

ChainLM: Empowering Large Language Models with Improved Chain-of-Thought Prompting

Xiaoxue Cheng¹, Junyi Li^{1,3}, Wayne Xin Zhao^{1*} and Ji-Rong Wen^{1,2}

¹Gaoling School of Artificial Intelligence, Renmin University of China

²School of Information, Renmin University of China

³DIRO, Université de Montréal

{chengxiaoxue, lijunyi}@ruc.edu.cn, batmanfly@gmail.com

Abstract

Chain-of-Thought (CoT) prompting can enhance the reasoning capabilities of large language models (LLMs), establishing itself as a primary approach to solving complex reasoning tasks. Existing CoT synthesis approaches usually focus on simpler reasoning tasks and thus result in low-quality and inconsistent CoT prompts. In response to this challenge, we present an empirical investigation of CoT prompting and introduce CoTGenius, a novel framework designed for the automatic generation of superior CoT prompts. CoTGenius is developed based on three major evolution strategies, *i.e.*, complicate, diversify, and specify—alongside two filtering mechanisms: evolutionary success judgement and correctness verification. We further employ CoTGenius to create an extensive CoT dataset, and subsequently fine-tune the Llama 2-Chat 7B and 13B models on this dataset. We call the resulting model ChainLM. To deal with the cumulative error issue in reasoning steps, we propose a step-level debating method, wherein multiple debaters discuss each reasoning step to arrive at the correct answer. Extensive experiments demonstrate that our ChainLM models exhibit enhanced proficiency in addressing a spectrum of complex reasoning problems compared to existing models. In addition, we conduct an in-depth analysis of the impact of data categories within CoTGenius on the model performance. We release our dataset and code at <https://github.com/RUCAIBox/ChainLM>.

Keywords: Chain-of-Thought, Large Language Models, Debating

1. Introduction

Large Language Models (LLMs) (Zhao et al., 2023) have recently made great progress on natural language understanding and generation, and have also shown great potential as general-purpose task solvers by following commands or prompts (Brown et al., 2020; Nakano et al., 2021; OpenAI, 2023). However, it is well known that LLMs often fall short in complex reasoning tasks, *e.g.*, mathematical reasoning and symbolic reasoning. Thus concerns about improving the reasoning capabilities of LLMs have drawn significant attention from the research community (Luo et al., 2023; Pan et al., 2023).

Amidst this backdrop, Chain-of-Thought (CoT) prompting has been proposed and emerged as an effective solution for complex reasoning (Wei et al., 2022), where LLMs incorporate a series of intermediate reasoning steps before inferring the final output. Kojima et al. (2022) simply add a phrase “*Let’s think step by step*” in prompts, enabling LLMs to conduct zero-shot CoT reasoning without any additional exemplars. However, the same benefits may not necessarily extend to relatively smaller LLMs. Although there have been some attempts towards *fine-tuning* LLMs with multi-step CoT prompting data to stimulate the step-by-step reasoning capacity in smaller models (Shridhar et al., 2022; Fu et al., 2023a; Kim et al., 2023), these approaches are

mainly focused on single reasoning tasks or rely on simple strategies to synthesize CoT prompting data. Fu et al. (2023b) found that existing datasets mostly focus on simple tasks involving fewer reasoning steps (*i.e.*, only 2 or 3 steps) or even omit certain intermediate steps, making the reasoning process incomplete and the model struggles in complex reasoning tasks. In addition, these automatic methods ignore considering the consistency between the reasoning process and the final answer (Wang et al., 2023b), resulting in low-quality and spurious CoT prompting data. Considering these issues, we aim to improve upon existing CoT prompting data for enhancing the complex reasoning capacities of LLMs. Different from existing efforts that primarily improve instructions with evolutionary algorithms (Xu et al., 2023; Guo et al., 2023), CoT improvement requires considering both the question and reasoning steps, as well as their consistency. However, the underlying mechanics and efficacy of CoT prompts for LLMs remain underexplored and we would like to investigate a more sophisticated and fundamental question: *what kind of CoT chains can more effectively elicit the potential capacities for LLMs?*

To delve into this problem, we conduct a series of empirical analysis experiments on GSM8K (Cobbe et al., 2021) to study the impact of CoT from three main aspects, *i.e.*, completeness, specificity, and logicity. First, we vary the number of reasoning steps to examine the performance of CoT prompts

*Corresponding author

Methods	Generation strategy	Task-agnostic	LLM fine-tuning	Correctness check
Few-shot CoT (Wei et al., 2022)	NA	✗	✗	✗
Zero-shot CoT (Kojima et al., 2022)	NA	✓	✗	✗
Auto-CoT (Zhang et al., 2022)	question clustering	✗	✗	✗
CoT Collection (Kim et al., 2023)	human-crafted & ICL	✗	✓	✗
CoTGenius (ours)	complicate, diversify, specify	✓	✓	✓

Table 1: A comparison between our CoTGenius and other CoT generation methods.

with different levels of reasoning completeness. Second, we iteratively refine the CoT prompts to make the reasoning processes more detailed and specific. Finally, we explore the logic of CoT prompts by interchanging the order of reasoning processes and the final answer. Based on the empirical results, we find that incorporating additional and more specific reasoning steps contributes to a highly complete chain-of-thought, resulting in more accurate answers and better performance in mathematical problems. Finally, the logic of reasoning before answering, used in the majority of prior studies, leads to substantial performance improvements compared to the logic of answering then explaining. Based on these findings, we aim to improve upon existing CoT prompting data and fine-tune LLMs on newly generated high-quality CoT data to enhance complex reasoning abilities of LLMs.

There are three major technical contributions in our work. First, we propose CoTGenius, a Chain-of-Thought prompting improvement framework to synthesize more complicated, diverse, and detailed CoT prompts, as shown in Figure 1. In this framework, we introduce three evolution strategies to improve CoT prompts, *i.e.*, *complicate*, *diversify*, and *specify*. Starting from a simple initial CoT prompt (including the question and reasoning steps), the complication and diversity strategies will first upgrade the initial question to a more complex one and create a new diverse question, and then generate reasoning steps for the evolved question. The specificity strategy will rewrite the initial CoT reasoning steps and incorporate more details. Since the CoT improvement is completely conducted by LLMs, the resulting CoT prompting data might be erroneous. Thus we design two filtering strategies to filter the failed CoT prompts, *i.e.*, *evolutionary success judgement* and *correctness verification*. We compared CoTGenius with previous CoT generation work in Table 1. We repeat the improvement process for four rounds using OpenAI ChatGPT API and finally obtain 44,335 CoT prompts.

