

# Improve Student’s Reasoning Generalizability through Cascading Decomposed CoTs Distillation

Anonymous ACL submission


## Abstract

Large language models (LLMs) exhibit enhanced reasoning at larger scales, driving efforts to distill these capabilities into smaller models via teacher-student learning. Previous works simply fine-tune student models on teachers’ generated Chain-of-Thoughts (CoTs) data. Although these methods enhance in-domain (IND) reasoning performance, they struggle to generalize to out-of-domain (OOD) tasks. We believe that the widespread spurious correlations between questions and answers may lead the model to preset a specific answer which restricts the diversity and generalizability of its reasoning process. In this paper, we propose **Cascading Decomposed CoTs Distillation (CasCoD)** to address these issues by decomposing the traditional single-step learning process into two cascaded learning steps. Specifically, by restructuring the training objectives—removing the answer from outputs and concatenating the question with the rationale as input—CasCoD’s two-step learning process ensures that students focus on learning rationales without interference from the preset answers, thus improving reasoning generalizability. Extensive experiments demonstrate the effectiveness of CasCoD on both IND and OOD benchmark reasoning datasets<sup>1</sup>.

## 1 Introduction

Recent developments in LLMs have brought remarkable improvements in reasoning tasks via CoT prompting (Wei et al., 2022b). However, these great reasoning capabilities are often associated with more parameters (Wei et al., 2022a), which is not practical to emergent in smaller language models (SLMs). Existing works (Magister et al., 2023; Ho et al., 2023; Fu et al., 2023; Zhou and Ai, 2024) try to make the reasoning capabilities isolated and distilled to student SLMs by simply fine-tuning

<sup>1</sup>Code can be found at <https://anonymous.4open.science/r/ef334sf-FB92>



(a) Answer SFT consistently outperform Std-CoT on OOD tasks.

Question: Why did someone bring a swimsuit to a ski resort?  
Options:  
(A) To swim in a heated pool.  
(B) To wear as an underlayer for warmth.  
(C) To use as a fashion statement.  
(D) To participate in a polar bear plunge event.  
Answer: (A) To swim in a heated pool.

(b) A case of spurious correlation between questions and answers.

Figure 1: (a) Empirical results of standard CoT distillation (Std-CoT) and directly fine-tuning on answer labels without CoTs (Answer SFT) on one in-domain (BBH-test) and the other four out-of-domain benchmark reasoning datasets. (b) In the given example, the semantic similarity between "swimsuit" in the question and "swim" in the answer demonstrates a high level of match, which could allow the model to predict the answer using simple keyword matching or certain rules.

on teacher LLMs generated CoTs data, known as standard CoTs distillation (Std-CoT). Although the method effectively leverages the LLMs’ CoTs to boost the reasoning performance of student models on seen tasks, it does not ensure effective reasoning in OOD settings, leading to weak generalization on unseen tasks. Our pioneer study demonstrates that, as shown in Figure 1 (a), when using the same IND training dataset, student models developed via the method Std-CoT perform better on IND tasks but significantly worse on OOD tasks compared

to models fine-tuned directly with question-answer pairs. The surprising findings indicate that students' CoTs do not effectively transfer to new domains and these SLMs seem to be more adept at learning to predict answers directly from questions.

We attribute these issues to the spurious correlations between questions and answers that are commonly found in implicit reasoning tasks (Gururangan et al., 2018; Zellers et al., 2019; Blodgett et al., 2020), as illustrated in Figure 1 (b). The Std-CoT approach requires models to learn both the rationale and the answer in a single step, where the learned spurious correlations in training stage can adversely affect the quality of rationale generation during inference. That is to say, upon reading a question, student models may fastly, unconsciously, and automatically formulate a "preset answer" (Hagendorff et al., 2022), which in turn may lead them to implicitly reduce the token generation space when producing CoT. This results in diminished diversity and generalizability of their rationales.

In this paper, we propose **Cascading decomposed CoTs Distillation (CasCoD)**, a straightforward yet effective method to address these issues. Specifically, we decompose the traditional single-step learning process of Std-CoT into two cascaded learning steps: a rationale learning step and an answer learning step. In the rationale learning step, the training objective, with the answer removed, is defined as <sup>2</sup>:  $q \rightarrow r$ . In the answer learning step, we concatenate the question with the target output from the rationale learning step and use this combined input for the answer learning step, setting the training objective as  $q, r \rightarrow a$ . This cascading two-step learning configuration mitigates the capture of spurious correlations between questions and answers during the training phase, ensuring that students focus on learning rationales without interference from the preset answers. Furthermore, the inference phase execution pipeline is aligned with the training phase; the model first generates a rationale when given a question, and then, based on the question-rationale pair, predicts the final answer, further alleviating potential reasoning biases caused by spurious correlations.

Extensive experiments demonstrate that CasCoD outperforms the baselines on both IND and OOD benchmark reasoning datasets (§4.3). Besides, we validate the generalizability of CasCoD across different models, model sizes, and training data sizes

(§4.4). Further analyses validate our hypothesis (§5.1) and confirm the significant impact of the two-step cascading learning process (§5.2) and the robustness of CasCoD (§5.3). The experiments on reasoning faithfulness (§5.4) and case studies (§5.5) indicate that models distilled by CasCoD can reason more consistently and demonstrate better generalization than baselines, effectively addressing interference from question-answer spurious correlations. Our contributions are as follows:

- We find that standard CoT distillation methods exhibit limited generalizability on OOD tasks, almost performing worse than methods fine-tuned directly with question-answer pairs.
- We decompose the traditional single-step learning process into two cascading learning steps to alleviate the impact of spurious correlations between questions and answers.
- Extensive experiments confirm the effectiveness of our method across both IND and OOD datasets, showing that CasCoD can generate more generalizable CoTs.

