrect. Because 'sadness' is a general emotion that doesn't precisely convey the sense of being undervalued or the specific discontent that led to considering divorce.

5. Where do you put your grapes just before checking out?

Options: (A) mouth (B) grocery cart (C) super market, (D) fruit basket (E) fruit market

Key Information: One would place grapes in the immediate moments before proceeding to the checkout in a shopping context, as part of the grocery purchasing process.

Explanations: A is incorrect. Because putting grapes in one's mouth before checkout implies eating them before purchase, which is inappropriate.

B is correct. Because the grocery cart is used for holding items to be purchased, right before checking out.

C is incorrect. Because the supermarket is the overall shopping location, not where one places items immediately before checkout.

D is incorrect. Because a fruit basket is typically for storing grapes at home, not for holding them just before checkout.

E is incorrect. Because the fruit market refers to the broader shopping venue, not the specific spot for grapes just before checkout.

6. Google Maps and other highway and street GPS services have replaced what?

Options: (A) United States (B) Mexico (C) countryside (D) atlas

Key Information: There is a transition from traditional navigation tools to modern navigation tools, focusing on what has been predominantly replaced by digital mapping and GPS services in terms of functionality.

Explanations: A is incorrect. Because the United States, being a country, cannot be replaced by digital mapping technologies.

B is incorrect. Because Mexico, as a country, cannot be replaced by digital mapping technologies, highlighting an incorrect match with navigation tools.

C is incorrect. Because the countryside, a type of geographic area, cannot be replaced by GPS services in the context of navigation tool replacement.

D is correct. Because an atlas, a book of maps, has been functionally replaced by digital mapping and GPS services, marking a shift in how people navigate.

7. What home entertainment equipment requires cable?

Options: (A) radio shack (B) substation (C) television (D) cabinet

Key Information: Devices used for entertainment purposes within a home setting that necessitate a connection through a physical cable for operation or functionality

Explanations: A is incorrect. Because radio shack refers to a place rather than a piece of home entertainment equipment requiring a cable.

B is incorrect. Because a substation pertains to electrical power distribution, not directly related to home entertainment or devices needing a cable in a home.

C is correct. Because televisions typically require a cable connection for signal reception, aligning with the key concept of home entertainment equipment needing a cable.

D is incorrect. Because a cabinet is used for storage and does not require a cable for home entertainment purposes, unrelated to the key concept.

8. The man laid on the soft moss and looked up at the trees, where was the man?

Options: (A) Niagra Falls (B) forest (C) waterfall (D) ground (E) tree

Key Information: The scenario describes a natural setting characterized by soft moss and an upward view of trees, suggesting a location where these features are prominent.

Explanations: A is incorrect. Because Niagra Falls is primarily associated with its waterfalls, not a setting of soft moss and tree views.

B is correct. Because a forest offers both soft moss to lie on and numerous trees to look up at, matching the scenario's description.

C is incorrect. Because while waterfalls can be found in forests, the specific mention of laying on soft moss and looking up at trees directly suggests a broader environment than just the area near a waterfall.

D is incorrect. Because while the man is technically on the ground, it is too general and misses the specific natural characteristics implied by the details of soft moss and trees.

E is incorrect. Because the scenario describes the man is looking up at trees, not situated in one, conflicting with the scenario given.

Table 7: The prompt for CSQA dataset.

PROMPT FOR OBQA

1. As a car approaches you in the night,

Options: (A) the headlights become more intense (B) the headlights recede into the dark (C) the headlights remain at a constant (D) the headlights turn off

Key Information: The intensity of headlights appears to increase as a car approaches due to the reduction in distance, allowing more light to reach the observer.

Explanations: A is correct. Because the intensity of headlights increases as the distance decreases, consistent with the concept of light appearing more intense as a car approaches.

B is incorrect. Because headlights become more visible, not less, as a car approaches, contradicting the concept of increasing intensity with proximity.

C is incorrect. Because the apparent brightness of headlights increases as the car comes closer, not remaining constant, diverging from the notion of constant intensity.

D is incorrect. Because headlights typically do not turn off as a car approaches, unrelated to the concept of light behavior with respect to distance.