Second, to validate the effectiveness of our CoTGenius framework, we fine-tune open-source LLMs (*i.e.*, Llama 2-Chat 7B and 13B) with our evolved CoT data, called *ChainLM*. We compare ChainLM to existing popular LLMs on several complex reasoning tasks and our model surpasses many open-source LLMs (*e.g.*, Alpaca, Vicuna) with significant improvements. Specially, ChainLM is superior to

directly fine-tuning LLMs on existing simple CoT datasets in large scale (Kim et al., 2023), which contain about 1.88 million CoT rationales. These datasets are simply built upon existing NLP tasks using prompts without considering the complication and specificity of CoT rationales.

Finally, to deal with the cumulative error issue in reasoning steps leading to inaccurate answers, we propose a CoT reasoning strategy, *step-level debating*. This method employs multiple agents (*i.e.*, LLMs) to debate about each intermediate reasoning step for a consensus, which can capitalize on the strengths of different models to improve the accuracy of intermediate steps. Compared to previous CoT reasoning strategies such as self-consistency and least-to-most, our step-level debating method based on ChainLM exhibits better performance on several reasoning tasks.

2. Empirical Analysis

As CoT prompting gradually becomes an effective means of solving complex reasoning tasks, we try to investigate what kind of CoT chains can more effectively elicit the potential capacities for LLMs from three aspects: inference completeness, prompt specificity, and reasoning logicity. All experiments are conducted by ChatGPT in GSM8K dataset.

2.1. Inference Completeness of CoT

As discovered in prior work (Fu et al., 2023b; Luo et al., 2023), the current CoT prompts usually consist of simpler and incomplete reasoning steps (*i.e.*, only 2 or 3 steps), which poses challenges for LLMs when tackling complex reasoning tasks. Since the step-by-step reasoning plays a pivotal role in guiding LLMs to arrive at answers, we aim to explore the influence of the completeness of reasoning steps (*i.e.*, the number of reasoning steps). Specially, we instruct ChatGPT with 2, 3 or 5 reasoning steps to solve mathematical problems in GSM8K.

As shown in Table 2, the results provide clear evidence of a positive correlation between the number of reasoning steps and model accuracy. Intuitively, as more intermediate reasoning steps are incorporated, LLMs are capable of inferring the final answer with greater ease and accuracy. This observation aligns well with our expectations and underscores

Steps	2	3	5
Accuracy	45.81	53.55	63.23

Table 2: Inference completeness: Accuracy (%) of different numbers of reasoning steps on GSM8K.

Iterations	0	1	2
Accuracy	76.49	79.15	79.37

Table 3: Prompt specificity: Accuracy (%) of varying refinement iterations on GSM8K.

the significance of completing step-wise reasoning in the CoT framework.

2.2. Prompt Specificity of CoT

Besides the number of reasoning steps, the level of specificity and details for each step also play pivot roles in CoT reasoning. If the CoT reasoning process misses some important details about the question, LLMs may conduct superficial reasoning, leading to inaccurate answers. To explore the efficacy of the specificity of CoT prompts, we employ ChatGPT to iteratively refine existing CoT reasoning steps to be more specific and detailed. We perform two refinement iterations and the results are presented in Table 3.

As we can see from the table, after one iteration of specificity refinement, the model could achieve satisfactory performance and additional iterations bring minimal accuracy benefits. The reasons behind might be that appropriately adding details to the reasoning process is able to assist LLMs in inferring answers more accurately, while more specificity iterations introduces little useful information to the model’s reasoning. These findings underscore the significance of increasing the specificity of CoT prompts, ultimately contributing to the enhancement of CoT reasoning and problem-solving proficiency.

2.3. Reasoning Logicality of CoT

When employing the CoT prompting to guide LLMs, most studies adopt the logic of first conducting step-by-step reasoning and then deriving the final answer based on the previous reasoning steps. However, in our empirical experiments, we explore another logic of first providing the answer at the outset, followed by detailed step-by-step explanations. This way can be formulated as an outcome explanation problem in explainable AI, which is called *rationalization* (Gurrapu et al., 2023).

In our experiments, we aim to investigate two logical patterns of CoT, *i.e.*, inferring the answer at the front or behind of the reasoning steps, and the results are present in Table 4. As can be seen from

Answer Position	front	behind
Accuracy	68.69	76.80

Table 4: Reasoning logicality: Accuracy (%) of different answer positions on GSM8K.

the results, the commonly used CoT logic of predicting the answer after rigorous reasoning, a structure more aligned with human thinking patterns, yields higher accuracy. This observation underscores the significance of mirroring human cognitive processes in the CoT framework to enhance model performance in complex reasoning tasks by using the logic of reasoning-then-answering.

In the above empirical experiments, we have observed that there are three factors that affect CoT reasoning for LLMs. **First**, increasing the number of reasoning steps plays a pivotal role in enhancing the completeness of CoT prompts and leading to accurate solutions. **Second**, improving the level of specificity and details for CoT prompts can assist LLMs in performing rigorous step-by-step reasoning and making accurate predictions. **Third**, the logic of reasoning-then-answering is an intuitive and more effective thinking pattern for CoT reasoning. In the next section, we will improve existing CoT prompts based on our empirical results.

3. CoT Improvement

3.1. The CoTGenius Framework

Based on the above empirical studies, we propose a CoT improvement framework, CoTGenius, which introduces three evolution strategies to retrofit existing CoT prompts. Besides, we propose evolutionary success judgement and correctness verification to filter out the failed evolved and erroneous CoT prompts. The overall illustration of CoTGenius is shown in Figure 1.

3.1.1. CoT Improvement Strategies

Existing work mostly focuses on improving instructions to enhance the instruction following capability of LLMs (Xu et al., 2023; Guo et al., 2023). Compared to them, CoT improvement requires considering both the question and reasoning steps, as well as the consistency between them. According to our empirical studies, we design three CoT improvement strategies using evolutionary methods including complicate, diversify, and specify. Detailed instructions of three strategies can be found in Table 8 in the Appendix A.