## 2 Related Works

**CoT Capability of Language Models.** LLMs have demonstrated a wide array of capabilities in numerous natural language processing tasks, underscored by various studies (Chowdhery et al., 2023; Wei et al., 2022a). One notable manifestation of this is the CoT prompting technique (Wei et al., 2022b), which facilitates models in articulating a series of deductive reasoning steps. This method has substantially enhanced LLMs' problem-solving abilities, as evidenced in several works (Kojima et al., 2022a; Wang et al., 2023b; Huang et al., 2023). Despite these advancements, the effectiveness of CoT prompting notably diminishes in smaller models (Wei et al., 2022a). Research by Chung et al. (2022) indicates that with targeted training on CoT data via instruction tuning, SLMs can unlock CoT capabilities. In our study, we show that SLMs' CoT performance can be further enhanced by decomposing the standard CoT distillation process into two cascaded learning steps.

**Distilling Knowledge from LLMs.** Numerous studies (Taori et al., 2023; Chiang et al., 2023; Peng et al., 2023) have explored the knowledge distillation from advanced LLMs like ChatGPT (OpenAI, 2023). These efforts typically focus on

<sup>2</sup> $q$ : the question,  $r$ : the rationale,  $a$ : the answers.

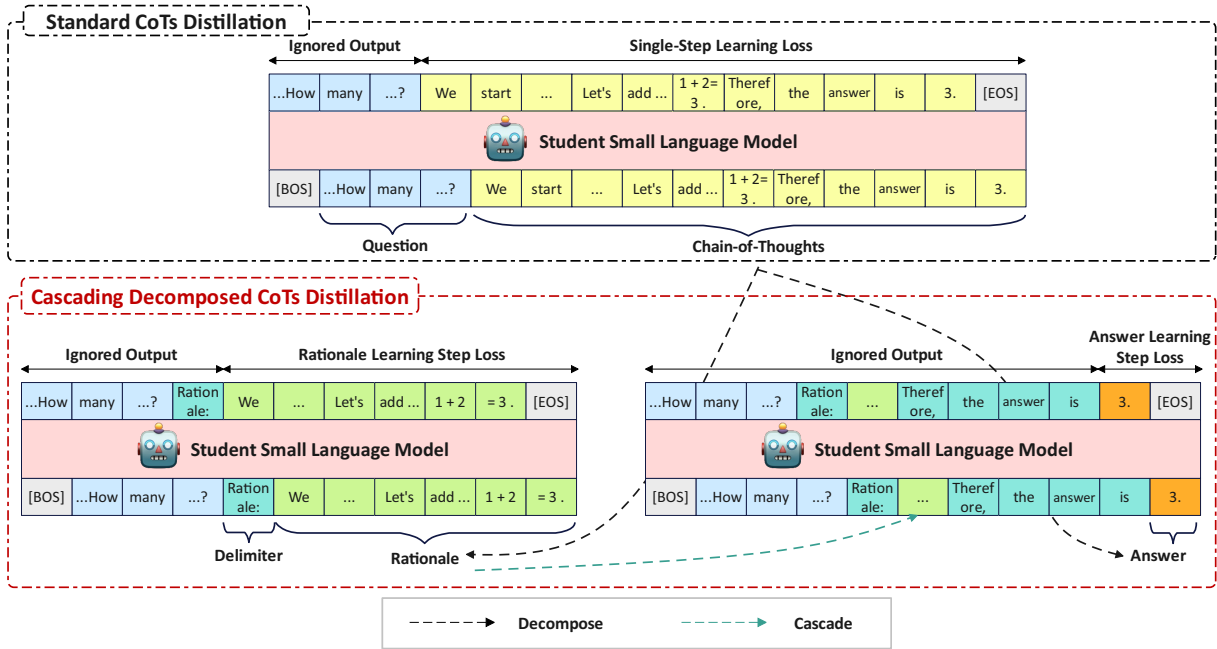


Figure 2: Overview of our proposed method **Cascading Decomposed CoTs Distillation (CasCoD)**. Different from the standard CoTs distillation, we decompose the single CoT learning step into two comprehensive learning steps including the rationale learning step and the answer learning step, and then learn them in a cascaded way.

149 distilling a broad range of abilities via instruction  
 150 tuning on extensive and varied datasets (Xu et al.,  
 151 2023; Wu et al., 2023; Jiang et al., 2023b; Li et al.,  
 152 2024). However, our work is aimed at distilling the  
 153 CoT reasoning capabilities from LLMs, in line  
 154 with Magister et al. (2023); Ho et al. (2023), who  
 155 propose standard CoTs distillation by directly fine-  
 156 tuning SLMs on teacher LLMs’ CoTs. Fu et al.  
 157 (2023) expands on this by using various reason-  
 158 ing data formats for specializing domain-specific  
 159 SLMs. Wang et al. (2023c) distill SLMs via learn-  
 160 ing from self-reflection and feedback in an inter-  
 161 active, multi-round paradigm with teacher LLMs.  
 162 Hsieh et al. (2023) propose to learn the rationale  
 163 and answers as separate goals for optimizing. Li  
 164 et al. (2022) propose learning the entire CoTs and  
 165 the single answers to enhance the reasoning of stu-  
 166 dents. Based on these, Liu et al. (2023) introduce  
 167 an additional distillation objective, self-evaluation,  
 168 aiming for SLMs to assess the accuracy of their  
 169 CoTs akin to LLMs’ evaluative processes. And  
 170 Chen et al. (2024) maximize the mutual informa-  
 171 tion between objectives. Ranaldi and Freitas (2024)  
 172 use in-family and out-family teachers to gener-  
 173 ate more CoTs for fine-tuning students. However,  
 174 these methods are affected by the spurious correla-  
 175 tions due to their isolated optimization objectives.  
 176 In contrast, we reorganize the training objectives  
 177 to effectively mitigate this issue.