2. What is the most likely to be an effect of acid rain on an aquatic environment?

Options: (A) decrease in plant life (B) increase in fish population (C) increase in plant growth (D) cleaner and clearer water

Key Information: Acid rain, containing harmful acidic compounds, detrimentally affects aquatic environments by altering water chemistry, harmful to aquatic life.

Explanations: A is correct. Because acid rain reduces water pH, harming aquatic plants, consistent with the concept that acid rain negatively affects aquatic life.

B is incorrect. Because acid rain decreases fish populations by introducing harmful chemicals, contrary to the idea of an increase.

C is incorrect. Because acid rain impedes plant growth due to its acidity, opposing the concept of promoting a healthy aquatic environment.

D is incorrect. Because acid rain pollutes water, reducing clarity, aligning with the concept of its harmful impacts on aquatic environments.

3. The moon's surface

Options: (A) is smooth on the entire surface (B) contains large cavities cause by explosions (C) contains an internal core of cheese (D) is filled with lakes

Key Information: The moon's surface is characterized by a variety of geological features, including craters from asteroid impacts.

Explanations: A is incorrect. Because the moon's surface is not smooth but covered in craters and rough terrain, contradicting the idea of uniform smoothness.

B is correct. Because the moon's surface features large craters caused by asteroid impacts, aligning with the concept of its varied geological features.

C is incorrect. Because the moon is made of rock, not cheese, contrasting with scientific evidence of its geological composition.

D is incorrect. Because the moon lacks lakes or liquid water, diverging from the key concept by misrepresenting its surface conditions.

4. When the weather changes as it does from Christmas to Easter,

Options: (A) the air may chill (B) the ground may freeze (C) the plants may die (D) the ground may warm

Key Information: The period from Christmas to Easter typically involves a transition from winter to spring in many parts of the world, marked by gradually increasing temperatures and thawing conditions.

Explanations: A is incorrect. Because the trend from Christmas to Easter is toward warmer weather, not colder, contrasting with the concept of seasonal warming.

B is incorrect. Because ground freezing aligns more with winter's onset; the period to Easter typically sees thawing, opposing the warming trend concept.

C is incorrect. Because the approach to Easter, signaling spring, is a time for plant life to begin thriving again, not dying, contradicting the concept of seasonal renewal.

D is correct. Because the ground warming from Christmas to Easter reflects the transition from winter to spring, consistent with the concept of increasing temperatures.

5. Heat and moisture in the ocean is a good recipe for

Options: (A) a violent storm (B) violent sea animals (C) condensation (D) inland storms

Key Information: Heat and moisture in the ocean contribute to the formation of storms by providing the energy and water vapor necessary for their development.

Explanations: A is correct. Because heat and moisture are key ingredients for storm development over the ocean, directly aligning with the concept of storm formation.

B is incorrect. Because the link between heat and moisture and 'violent sea animals' is not relevant to the meteorological focus of the key concept.

C is incorrect. Because while condensation is a component of storm formation, it does not fully encompass the broader process of storm development indicated by the key concept.

D is incorrect. Because heat and moisture contribute to storms that can affect inland areas, demonstrating the widespread impact of these conditions beyond the immediate ocean environment.

6. Hummingbirds take what with them

Options: (A) Bees (B) energy (C) Pollen (D) Honey

Key Information: As pollinators, hummingbirds interact with flowers for nectar, during which they inadvertently collect and transfer pollen.

Explanations: A is incorrect. Because hummingbirds do not transport bees, unrelated to their role as pollinators.

B is incorrect. Because while hummingbirds expend energy, saying they 'take energy' is too abstract and unrelated to their pollination activities.

C is correct. Because pollen is transferred by hummingbirds as they feed on nectar, aligning with their role as pollinators.