Complicate. In general, solving complex problems requires a number of reasoning steps, which can

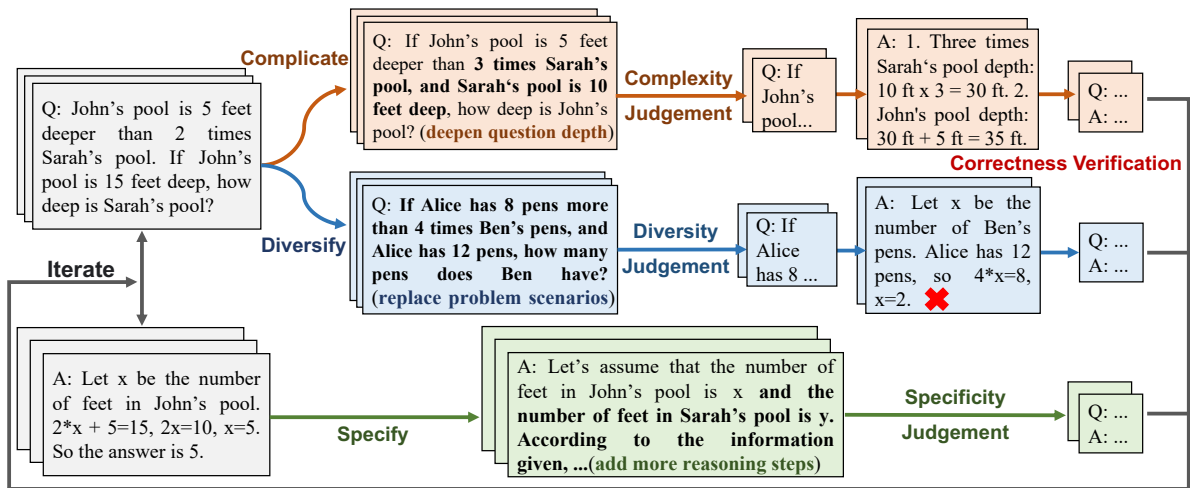


Figure 1: Illustration of our proposed CoTGenius framework for improving CoT prompts.

better demonstrate the step-by-step reasoning process and is beneficial to stimulate the model's reasoning ability during fine-tuning. Thus, we propose a *complication* evolution strategy, aimed at upgrading the question in CoT prompts into a complex and challenging one involving more reasoning steps. Specifically, we use ChatGPT to complicate the initial question based on complication instructions paired with few-shot in-context demonstrations. In our complication instructions, we adopt two methods to increase the difficulty level of the question:

- **Adding Conditions and Constraints:** Simply adding conditions and constraints to the question can directly lead to an increase of reasoning steps. For example, a math question " $x^2 = 16$, what is x ?" can be complicated by adding constraints $x > 0$.

- **Deepening Question Depth:** In addition to adding constraints, the second method focuses on increasing the depth of the given question, expecting to transform superficial questions into profound ones and further fostering the complexity of CoT. For example, we can deepen the math question " $x^2 = 16$, what is x ?" into a challenging one " $x^2 = 16$, $x > 0$, is x a prime number?".

After complicating the initial questions in CoT prompts, we further use ChatGPT to generate their corresponding CoT reasoning process with more steps compared to the initial reasoning process.

Diversify. Many studies have reported that increasing the diversity of training data can enhance the generalization ability of LLMs (YuLan-Chat-Team, 2023). Therefore, we propose a *diversity* evolution strategy to expand the question topics of CoT prompts, making our model applicable to more general scenarios. Similarly, we also leverage ChatGPT to execute the diversity evolution by the following two methods:

- **Replacing Problem Scenarios:** This method can totally change the background of the original

problems with the aim of increasing the topic diversity of the problem formally. For example, a problem scenario "A pipe takes an hour to fill the tank" can be changed to "A car travels from point A to point B at a speed of 60 km/h".

- **Drawing Inspiration from the Given Question:** Besides changing the background, it is more important to diversify the core of a given problem. For this purpose, ChatGPT is employed to leverage the content of given questions as a source of inspiration to craft completely new questions. For example, the math question " $x^2 = 16$, what is x ?" can be reformulated into "The radius of a circle x satisfies $x^2 = 16$, what is the radius of the circle?".

After diversifying the questions in the initial CoT prompts, we also utilize ChatGPT to generate their reasoning processes correspondingly.

Specify. As observed in our empirical study about the prompt specificity, CoT prompts with sufficient details can significantly improve the performance of the model. Therefore, we introduce a *specificity* evolution strategy to insert more details into the CoT reasoning steps while keeping the underlying question unchanged. Specifically, this strategy involves ChatGPT to rewrite the CoT reasoning process corresponding to each question using two methods:

- **Adding More Reasoning Steps:** The first approach incorporates additional reasoning steps into the original CoT prompts, making the reasoning process more detailed and specified. This enhancement aims to provide LLMs with clearer and more comprehensive step-by-step reasoning guidance.

- **Rewriting Existing Reasoning Steps:** The second method will revise existing reasoning steps within CoT prompts to make them more standardized and smooth. In particular, the standardization operation aims to make the CoT prompts consistent with the logic of reasoning-then-answering, while the smoothing operation is to make the reasoning

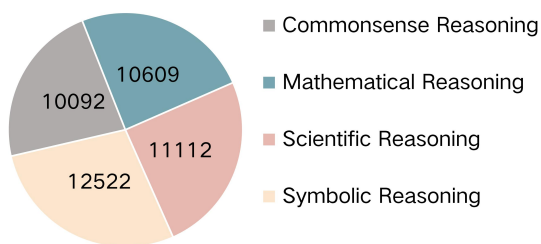


Figure 2: Statistics of synthesized CoT samples.

conducted step-by-step rigorously.

In the specificity process, we remain the input question unchanged but improve CoT reasoning steps. By specifying the CoT rationales, we endeavor to enhance their clarity, details, and effectiveness to guide LLMs in complex reasoning tasks.