### 3 Methodology

178  
 179 In this section, we introduce our new distillation  
 180 method that decomposes the single-step learning  
 181 process of standard CoTs distillation into two cas-  
 182 caded learning steps, as illustrated in Figure 2.  
 183 Formally, the standard CoTs distillation objective  
 184  $q \rightarrow CoT$  is split into two learning processes, ra-  
 185 tionale step learning with the objective  $q \rightarrow r$  and  
 186 answer step learning with the objective  $q, r \rightarrow a$ .  
 187 Below we first describe how to extract CoTs from  
 188 teacher LLMs in §3.1. Then we describe the stan-  
 189 dard CoTs distillation method and discuss its lim-  
 190 itations in §3.2. Finally, we provide a detailed  
 191 presentation of our method in §3.3.

#### 3.1 Extract CoTs From Teacher LLMs

192  
 193 The initial phase of the distillation is to extract  
 194 CoTs from teacher LLMs for each question-answer  
 195 pair  $\{q, a\}$  in a raw dataset. This involves em-  
 196 ploying the CoT prompting technique (Wei et al.,  
 197 2022b), which guides the teacher LLMs to gener-  
 198 ate CoTs that follow a prescribed format with multiple  
 199 reasoning steps. The prompt template is shown in  
 200 Appendix C.1. Note that CoTs produced by LLMs  
 201 may not always be correct. To maintain CoTs qual-  
 202 ity, following the previous work (Magister et al.,  
 203 2023; Hsieh et al., 2023), we retain only those that  
 204 match the ground truth answer in the dataset, build-

ing a CoT dataset  $\mathcal{D} = \{q, CoT\}$  for training the student model. Additionally, to facilitate the introduction of CasCoD, we explicitly split the extracted CoTs into two parts based on predefined rules in CoT prompting, formalizing this as  $CoT = r \oplus a$ . For instance, we use the phrase "Therefore, the answer is" to divide the CoT, categorizing the text before this delimiter as the rationale  $r$  and the text after it as the answer  $a$ .

### 3.2 Preliminaries for CoTs Distillation

Previous standard CoTs distillation (Magister et al., 2023; Ho et al., 2023), referred to as single-step learning, is to teach SLMs to generate the CoT in one time as follows:

$$\mathcal{L}_{\text{Std-CoT}} = \mathbb{E}_{q, CoT \sim \mathcal{D}} [\ell(q, CoT)] \quad (1)$$

where  $\ell$  signifies the negative log-likelihood loss function, expressed as:

$$\ell(x, y) = - \sum_{y_t \in y} \log P(y_t | x, y_{<t}) \quad (2)$$

However, this method requires the model to simultaneously learn both rationales and answers in a single step, readily leading to the capture of question-answer spurious correlations in widespread implicit reasoning datasets (Blodgett et al., 2020). These correlations degrade the quality of CoT generation during inference, resulting in weak reasoning generalization. In other words, this implicit learning of correlations might lead student models to preset answers after reading the questions, potentially causing a state reduction in the token generation space when producing CoTs.

### 3.3 Cascading Decomposed CoTs Distillation

Different from the training strategy in standard CoTs distillation, our method decomposes its single-step learning process into two cascaded learning steps, one for the rationale learning step and the other for the answer learning step.

For the rationale learning step, each question is combined with a rationale learning delimiter "Rationale:" as the input  $q$ , with the rationale  $r$  produced by the teacher serving as the label for distilling the rationale. With the answer objective removed, this training strategy allows models to engage in learning rationales without the interference of spurious correlations. The loss function of rationale step learning is as follows:

$$\mathcal{L}_{\text{rationale}} = \mathbb{E}_{q, r, a \sim \mathcal{D}} [\ell(q, r)] \quad (3)$$

For the answer learning step, we concatenate both the input and output of the rationale learning step with an answer learning delimiter "Therefore, the answer is" as the input, and the answer  $a$  serves as the label for distilling the answer. This strategy helps students learn to reason consistently from the question-rationale pair rather than merely presetting spurious answers based solely on the question. The loss function of answer learning step is thus:

$$\mathcal{L}_{\text{answer}} = \mathbb{E}_{q, r, a \sim \mathcal{D}} [\ell(q \oplus r, a)] \quad (4)$$

Due to the inherent tight connection between rationale learning and answer learning, for each instance in the dataset, we optimize both learning objectives simultaneously for the CoTs distillation:

$$\mathcal{L}_{\text{CasCoD}} = (1 - \alpha)\mathcal{L}_{\text{rationale}} + \alpha\mathcal{L}_{\text{answer}} \quad (5)$$

where  $\alpha$  is a hyperparameter used to weight the loss in the two learning steps.

During inference, student models follow the same pipeline as in training: first, generate a rationale based on the question, and then predict the final answer using the question-rationale pair. The cascading training objectives reduce the probability of student models capturing spurious correlations between questions and answers in the training phase, thereby alleviating potential reasoning biases caused by spurious correlations in the inference stage, thus enhancing CoTs generalizability.

## 4 Experiments

In this section, we conduct extensive experiments and analysis to evaluate the effectiveness of our method across both IND and OOD datasets.