D is incorrect. Because hummingbirds consume nectar, not honey, which is unrelated to their interaction with flowers

7. What covers over 90% of the Earth's surface and 0% of the moon's surface

Options: (A) a magnesium iron silicate mineral (B) chemical element with the symbol S (C) the element with the symbol Fe (D) that which contains 2 hydrogen and 1 oxygen molecules

Key Information: A substance that is abundant on Earth's surface but nonexistent on the moon's

surface, pointing towards Earth's unique characteristic related to its surface coverage.
 Explanations: A is incorrect. Because magnesium iron silicate minerals are primarily found in Earth's mantle, not its surface, and have no relevance to the moon's surface.
 B is incorrect. Because Sulfur does not cover Earth's surface, making it irrelevant to the comparison with the moon.
 C is incorrect. Because iron (Fe) is common in both Earth's and the moon's compositions but does not cover Earth's surface in the context required by the question.
 D is correct. Because water, composed of 2 hydrogen and 1 oxygen molecules (H₂O), covers over 70% of Earth's surface but 0% of the moon's surface, aligning with Earth's unique surface characteristics.

Table 8: The prompt for OBQA dataset.

PROMPT FOR STRATEGYQA

1. Do hamsters provide food for any animals?

(A) yes (B) no

Key Information: Hamsters serve as prey for other animals in the food chain.

Explanations: A is correct. Because hamsters are prey for various predators, reflecting their role in the food chain and ecosystem interconnectedness.

B is incorrect. Because denying this ignores the reality of food chain dynamics and hamsters' role in supporting habitat biodiversity.

2. Could Brooke Shields succeed at University of Pennsylvania?

(A) yes (B) no

Key Information: The potential for success of an individual, Brooke Shields, at a specific academic institution, University of Pennsylvania, implying considerations of her capability, ambition, and the university's environment.

Explanations: A is correct. Because Brooke Shields' academic and career accomplishments suggest she would likely succeed at the University of Pennsylvania, reflecting her capability to thrive in demanding environments.

B is incorrect. Because there's ample evidence of Brooke Shields' ability and ambition, making her success at the University of Pennsylvania plausible, thus dismissing this ignores her proven capability.

3. Hydrogen's atomic number squared exceeds number of Spice Girls?

(A) yes (B) no

Key Information: Hydrogen's atomic number is 1, and the Spice Girls group consists of 5 members.

Explanations: A is incorrect. Because hydrogen's atomic number squared (1) does not exceed the Spice Girls' member count (5).

B is correct. Because since hydrogen's atomic number squared equals 1, it is less than the Spice Girls' 5 members, aligning with the key concept.

4. Is it common to see frost during some college commencements?

(A) yes (B) no

Key Information: College commencements in certain regions or during certain times of the year might coincide with colder weather conditions, which can lead to the formation of frost.

Explanations: A is correct. Because frost can occur during college commencements in colder

climates or seasonal transitions, reflecting geographic and seasonal weather variations.
B is incorrect. Because disregarding frost ignores the reality of seasonal cold weather affecting commencements in many regions.

5. Could a llama birth twice during War in Vietnam (1945-46)?

(A) yes (B) no

Key Information: The natural gestation period of a llama is approximately 11 months, and the specific duration of the War in Vietnam stated to be 6 months.

Explanations: A is incorrect. Because the llama's 11-month gestation period surpasses the 6-month Vietnam War duration, making two births impossible within this time frame.

B is correct. Because given the llama's gestation period and the war's duration, it's biologically impossible for a llama to give birth twice during the conflict.

6. Would a pear sink in water?

(A) yes (B) no

Key Information: The density of pear would determine if it sinks or floats in water.

Explanations: A is incorrect. Because pears float due to a density less than water, attributed to their fibrous and airy composition.

B is correct. Because pears float in water because of their air-filled fibrous structure, making them less dense than water.

7. Is Albany, Georgia the most populous US Albany?

(A) yes (B) no

Key Information: Comparing the population size between Albany, Georgia and Albany, New York.

Explanations: A is incorrect. Because Albany, New York, has a larger population than Albany, Georgia, making the latter not the most populous.

B is correct. Because Albany, New York, surpasses Albany, Georgia, in population, thereby making the Georgia city not the most populous Albany.

8. Can the Great Depression be treated with Prozac?

(A) yes (B) no

Key Information: Great Depression is a historical event and Prozac is a medication used to treat clinical depression.

Explanations: A is incorrect. Because Prozac treats clinical depression, not historical events like the Great Depression.

B is correct. Because the Great Depression, an economic crisis, cannot be treated with Prozac, a medication for clinical depression.

Table 9: The prompt for StrategyQA dataset.