3.1.2. CoT Filtering

Although we employ ChatGPT to conduct the CoT improvement automatically, it is imperative to ensure that the synthesized CoT prompts align well with the evolution strategies. Besides, the evolving CoT reasoning steps might be inconsistent with the question and contain errors, so we need to verify the correctness of CoT prompts and filter those failed evolved data. To achieve these goals, we introduce two CoT filtering methods: evolutionary success judgement and correctness verification.

Evolutionary Success Judgement. In this method, we adopt three powerful LLMs, *i.e.*, ChatGPT, Claude, and PaLM, to assess whether the rewritten questions in the complication and diversity evolution, as well as the rewritten CoT reasoning steps in the specificity evolution, successfully achieve the evolution objectives compared to the input data. We only retain those CoT prompts that have undergone successful evolution as our synthesized data via max-voting among the three LLMs.

Correctness Verification. During the process of complication and diversity evolution, we first upgrade the CoT questions and then utilize ChatGPT to generate reasoning steps for these evolved new questions. Thus, we design correctness verification to assess the consistency between the questions and generated reasoning steps. We also use the three aforementioned LLMs to provide judgments regarding the reasoning correctness, and use the max-voting strategy to retain the correct ones.

The instructions of evolutionary success judgement and correctness verification can be found in Table 9 and Table 10 in the Appendix B.

3.2. Improved CoT Dataset

We categorize complex reasoning tasks that benefit from the CoT technique into four distinct types, *i.e.*,

commonsense reasoning, mathematical reasoning, scientific reasoning, and symbolic reasoning. For each task type, we carefully select seed datasets with CoT prompts to serve as the initialization of our improvement process:

- **Commonsense Reasoning:** We adopt StrategyQA (Geva et al., 2021) and Date Understanding (bench authors, 2023) as seed datasets.
- **Mathematical Reasoning:** We consider AQUA-RAT (Ling et al., 2017) and GSM8K (Cobbe et al., 2021) as seed datasets.
- **Scientific Reasoning:** We employ ARC-Challenge (Clark et al., 2018), OpenbookQA (Mihaylov et al., 2018), and WorldTree (Xie et al., 2020) as seed datasets.
- **Symbolic Reasoning:** We include Colored Objects, Tracking Shuffled Objects, and Word Sorting from BIGbench (bench authors, 2023) as seed datasets.

We mix up the training sets of the above seed datasets as the input data of our evolution process and perform CoT improvement following the CoT-Genius framework. These seed datasets are iteratively evolved into new data, with each round of evolution building upon the previous results. This iterative process is repeated for four epochs to obtain sufficient data of varying complexity, diversity, and specificity. Note that we only keep the newly generated CoT prompts without the seed datasets and randomly shuffle the samples to create the final dataset. In the end, we successfully generate a total of 44,335 samples. The statistics of our dataset are presented in Figure 2.

3.3. ChainLM: Fine-tuning LLM on CoT

To validate the effectiveness of our improved CoT data in enhancing the reasoning capabilities of LLMs, we fine-tune Llama 2-Chat 7B and 13B models (Touvron et al., 2023b) with our dataset for three epochs. We call the fine-tuned model as *ChainLM*. Based on our improved CoT data, ChainLM possesses exceptional capabilities in complex reasoning tasks. However, we discover that most inaccurate answers usually stem from the errors in intermediate steps and these errors tend to accumulate over steps. Thus, to increase the accuracy of intermediate steps, we propose a step-level debating method for improving our ChainLM model.

Step-level Debating. Due to the absence of correctness labels for intermediate steps, it is challenging to fine-tune LLMs with process supervision. Inspired by previous work employing multiple LLMs to discuss for output consensus (Chen et al., 2023; Liang et al., 2023; Du et al., 2023), we propose a step-level debating method to improve the CoT reasoning at each step. Different from prior work focused on discussing about the final solution, our

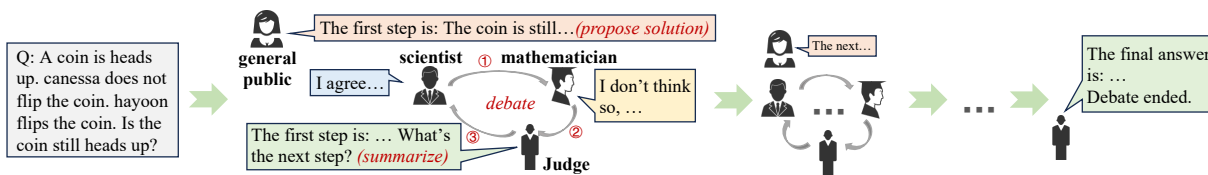


Figure 3: Illustration of our proposed step-level debating strategy.

method employs multiple LLMs to engage in debating about each reasoning step for fine-grained consensus. Specifically, we set one general public and three debaters roles, *i.e.*, scientist, mathematician, and judge, which are all built upon our ChainLM model. The instructions for each role are provided in Table 11 in the Appendix C. For each question, general public is asked to answer it step by step, while scientist and mathematician are required to debate for each reasoning step. The judge is responsible for summarizing the debate between scientist and mathematician and finally giving an outcome of each reasoning step. Once a solution consensus is reached at the current reasoning step, general public proceeds to the next step of reasoning, as depicted in Figure 3. To illustrate the entire process, we present an example of step-level debating in Table 12 in Appendix C.

4. Experiment

In this part, we detail the experimental setup and then highlight the main takeaways of our results.

4.1. Experimental Setup

Datasets and Metrics. As shown in section 3.2, we categorize complex reasoning tasks into four types. To further evaluate the performance of our model on these tasks, we select two widely-used datasets for each task type:

- **Commonsense Reasoning:** We select CommonsenseQA (Talmor et al., 2019), a multiple-choice QA dataset requiring commonsense knowledge to predict the answer, and SocialIQA (Sap et al., 2019), containing multiple-choice questions for probing emotional and social intelligence in a number of daily situations. We conduct 3-shot evaluation on validation sets of CommonsenseQA (1221 samples) and SocialIQA (1954 samples).

- **Mathematical Reasoning:** We choose MATH (Hendrycks et al., 2021), a challenging dataset about competition mathematical problems, and Elementary Mathematics from MMLU (Hendrycks et al., 2020), containing elementary mathematical problems. We perform 3-shot evaluation on test set of MATH (5000 samples) since it is pretty challenging for LLMs and zero-shot on test set of Elementary Mathematics (378 samples).