### 4.1 Datasets

**In-domain Dataset: BIG-Bench Hard (BBH)** (Suzgun et al., 2023) comprises 27 challenging tasks covering arithmetic, symbolic reasoning et al. from BIG-Bench (BB) (Guo et al., 2023). The majority of the data involve multiple-choice questions, with a few being open-ended. To underscore the superiority of our approach, we chose to perform distillation on this most challenging dataset. We randomly divide the BBH dataset into a training set (BBH-train) for distillation and a test set (BBH-test) for IND evaluation, in a 4:1 ratio.

**Out-of-domain Datasets: (1) BIG-Bench Sub (BB-sub).** BB is a popular benchmark consisting of 203 tasks covering a wide range of topics, including



mathematics, common-sense reasoning, and various other domains. For ease of evaluation, we filter the subtasks within BB based on subtask keywords, focusing on tasks related to "multiple-choice" and "reasoning", and ensure that tasks from BBH were excluded, resulting in 61 subtasks. Then we randomly sample up to 100 instances for each subtask, resulting in BB-sub. (2) **AGIEval** (Zhong et al., 2023) is a renowned human-centric benchmark used to assess LMs’ reasoning abilities, whose tasks span various domains, including college entrance exams (English / Math / Law), logic tests et al. We evaluate our method on the subtasks that are related to multiple-choice questions in the English language. (3) **AI2 Reasoning Challenge (ARC)** (Clark et al., 2018) consists of ARC-Easy (ARC-E) and ARC-Challenge (ARC-C). The distinction lies in ARC-E consisting of relatively simpler questions from middle and high school science exams, while ARC-C comprises more complex and challenging questions. We utilize the testing set of the ARC dataset for evaluation. The statistics of all above datasets can be found in Appendix B.1.

## 4.2 Models & Baselines & Setup

**Models.** We employ the popular open-source language model LLaMA2-7B (Touvron et al., 2023) as the student SLM in the main experiment and also explore different student models in §4.4. Considering the pricing and capabilities, we utilize OpenAI’s powerful black-box LLM, gpt-3.5-turbo-0613, as teacher LLMs to extract CoTs with the same manual prompt used in the previous work (Suzgun et al., 2023).

**Baselines.** We compare our method with the following baselines: (1) **Teacher & Vanilla Student** under various settings, e.g., Zero-shot (+CoT) or Few-shot (+CoT), for showing the impact of distilling reasoning ability from LLMs. (2) **Std-CoT** (Magister et al., 2023; Ho et al., 2023), which is the standard CoTs distillation method that directly fine-tune student models on the CoTs data. (3) **Step-by-step** (Hsieh et al., 2023) is a multi-task CoTs distillation method that distills rationales and answers separately. (4) **MT-CoT** (Li et al., 2022) is also a multi-task CoTs distillation method, but unlike Step-by-step, it simultaneously optimizes the objectives of answer prediction and entire CoTs learning. (5) **SCOTT** (Wang et al., 2023a) that enhances the reasoning consistency of the student model by introducing additional counterfactual data.

**Setup.** We employ LoRA (Hu et al., 2022) for parameter-efficient fine-tuning of the student SLMs. In §5.3, our empirical results indicate that the optimal weight is set  $\alpha$  at 0.3. However, to mitigate the effects of unbalanced weighting, we include an additional method setup for comparison against the baselines in Table 1, labeled CasCoD ( $\alpha = 0.5$ ). All experiments are conducted using a mixed-precision training strategy on  $4 \times$  A100 GPUs. For the inference stage, vLLM<sup>3</sup>(Kwon et al., 2023) is utilized to accelerate inference, employing a greedy decoding strategy to generate text on one single A100 GPU. More details on training and hyperparameters can be found in Appendix B.2.

## 4.3 Main Results

Table 1 presents the automatic evaluation results of our proposed CasCoD and baselines.

**CoTs distillation enhances the reasoning performance of students.** Comparing with the Zero-shot-CoT and Few-shot-CoT settings of student models, the performance of those with distillation is significantly improved by learning CoTs. Except for BB-sub, the student model has 3-4 times improvement compared to vanilla ones across all datasets.

**CasCoD overcomes limitations of distillation baselines in OOD performance.** From the Table 1, we can find that Answer-SFT on the OOD datasets outperforms all the distillation baselines by an average of 5%, which indicates that it seems student models’ performance decreases when learning the CoTs. This pattern is also noticeable in models without distillation, as evidenced by the comparison between Zero-shot and Zero-shot-CoT (or Few-shot and Few-shot-CoT) settings. We attribute this to spurious correlations between questions and answers as introduced in Figure 1 (b), which students can easily learn. The distillation baselines that require students to consider predicting answers while generating the rationale, inadvertently make the simpler task of answer prediction interfere with the rationale learning, thus reducing the generalization of CoTs. In contrast, CasCoD\* not only surpasses Answer-SFT by 7.9% in IND datasets but also achieves comparable results in OOD scenarios. This underscores the effectiveness of our cascade two step learning strategy, which restructures training objectives to mitigate the impact of spurious correlations, in enhancing reasoning capabilities across diverse datasets.