- **Scientific Reasoning:** We choose a multi-modal multiple-choice science question dataset ScienceQA (Lu et al., 2022) and a science exam questions dataset SciQ (Welbl et al., 2017). These two datasets center around the topics like physics, chemistry, and biology. We evaluate our model in zero-shot manner on test sets of ScienceQA (2224 samples) and SciQ (1000 samples).

- **Symbolic Reasoning:** We select Penguins in a Table and Object Counting from BIGbench (bench authors, 2023). The former is to answer questions about animals in a table, and the latter is to count different types of objects without any choice provided. We also conduct zero-shot evaluation on validation set of Penguins in a Table (29 samples) and 3-shot on Object Counting (200 samples).

Moreover, we also select an out-domain dataset, *i.e.*, **Phrase Relatedness** (bench authors, 2023), which does not fall into the above four categories. This dataset will present models with a phrase (n -gram) and ask them to select the most related phrase (n -gram) among four choices. Therefore, it can be used to evaluate the semantic reasoning performance of our model. We adopt the training set of Phrase Relatedness (80 samples) for zero-shot evaluation. In our experiments, we compute the accuracy (%) for all evaluation datasets.

Baselines. We compare our ChainLM models with a wide range of existing powerful closed-source and open-source models. Specially, we include four representative closed-source LLMs, including ChatGPT (gpt-3.5-turbo)¹, InstructGPT (text-davinci-002/003) (Ouyang et al., 2022), and Anthropic’s Claude model. For open-source LLMs, we choose LLaMA (7B) (Touvron et al., 2023a), Llama 2-Chat (7B) (Touvron et al., 2023b), Falcon (7B) (Almazrouei et al., 2023), ChatGLM (6B) (Du et al., 2022; Zeng et al., 2022), Alpaca (7B) (Taori et al., 2023), Vicuna (7B) (Chiang et al., 2023), WizardLM (7B and 13B) (Xu et al., 2023) and CoT-T5 (11B) (Kim et al., 2023) for comparison. It is worth noting that CoT-T5 is a model obtained by fine-tuning T5 on a large-scale CoT collection, which contains 1.88 million CoT rationales across 1060 tasks generated by LLMs. To validate the effectiveness of our synthesized CoT data compared to the original seed data, we directly fine-tune Llama 2-Chat (7B) on the mixture of seed datasets as a

¹<https://openai.com/blog/chatgpt/>

Models	Commonsense		Math		Science		Symbol		other domain
	CSQA	SIQA	Math	EleMath	ScienceQA	SciQ	Penguins	O_Counting	Phrase
ChatGPT	73.55	66.84	34.10	68.52	61.92	62.80	58.62	82.50	67.50
Claude	-	-	20.18	-	-	-	47.65	-	-
Davinci003	74.61	70.62	17.66	61.11	54.32	62.80	31.03	78.00	87.50
Davinci002	62.98	61.87	19.10	40.48	42.81	55.70	65.52	70.00	96.25
LLaMA	41.28	39.25	2.90	23.81	31.61	32.50	31.03	22.50	16.25
Llama 2-Chat	65.35	52.19	2.50	24.88	28.87	28.70	38.93	33.50	56.25
Falcon	-	-	2.30	-	-	-	24.16	-	-
ChatGLM	-	-	1.10	-	-	-	14.09	-	-
Alpaca	46.85	45.80	3.55	20.90	31.88	34.70	37.93	25.00	45.00
Vicuna	45.95	34.49	2.60	23.54	35.66	35.10	27.59	36.00	36.25
CoT-T5 11B	81.99	69.65	1.50	26.19	<u>43.03</u>	44.00	<u>44.83</u>	32.50	77.50
WizardLM 7B	43.41	38.69	1.42	20.63	27.56	45.30	24.14	15.00	33.75
WizardLM 13B	59.87	48.93	<u>5.30</u>	<u>33.33</u>	40.83	<u>47.00</u>	20.69	38.00	57.50
ChainLM seed	62.33	55.17	3.80	31.22	41.01	33.40	17.24	33.50	32.50
ChainLM 7B	65.75	54.25	3.92	<u>33.33</u>	42.22	42.30	41.38	<u>46.00</u>	<u>82.50</u>
ChainLM 13B	<u>68.22</u>	<u>55.27</u>	5.38	34.13	43.07	49.40	48.28	52.50	83.75

Table 5: Evaluation results at accuracy on nine complex reasoning datasets. **Bold** and underline fonts denote the best and second best methods among open-source models. “-” means the model has not been tested on this dataset on their original papers.

Variants	Commonsense		Math		Science		Symbol	
	CSQA	SIQA	Math	EleMath	ScienceQA	SciQ	Penguins	O_Counting
ChainLM 7B	65.75	54.25	3.92	33.33	42.22	42.30	41.38	46.00
w/o CS	59.21	50.61	3.84	32.54	38.53	41.90	44.83	48.50
w/o Math	52.83	54.91	3.30	30.69	39.34	43.50	27.59	47.00
w/o Sci	56.67	52.87	4.42	28.84	35.75	42.30	44.83	39.00
w/o Sym	59.79	53.68	4.42	30.16	37.77	41.40	31.03	41.50

Table 6: Results of ablation study.

baseline for evaluation, called ChainLM seed.

Implementation Details. We adopt the code of Alpaca (Taori et al., 2023) to fine-tune our models. Specifically, we fine-tune Llama 2-Chat on our improved CoT data for three epochs with a learning rate of $2e-5$. We set the maximum number of tokens to 4096, and batch size to 128, 256 for Llama 2-Chat 7B and 13B, respectively. For evaluation, we set the temperature to 0.1 for all models to reduce output randomness and set the maximum number of tokens for generation to 512.

4.2. Results and Analysis

The performance results of our model and the comparison models are presented in Table 5. First, we can observe that our ChainLM model substantially outperforms other open-source models. For example, compared to Llama 2-Chat, we achieve significant improvement at accuracy from 24.88 to 33.33 in Elementary Mathematics, 28.87 to 42.22 in ScienceQA, and 33.50 to 46.00 in Object Counting, which demonstrates the effectiveness of fine-tuning on our improved CoT data. Since MATH is composed of exceptionally challenging math problems,

open-source models only achieve low performance compared to closed-source models. Second, CoT-T5 is the most relevant model to us and our model achieves much better results than CoT-T5 except CSQA and SIQA. We hypothesize that commonsense reasoning relies more on world knowledge instead of the reasoning ability of models, while CoT-T5 is trained on extensive NLP tasks to gain knowledge. It is particularly noteworthy that ChainLM outperforms the baseline fine-tuning on the original seed datasets, which clearly underscores the effectiveness of our CoT improvement framework. Finally, scaling our fine-tuned model size to 13B obtains a closer or even better results compared to closed-source models on certain datasets. In Phrase Relatedness, our model achieves 83.75 accuracy better than ChatGPT.

4.3. Ablation Study

Our ChainLM model is fine-tuned on four types of complex reasoning tasks with CoT prompts, thus it is non-trivial to deeply analyze the impact of each task type on the final model performance. In this part, we conduct ablation study by removing one task type from synthesized CoT data (Section 3.2)

Variants	Commonsense		Math		Science		Symbol	
	CSQA	SIQA	Math	EleMath	ScienceQA	SciQ	Penguins	O_Counting
ChainLM 7B	65.75	54.25	3.92	33.33	42.22	42.30	41.38	46.00
w/ SC	66.09	54.55	4.59	38.88	51.17	44.30	51.72	47.00
w/ LtM	63.72	48.36	3.54	32.80	46.58	42.50	48.28	40.50
w/ Debating	66.50	54.60	4.76	44.44	52.47	50.50	55.17	47.50

Table 7: A comparison between our step-level debating method and other reasoning methods.

and then fine-tuning Llama 2-Chat 7B.

Table 6 presents the results of four model variants. It can be seen that removing any task type generally leads to varying degrees of performance degradation. In most cases, the performance on target task closely correlates with the task type in our fine-tuning data. For example, the model fine-tuned without scientific reasoning data exhibits the lowest accuracy on the ScienceQA dataset among the four variants. It is interesting that we notice an implicit relation between commonsense reasoning and scientific reasoning—removing commonsense reasoning data results in a dramatic accuracy degradation in scientific reasoning, and vice versa. The reasons might be that some commonsense reasoning problems require scientific knowledge to assist in their solution, and in scientific problems there are numerous commonsense involved in reasoning. Furthermore, we observe that the performance in symbolic reasoning actually improves after removing some data, which could be attributed to synthesized data interfering with the comprehension of symbols and rules.

4.4. CoT Reasoning Strategies

In Section 3.3, we propose a CoT reasoning strategy, step-level debating, employing multiple LLMs to discuss about each reasoning step for consensus. Therefore, to validate its effectiveness, we compare our method to previous CoT reasoning strategies, *i.e.*, self-consistency (Wang et al., 2022) and least-to-most (Zhou et al., 2022). We set 10 paths for self-consistency and three rounds for our step-level debating.

As we can see from Table 7, our step-level debating method outperforms baselines consistently. Although self-consistency also adopts max-voting to achieve consensus among multiple solutions, it only focuses on the final answer without considering intermediate steps. In contrast, our method proceeds to the next reasoning step only when a consensus is reached at the current step. The failure of least-to-most lies in that many problems cannot be accurately decomposed into simpler subproblems. Our step-level debating method performs well especially in tasks where errors are prone to occur in intermediate steps, *i.e.*, Elementary Mathematics

and Penguins in a Table. This significantly illustrates the necessity of considering the consistency during the reasoning process.

5. Related Work

Chain-of-Thought Prompting. Chain-of-Thought (CoT) (Wei et al., 2022; Kojima et al., 2022) prompting is an effective solution for solving complex problems by explicitly generating reasoning steps. As CoT prompting has a critical effect on improving model performance, many studies (Saparov and He, 2022; Wang et al., 2023a) are proposed to further improve CoT prompting technique. For example, self-consistency (Wang et al., 2023b) proposes generating several reasoning paths and then selecting the most consistent answer by voting. Self-verification (Weng et al., 2022) lets LLMs themselves verify their prediction results. Tree-of-Thought (Yao et al., 2023) is a paradigm that allows LLMs to explore multiple reasoning paths over thoughts by framing the problem as a search over a tree. Instead of improving CoT prompting from the perspective of generation methods, our work explores the factors about why CoT prompting works through empirical analysis and designs a data augmentation framework to improve CoT prompting.

Instruction Tuning. Instruction Tuning is an approach to fine-tuning LLMs on a collection of instructions and responses. Early work focuses on fine-tuning LLMs on specific NLP tasks. T5 (Raffel et al., 2020) first introduces text-to-text framework to fine-tune models on multiple tasks. Subsequent work such as FLAN (Wei et al., 2021), FLAN-T5 (Chung et al., 2022), and ZeroPrompt (Xu et al., 2022) improve LLMs by increasing the number of tasks and carefully designing instructions for the tasks. To bridge the gap between human queries and synthesized instructions, many studies propose to fine-tune LLMs with open-domain instructions. Alpaca (Taori et al., 2023) is fine-tuned on 52K instruction data generated by Self-Instruct (Wang et al., 2022). Vicuna (Chiang et al., 2023) collects 70K user-shared ChatGPT conversations from ShareGPT.com for fine-tuning. WizardLM (Xu et al., 2023) proposes Evol-Instruct to

evolve instructions. Our method is most similar to WizardLM, except that we focus on enhancing CoT data and improving the complex reasoning capabilities of LLMs.

6. Conclusion

This paper presented a CoT improvement framework, CoTGenius, which encompasses three evolution strategies, *i.e.*, complicate, diversify, and specify, alongside two filtering mechanisms: evolutionary success judgement and correctness verification. We fine-tuned Llama 2-Chat on superior CoT prompts synthesized by our CoTGenius framework, called ChainLM. Through a series of rigorous experiments, we demonstrate that our model exhibited better performance when confronted with complex reasoning tasks. To further deal with the accumulative error issue in intermediate reasoning steps, we proposed step-level debating, a collaborative approach where multiple agents engage in discussion for each CoT reasoning step to yield correct answers. We believe that our data and model can facilitate future work towards powerful LLMs.