<sup>3</sup><https://github.com/vllm-project/vllm>

Method	Distill?	Gen CoT?	BBH-test	BB-sub	AGIEval	ARC-E	ARC-C	AVG
In-domain?			✓	×	×	×	×	
<b>Teacher: ChatGPT (gpt-3.5-turbo)</b>								
Zero-shot-CoT	×	✓	42.7	44.1	49.5	91.9	81.1	61.9
Few-shot-CoT	×	✓	73.1	-	-	-	-	-
<b>Student: LLaMA2-7B</b>								
Zero-shot	×	×	14.8	15.5	6.9	18.2	13.9	13.9
Zero-shot-CoT	×	✓	10.6	7.7	7.1	18.4	14.8	11.7
Few-shot	×	×	15.1	28.5	25.5	25.5	25.4	24.0
Few-shot-CoT	×	✓	16.3	25.3	9.9	17.2	17.2	17.2
Answer-SFT	×	×	51.5	33.2	31.2	71.6	53.7	48.2
Std-CoT (Magister et al., 2023)	✓	✓	54.2	28.7	21.6	59.6	45.1	41.8
SCOTT (Wang et al., 2023a)	✓	✓	42.4	18.8	13.0	45.7	34.1	30.8
MT-CoT (Li et al., 2022)	✓	✓	<u>56.8</u>	30.3	22.0	49.4	38.2	39.3
Step-by-step (Hsieh et al., 2023)	✓	✓	42.4	27.7	<b>28.8</b>	68.5	48.6	43.2
CasCoD (ours, $\alpha = 0.5$ )	✓	✓	52.5	<u>36.4</u>	28.1	<b>71.8</b>	<b>54.7</b>	<u>48.7</u>
CasCoD* (ours, $\alpha = 0.3$ )	✓	✓	<b>59.4</b>	<b>37.0</b>	<u>28.3</u>	<u>70.6</u>	<u>52.7</u>	<b>49.6</b>

Table 1: Accuracy (%) on in-domain and out-of-domain datasets with different methods. We employ "Let’s think step by step" (Kojima et al., 2022b) for Zero-shot-CoT settings and the manually curated prompt (Suzgun et al., 2023) for Few-shot-CoT settings. The best performance among distilled student models is marked in **bold**, and the second-best performance is indicated by an underline.

#### CasCoD significantly outperforms the distillation baselines across IND and OOD datasets.

From Table 1, it can be observed that CasCoD significantly outperforms baselines on both IND and OOD datasets in most cases, regardless of whether the loss is weighted. Specifically, CasCoD\* secures an average in-domain improvement of 5.2% and an out-of-domain enhancement of 8.4% over the Std-CoT, along with an overall 6.4% improvement compared to the multi-task learning (Step-by-step) approach. Impressively, CasCoD\* achieves 80.1% of the teacher LLM’s performance in Zero-shot-CoT settings. These results underscore the efficacy of CasCoD, significantly boosting the generative capabilities of CoTs on unseen tasks.

#### 4.4 Ablation Study on Model & Data Sizes

**CasCoD is universally applicable to models of varying sizes.** We perform model distillation on TinyLLaMA-1.1B<sup>4</sup> (Zhang et al., 2024), LLaMA2-7B, and LLaMA2-13B, respectively and compare with standard CoTs distillation (Std-CoT) and multi-task distillation (MT-CoT & Step-by-step). In Figure 3 and 7, we can find that CasCoD consistently outperforms the baselines on both IND and OOD datasets across various sizes of student models. Notably, the performance improvement of our method is the most obvious in the BB-sub, where

the performance of the 13B student model reaches 92.7% of the teacher LLM’s performance. Furthermore, as model sizes increase, the performance gap between CasCoD and the baselines widens on OOD datasets, highlighting CasCoD’s superior efficiency in distilling CoTs for larger models.

**CasCoD is universally applicable to models of different architectures.** We perform model distillation on CodeLLaMA-7B (Touvron et al., 2023), LLaMA3-8B (AI@Meta, 2024) and Mistral-7B-v0.2 (Jiang et al., 2023a), respectively, and compare with Std-CoT and Step-by-step. From the Table 3, we can see that regardless of whether it’s CodeLLaMA, LLaMA3, or Mistral, CasCoD significantly outperforms the baselines on OOD tasks, demonstrating its high effectiveness and scalability. Particularly, on the powerful base model Mistral, the superiority of our method is further amplified.

**CasCoD significantly outperforms standard CoTs distillation on OOD with much less training data.** In Figure 6, CasCoD achieves a 6.3% improvement over Std-CoT on the BB-sub dataset, using only 25% of the full BBH-train data. In the case of other OOD datasets, CasCoD requires merely 12.5% of the full training data to surpass the Std-CoT trained with the full dataset by 5% to 7% in performance. These results demonstrate the efficiency of CasCoD, capable of enhancing CoTs generalization with a smaller amount of CoTs data.

<sup>4</sup><https://huggingface.co/TinyLlama/TinyLlama-1.1B-intermediate-step-1431k-3T>




Figure 3: Ablation study on model size for four OOD datasets. The dotted line indicates the performance of the teacher LLM under the Zero-shot-CoT setting. The results in IND dataset can be found in Appendix A.2.1.

## 5 Analysis

### 5.1 Hypothesis Validation

In this subsection, we aim to validate our hypothesis that student learning of spurious correlations affects the quality of their generated rationales. Due to the nature of causal language models—autoregression, we incorporate "The answer is" into input prompts to directly prompt models to provide the answer immediately (which we treat as the "preset answer" referred to §1) after reading the question, rather than generating rationales before providing an answer. We define the following metrics to assist the validation: (1) **A**: The proportion of incorrect CoT reasoning when the preset answer is wrong. A higher ratio indicates a greater negative impact of the preset answer on CoT reasoning; (2) **B**: The proportion of correct CoT reasoning when the preset answer is correct. A higher ratio indicates a greater positive impact of the preset answer; (3) **C** = **A** − **B** (combine **A** and **B**). We compare CasCoD with Std-CoT using these metrics on four OOD tasks. The results are shown in Table 2. We observe that Std-CoT significantly outperforms CasCoD in metrics 'A' and 'C', while significantly underperforms CasCoD in metric 'B' on all four OOD tasks. This indicates that our method can selectively utilize spurious correlations to some extent, suppressing the negative effects of incorrect preset answers on reasoning and reinforcing the positive effects of correct preset answers on reasoning, thereby enhancing performance on OOD tasks, which experimentally validate our hypothesis.