7. Acknowledgements

This work was partially supported by National Natural Science Foundation of China under Grant No.62222215, Beijing Natural Science Foundation under Grant No. 4222027 and L233008. Xin Zhao is the corresponding author.

8. Bibliographical References

- Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Merouane Debbah, Etienne Goffinet, Daniel Heslow, Julien Launay, Quentin Malaric, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. Falcon-40B: an open large language model with state-of-the-art performance.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared 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 M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher 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 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Justin Chih-Yao Chen, Swarnadeep Saha, and Mohit Bansal. 2023. Reconcile: Round-table conference improves reasoning via consensus among diverse llms. *CoRR*, abs/2309.13007.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- Yilun Du, Shuang Li, Antonio Torralba, Joshua B Tenenbaum, and Igor Mordatch. 2023. Improving factuality and reasoning in language models through multiagent debate. *arXiv preprint arXiv:2305.14325*.
- Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. 2022. Glm: General language model pretraining with autoregressive blank infilling. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 320–335.
- Yao Fu, Hao Peng, Litu Ou, Ashish Sabharwal, and Tushar Khot. 2023a. Specializing smaller language models towards multi-step reasoning. In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 10421–10430. PMLR.
- Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. 2023b. Complexity-based prompting for multi-step reasoning. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian, and Yujiu Yang. 2023. Connecting large language models with evolutionary algorithms yields powerful prompt optimizers. *CoRR*, abs/2309.08532.
- Sai Gurrupu, Ajay Kulkarni, Lifu Huang, Ismini Lourentzou, Laura J. Freeman, and Feras A.

- Batarseh. 2023. Rationalization for explainable NLP: A survey. *CoRR*, abs/2301.08912.
- Seungone Kim, Se June Joo, Doyoung Kim, Joel Jang, Seonghyeon Ye, Jamin Shin, and Minjoon Seo. 2023. The cot collection: Improving zero-shot and few-shot learning of language models via chain-of-thought fine-tuning. *arXiv preprint arXiv:2305.14045*.
- 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.
- Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujia Yang, Zhaopeng Tu, and Shuming Shi. 2023. Encouraging divergent thinking in large language models through multi-agent debate. *CoRR*, abs/2305.19118.
- Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng, Qingwei Lin, Shifeng Chen, and Dongmei Zhang. 2023. Wizardmath: Empowering mathematical reasoning for large language models via reinforced evol-instruct. *CoRR*, abs/2308.09583.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, Xu Jiang, Karl Cobbe, Tyna Eloundou, Gretchen Krueger, Kevin Button, Matthew Knight, Benjamin Chess, and John Schulman. 2021. Webgpt: Browser-assisted question-answering with human feedback. *CoRR*, abs/2112.09332.
- OpenAI. 2023. Gpt-4 technical report. *OpenAI*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Liangming Pan, Alon Albalak, Xinyi Wang, and William Yang Wang. 2023. Logic-lm: Empowering large language models with symbolic solvers for faithful logical reasoning. *CoRR*, abs/2305.12295.
- 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. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Abulhair Saparov and He He. 2022. Language models are greedy reasoners: A systematic formal analysis of chain-of-thought. In *The Eleventh International Conference on Learning Representations*.
- Kumar Shridhar, Alessandro Stolfo, and Mrinmaya Sachan. 2022. Distilling multi-step reasoning capabilities of large language models into smaller models via semantic decompositions. *CoRR*, abs/2212.00193.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. Llama: Open and efficient foundation language models. *CoRR*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Boshi Wang, Sewon Min, Xiang Deng, Jiaming Shen, You Wu, Luke Zettlemoyer, and Huan Sun. 2023a. Towards understanding chain-of-thought prompting: An empirical study of what matters. In *ICLR 2023 Workshop on Mathematical and Empirical Understanding of Foundation Models*.
- 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, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khoshabi, and Hannaneh Hajishirzi. 2022. Self-instruct: Aligning language model with self generated instructions. *arXiv preprint arXiv:2212.10560*.
- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.

- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Yixuan Weng, Minjun Zhu, Shizhu He, Kang Liu, and Jun Zhao. 2022. Large language models are reasoners with self-verification. *arXiv preprint arXiv:2212.09561*.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. 2023. Wizardlm: Empowering large language models to follow complex instructions. *arXiv preprint arXiv:2304.12244*.
- Hanwei Xu, Yujun Chen, Yulun Du, Nan Shao, Wang Yanggang, Haiyu Li, and Zhilin Yang. 2022. Zeroprompt: Scaling prompt-based pretraining to 1,000 tasks improves zero-shot generalization. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4235–4252.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. *arXiv preprint arXiv:2305.10601*.
- YuLan-Chat-Team. 2023. Yulan-chat: An open-source bilingual chatbot. <https://github.com/RUC-GSAI/YuLan-Chat>.
- Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, et al. 2022. Glm-130b: An open bilingual pre-trained model. *arXiv preprint arXiv:2210.02414*.
- Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2022. Automatic chain of thought prompting in large language models. In *The Eleventh International Conference on Learning Representations*.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. [A survey of large language models](#). *CoRR*, abs/2303.18223.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V Le, et al. 2022. Least-to-most prompting enables complex reasoning in large language models. In *The Eleventh International Conference on Learning Representations*.
- 9. Language Resource References**
- BIG bench authors. 2023. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *Transactions on Machine Learning Research*.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *CoRR*, abs/2110.14168.
- 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.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. In *International Conference on Learning Representations*.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. 2017. Program induction by rationale generation: Learning to solve and explain algebraic word problems. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 158–167.
- Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35:2507–2521.

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.

Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. Social iqa: Commonsense reasoning about social interactions. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4463–4473.

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.

Johannes Welbl, Nelson F Liu, and Matt Gardner. 2017. Crowdsourcing multiple choice science questions. In *Proceedings of the 3rd Workshop on Noisy User-generated Text*, pages 94–106.

Zhengnan Xie, Sebastian Thiem, Jaycie Martin, Elizabeth Wainwright, Steven Marmorstein, and Peter Jansen. 2020. [WorldTree v2: A corpus of science-domain structured explanations and inference patterns supporting multi-hop inference](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 5456–5473, Marseille, France. European Language Resources Association.