### 5.2 Two-Step vs. Single-Step Implementation

In this subsection, we explore whether CasCoD's two-step training objectives can be achieved in a single-step computation. Upon analysis of the two

Method & Metric	BB-sub	AGIEval	ARC-E	ARC-C
Std-CoT (A) ↓	79.7	82.2	63.7	50.8
Std-CoT (B) ↑	47.5	30.6	58.1	69.7
Std-CoT (C) ↓	32.2	51.6	5.6	-18.9
CasCoD (A) ↓	73.3	77.9	61.8	50.9
CasCoD (B) ↑	57.4	36.7	71.4	83.9
CasCoD (C) ↓	<b>15.9</b>	<b>41.2</b>	<b>-9.6</b>	<b>-33.0</b>

Table 2: Hypothesis validation results (%) on four OOD tasks. The hypothesis is better supported when 'A' is lower, 'B' is higher, and 'C' is lower.

cascaded learning steps, we find that under teacher-forcing (Goodfellow et al., 2016), CasCoD closely mirrors Std-CoT, with key distinctions including adjustable token-level weights and the omission of delimiters in loss calculations. Each sample in CasCoD's original framework undergoes two forward calculation, raising the question of whether a similar outcome is possible with only one. To investigate this, we introduce a variant, CasCoD-single, which is designed to fulfill the two-step training objectives through a single forward computation. Figure 4 indicates that the two-step CasCoD consistently surpasses the single-step variant across all




Figure 4: Comparison between two-step and single-step training implementations of CasCoD.




Figure 5: Ablation study on task weights  $\alpha$ . The results are reported by **IND-AVG** and **OOD-AVG** that respectively denote average accuracy on IND and OOD datasets. The best performance among weights are marked with "☆".

501 datasets. This underscores that a single forward calculation does not suffice to meet CasCoD’s training objectives, emphasizing the critical importance of the cascading two-step learning process.

### 505 5.3 Impact of Weights

506 In this subsection, we explore how variations in weights affect the performance of models with different parameter sizes on both IND and OOD datasets, as shown in Figure 5.

510 **Students’ performance is not sensitive to weights on OOD datasets.** From the figure, we observe that regardless of weight changes, CasCoD consistently outperforms Std-CoT in OOD by average, even at  $\alpha = 0.9$  (meaning the model allocated only 10% of its attention to rationales generation). This demonstrates that CasCoD exhibits robust generalization in OOD and also underscores the effectiveness of decomposing CoTs for distillation.

519 **CasCoD is more robust for smaller student models.** We observe that the 1.1B model shows less variation in performance compared to the 7B and 13B models in IND. Notably, the performance of the 13B model drops sharply as  $\alpha$  changes from 0.5 to 0.9, indicating that larger models are more susceptible to weight adjustments in the IND dataset.

526 **Prioritizing the rationale over the answer yields better results.** It is evident that across different model sizes, the optimal weights on both IND and OOD datasets range approximately from 0.01 to 0.3, indicating that focusing on the rationale help improve the generalizability of CoTs.

### 532 5.4 Faithfulness of Students

533 To ensure that the rationale provided by students supports their predicted answers, another metric

535 for evaluating CoTs distillation is the faithfulness of students. Following the previous work (Wang et al., 2023a), we use the LAS metric (Hase et al., 2020), whose core idea is to measure the extent that the rationales  $r'$  aid a simulator in predicting the answers  $a'$ , defined as:

$$541 \quad LAS = Acc(q, r' \rightarrow a') - Acc(q \rightarrow a') \quad (6)$$

542 where we employ ChatGPT and GPT4 as the simulator, respectively. The results are shown in Table 4. CasCoD is observed to generate rationales that are more consistent with answers than baselines. This suggests that despite CasCoD’s multi-step learning process, the introduction of cascading learning ensures that students can faithfully reason.

### 549 5.5 Case Study

550 Due to page limitations, we provide a systematic case study in Appendix D to illustrate our hypothesis and the improvement in CoT generalizability.

## 553 6 Conclusion

554 We propose a simple yet effective CoTs distillation method CasCoD to address the issue of question-answer spurious correlations that previous CoTs distillation methods suffer from. Specifically, we decompose the traditional single-step learning process into two cascaded learning steps and restructure their training objectives. Extensive experiments show that CasCoD significantly outperforms the baselines across both IND and OOD datasets. Further analysis reveals that CasCoD is robust to model size, training data size, different models, and weights and can lead to faithful student models.



## 566 Limitations

567 In our study, we explore distilling CoTs into two  
568 cascading steps, which is an initial step toward un-  
569 derstanding finer decompositions. Research (Scha-  
570 effer et al., 2023a) suggests that the emergent abili-  
571 ties of LLMs result from managing multiple sub-  
572 tasks simultaneously, hinting at the potential for  
573 more intricate cascading steps in CoTs. Our current  
574 work does not yet define the precise rules for such  
575 more steps decomposition, nor the optimal timing  
576 and methods for focusing learning on specific steps.  
577 We hope our work can inspire the community and  
578 leave these aspects for future exploration, intended  
579 to refine and extend the CoT reasoning capabilities  
580 of SLMs as suggested by our findings.

## 581 Ethics Statement

582 Our work utilizes CoT data extracted from Chat-  
583 GPT for distillation, which may result in inheriting  
584 the social biases (Schaeffer et al., 2023b) and hal-  
585 lucination (Zhang et al., 2023) present in LLMs.  
586 However, we are optimistic that future advance-  
587 ments in resolving these issues in LLMs will nat-  
588 urally lead to the development of student models  
589 with reduced toxicity.

## 590 References

591 AI@Meta. 2024. [Llama 3 model card](#).

592 Su Lin Blodgett, Solon Barocas, Hal Daumé III, and  
593 Hanna M. Wallach. 2020. Language (technology) is  
594 power: A critical survey of "bias" in NLP. In *ACL*,  
595 pages 5454–5476. Association for Computational  
596 Linguistics.

597 Xin Chen, Hanxian Huang, Yanjun Gao, Yi Wang,  
598 Jishen Zhao, and Ke Ding. 2024. Learning to maxi-  
599 mize mutual information for chain-of-thought distil-  
600 lation. In *Findings of the Association for Computa-  
601 tional Linguistics: ACL 2024*.

602 Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng,  
603 Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan  
604 Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion  
605 Stoica, and Eric P. Xing. 2023. [Vicuna: An open-  
606 source chatbot impressing gpt-4 with 90%\\* chatgpt  
607 quality](#).

608 Aakanksha Chowdhery, Sharan Narang, Jacob Devlin,  
609 Maarten Bosma, Gaurav Mishra, Adam Roberts,  
610 Paul Barham, Hyung Won Chung, Charles Sutton,  
611 Sebastian Gehrmann, Parker Schuh, Kensen Shi,  
612 Sasha Tsvyashchenko, Joshua Maynez, Abhishek  
613 Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vin-  
614 odkumar Prabhakaran, Emily Reif, Nan Du, Ben  
615 Hutchinson, Reiner Pope, James Bradbury, Jacob

Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, 616  
Toju Duke, Anselm Levskaya, Sanjay Ghemawat, 617  
Sunipa Dev, Henryk Michalewski, Xavier Garcia, 618  
Vedant Misra, Kevin Robinson, Liam Fedus, Denny 619  
Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, 620  
Barret Zoph, Alexander Spiridonov, Ryan Sepassi, 621  
David Dohan, Shivani Agrawal, Mark Omernick, An- 622  
drew M. Dai, Thanumalayan Sankaranarayanan Pilla- 623  
i, Marie Pellat, Aitor Lewkowycz, Erica Moreira, 624  
Rewon Child, Oleksandr Polozov, Katherine Lee, 625  
Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark 626  
Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy 627  
Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, 628  
and Noah Fiedel. 2023. Palm: Scaling language mod- 629  
eling with pathways. *J. Mach. Learn. Res.*, 24:240:1– 630  
240:113. 631

Hyung Won Chung, Le Hou, Shayne Longpre, Barret 632  
Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, 633  
Mostafa Dehghani, Siddhartha Brahma, Albert Web- 634  
son, Shixiang Shane Gu, Zhuyun Dai, Mirac Suz- 635  
gun, Xinyun Chen, Aakanksha Chowdhery, Sharan 636  
Narang, Gaurav Mishra, Adams Yu, Vincent Y. Zhao, 637  
Yanping Huang, Andrew M. Dai, Hongkun Yu, Slav 638  
Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam 639  
Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 640  
2022. Scaling instruction-finetuned language models. 641  
*CoRR*, abs/2210.11416. 642

Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, 643  
Ashish Sabharwal, Carissa Schoenick, and Oyvind 644  
Tafjord. 2018. Think you have solved question an- 645  
swering? try arc, the AI2 reasoning challenge. *CoRR*, 646  
abs/1803.05457. 647

Yao Fu, Hao Peng, Litu Ou, Ashish Sabharwal, and 648  
Tushar Khot. 2023. Specializing smaller language 649  
models towards multi-step reasoning. In *ICML*, vol- 650  
ume 202 of *Proceedings of Machine Learning Re- 651  
search*, pages 10421–10430. PMLR. 652

Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 653  
2016. *Deep Learning*. MIT Press. [http://www.  
654 deeplearningbook.org](http://www.deeplearningbook.org). 655

Geyang Guo, Ranchi Zhao, Tianyi Tang, Wayne Xin 656  
Zhao, and Ji-Rong Wen. 2023. Beyond imitation: 657  
Leveraging fine-grained quality signals for alignment. 658  
*CoRR*, abs/2311.04072. 659

Suchin Gururangan, Swabha Swayamdipta, Omer Levy, 660  
Roy Schwartz, Samuel R. Bowman, and Noah A. 661  
Smith. 2018. Annotation artifacts in natural language 662  
inference data. In *NAACL-HLT (2)*, pages 107–112. 663  
Association for Computational Linguistics. 664

Thilo Hagendorff, Sarah Fabi, and Michal Kosinski. 665  
2022. Thinking fast and slow in large language mod- 666  
els. *arXiv preprint arXiv:2212.05206*. 667

Peter Hase, Shiyue Zhang, Harry Xie, and Mohit Bansal. 668  
2020. Leakage-adjusted simulatability: Can models 669  
generate non-trivial explanations of their behavior in 670  
natural language? In *EMNLP (Findings)*, volume 671  
EMNLP 2020 of *Findings of ACL*, pages 4351–4367. 672  
Association for Computational Linguistics. 673