Appendix

We provide the implementation details of the CoT-Genius framework as supplementary materials.

A. CoT Improvement Strategies

We present the instructions for the three CoT improvement strategies including complicate, diversify, and specify in Table 8.

B. CoT Filtering

We present the instructions for evolutionary success judgement including complicate, diversify, and specify in Table 9. The instruction for correctness verification is shown in Table 10.

C. Step-level Debating

We provide descriptions for the roles of general public, scientist, mathematician, and judge for debater initialization in Table 11. To clarify the step-level debating process, we present a step-level debating example of Object Counting in Table 12.

Complicate	<p>I want you to act as a Question Rewriter. Your objective is to rewrite a given question into a more complex version to make it require more reasoning steps. But the rewritten question must be reasonable, understandable, and answerable by humans.</p> <p>You SHOULD complicate the given question using the following methods:</p> <ol style="list-style-type: none"> 1. Add some more constraints/requirements into #Given Question#. 2. Increase the depth of the #Given Question#. <p>#Rewritten Question# must be a solvable problem independent of the #Given Question#.</p> <p>#Given Question#: #Rewritten Question#:</p>
Diversify	<p>I want you to act as a Question Rewriter. Your objective is to rewrite a given question into a more diverse version. But the rewritten question must be reasonable, understandable, and answerable by humans.</p> <p>You SHOULD diversify the given question using the following methods:</p> <ol style="list-style-type: none"> 1. Replace problem scenarios. 2. Draw inspiration from the #Given Question# to create a brand new question. <p>#Rewritten Question# must be a solvable problem independent of the #Given Question#.</p> <p>#Given Question#: #Rewritten Question#:</p>
Specify	<p>I want you to act as a Chain-of-Thought Rewriter. Given a question and its Chain-of-Thought answer, your objective is to rewrite the given Chain-of-Thought answer into a more specific version. But the rewritten CoT must be reasonable and have the same answer as the given CoT.</p> <p>You SHOULD specify the given CoT using the following methods:</p> <ol style="list-style-type: none"> 1. Add more reasoning steps to make the reasoning progress more detailed. 2. Rewrite existing reasoning steps to make them more standardized. <p>#Given Question#: #Given CoT#: #Rewritten CoT#:</p>

Table 8: Instructions of three evolutionary strategies.

Complicate	<p>Given two questions, try your best to judge whether #Question 2# is more difficult than #Question 1#. If #Question 1# is more difficult, write 'No'. If #Question 2# is more difficult, write 'Yes'.</p> <p>#Question 1#: #Question 2#: #Your Judgement#:</p>
Diversify	<p>Given two questions, try your best to judge whether #Question 2# is different from #Question 1#. If the two questions are different, write 'Yes'. Otherwise, write 'No'.</p> <p>#Question 1#: #Question 2#: #Your Judgement#:</p>
Specify	<p>Given a question and two Chain-of-Thought answers to the question, try your best to judge whether #CoT 2# is better than #CoT 1#. If #CoT 2# is better than #CoT 1#, write 'Yes'. If #CoT 1# is better than #CoT 2#, write 'No'.</p> <p>#Question#: #CoT 1#: #CoT 2#: #Your Judgement#:</p>

Table 9: Instructions of evolutionary success judgement.

<p>Given a question and an answer to the question, try your best to judge whether the answer is right or wrong. If it's right, write 'Yes'. If it's wrong, write 'No'.</p> <p>#Question#: #Answer#: #Your Judgement#:</p>

Table 10: Instructions of correctness verification.

general public	<p>You are the general public. Given a problem, you should give one step of your Chain-of-Thought answer. For each step, we will have a debating and the judge will decide the final answer for this step. You need to give the next step based on the previous steps until the judge gives the final answer. It must be noted that you can only give one step at a time.</p>
scientist	<p>You are the scientist. Given a problem and one solving step, you should judge whether the step and discussion are correct. If they are not correct, you should give your reason and your opinion of the correct step.</p>
mathematician	<p>You are the mathematician. Given a problem and one solving step, you should judge whether the step and discussion are correct. If they are not correct, you should give your reason and your opinion of the correct step.</p>
judge	<p>You are the judge. Given a problem and the debating process of one solving step, you should judge which opinion is correct and give the answer of the very step. If you can conclude the final answer directly, repeat the final answer with 'Debate ended.' in the end. If there are choices in the question, give the right choice.</p>

Table 11: Descriptions for the roles of general public, scientist, mathematician, and judge in CoT debating.

user	I have a potato, three heads of broccoli, a cabbage, and two stalks of celery. How many vegetables do I have?
general public	First, potatoes, broccoli, cabbage, and celery are all vegetables.
scientist	I agree.
mathematician	I think the first step of reasoning is correct.
judge	Potatoes, broccoli, cabbage, and celery are all vegetables.
general public	I will count the three heads of broccoli as one vegetable. Then, I will count the cabbage as another vegetable. Finally, I will count the two stalks of celery as another vegetable. So, in total, I have three vegetables.
scientist	The count is incorrect. The three heads of broccoli should be counted as separate vegetables, resulting in a total of 4 vegetables. Additionally, the cabbage and celery should also be counted separately, resulting in a total of 6 vegetables. Therefore, the correct count of vegetables is 4 (broccoli heads) + 1 (cabbage) + 2 (celery stalks) = 7 vegetables.
mathematician	I disagree with the count itself. By considering each broccoli head as a separate vegetable, we arrive at 3 vegetables. Additionally, counting the potato, cabbage, and celery separately, we have a total of 7 vegetables. Therefore, the correct count of vegetables is 7, which includes 1 potato, 3 broccoli heads, 1 cabbage, and 2 celery stalks.
judge	The last reasoning step is correct in recognizing the need to count the number of vegetables. However, the counts of general public and scientist are incorrect. By considering each broccoli head as a separate vegetable, we arrive at 3 vegetables. Additionally, counting the potato, cabbage, and celery separately, we have a total of 7 vegetables. Therefore, the correct count of vegetables is 7, which includes 1 potato, 3 broccoli heads, 1 cabbage, and 2 celery stalks. Debate ended.

Table 12: An example of step-level CoT debating.