674	Namgyu Ho, Laura Schmid, and Se-Young Yun. 2023.	Weize Liu, Guocong Li, Kai Zhang, Bang Du, Qiyuan	730
675	Large language models are reasoning teachers. In	Chen, Xuming Hu, Hongxia Xu, Jintai Chen, and Jian	731
676	<i>ACL (I)</i> , pages 14852–14882. Association for Com-	Wu. 2023. Mind’s mirror: Distilling self-evaluation	732
677	putational Linguistics.	capability and comprehensive thinking from large	733
		language models. <i>CoRR</i> , abs/2311.09214.	734
678	Cheng-Yu Hsieh, Chun-Liang Li, Chih-Kuan Yeh,	Lucie Charlotte Magister, Jonathan Mallinson, Jakub	735
679	Hootan Nakhost, Yasuhisa Fujii, Alex Ratner, Ranjay	Adámek, Eric Malmi, and Aliaksei Severyn. 2023.	736
680	Krishna, Chen-Yu Lee, and Tomas Pfister. 2023. Dis-	<b>Teaching small language models to reason</b> . In <i>Pro-</i>	737
681	tilling step-by-step! outperforming larger language	<i>ceedings of the 61st Annual Meeting of the Asso-</i>	738
682	models with less training data and smaller model	<i>ciation for Computational Linguistics (Volume 2:</i>	739
683	sizes. In <i>ACL (Findings)</i> , pages 8003–8017. Associa-	<i>Short Papers)</i> , <i>ACL 2023, Toronto, Canada, July</i>	740
684	tion for Computational Linguistics.	<i>9-14, 2023</i> , pages 1773–1781. Association for Com-	741
		putational Linguistics.	742
685	Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan	OpenAI. 2023. Chatgpt (June 13 version). <a href="https://chat.openai.com">https://chat.openai.com</a> .	743
686	Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and		744
687	Weizhu Chen. 2022. Lora: Low-rank adaptation of	Baolin Peng, Chunyuan Li, Pengcheng He, Michel Gal-	745
688	large language models. In <i>ICLR</i> . OpenReview.net.	ley, and Jianfeng Gao. 2023. <b>Instruction tuning with</b>	746
		<b>GPT-4</b> . <i>CoRR</i> , abs/2304.03277.	747
689	Jiaxin Huang, Shixiang Gu, Le Hou, Yuxin Wu, Xuezhi	Leonardo Ranaldi and André Freitas. 2024. Aligning	748
690	Wang, Hongkun Yu, and Jiawei Han. 2023. Large	large and small language models via chain-of-thought	749
691	language models can self-improve. In <i>EMNLP</i> , pages	reasoning. In <i>EACL (I)</i> , pages 1812–1827. Associa-	750
692	1051–1068. Association for Computational Linguistics.	tion for Computational Linguistics.	751
693			
694	Albert Q. Jiang, Alexandre Sablayrolles, Arthur Men-	Rylan Schaeffer, Brando Miranda, and Sanmi Koyejo.	752
695	sch, Chris Bamford, Devendra Singh Chaplot, Diego	2023a. Are emergent abilities of large language mod-	753
696	de Las Casas, Florian Bressand, Gianna Lengyel,	els a mirage? <i>CoRR</i> , abs/2304.15004.	754
697	Guillaume Lample, Lucile Saulnier, L�el�io Ren-		
698	ard Lavaud, Marie-Anne Lachaux, Pierre Stock,	Rylan Schaeffer, Brando Miranda, and Sanmi Koyejo.	755
699	Teven Le Scao, Thibaut Lavril, Thomas Wang, Timo-	2023b. Are emergent abilities of large language mod-	756
700	th�ee Lacroix, and William El Sayed. 2023a. Mistral	els a mirage? <i>arXiv preprint arXiv:2304.15004</i> .	757
701	7b. <i>CoRR</i> , abs/2310.06825.		
702	Yuxin Jiang, Chunkit Chan, Mingyang Chen, and Wei	Mirac Suzgun, Nathan Scales, Nathanael Sch�arli, Se-	758
703	Wang. 2023b. <b>Lion: Adversarial distillation of pro-</b>	bastian Gehrmann, Yi Tay, Hyung Won Chung,	759
704	<b>proprietary large language models</b> . In <i>Proceedings of</i>	Aakanksha Chowdhery, Quoc V. Le, Ed Chi, Denny	760
705	<i>the 2023 Conference on Empirical Methods in Natu-</i>	Zhou, and Jason Wei. 2023. Challenging big-bench	761
706	<i>ral Language Processing, EMNLP 2023, Singapore,</i>	tasks and whether chain-of-thought can solve them.	762
707	<i>December 6-10, 2023</i> , pages 3134–3154. Association	In <i>ACL (Findings)</i> , pages 13003–13051. Association	763
708	for Computational Linguistics.	for Computational Linguistics.	764
709	Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yu-	Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann	765
710	taka Matsuo, and Yusuke Iwasawa. 2022a. Large lan-	Dubois, Xuechen Li, Carlos Guestrin, Percy Liang,	766
711	guage models are zero-shot reasoners. In <i>NeurIPS</i> .	and Tatsunori B. Hashimoto. 2023. Stanford alpaca:	767
		An instruction-following llama model. <a href="https://github.com/tatsu-lab/stanford_alpaca">https://</a>	768
712	Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yu-	<a href="https://github.com/tatsu-lab/stanford_alpaca">github.com/tatsu-lab/stanford_alpaca</a> .	769
713	taka Matsuo, and Yusuke Iwasawa. 2022b. Large lan-		
714	guage models are zero-shot reasoners. In <i>NeurIPS</i> .	Hugo Touvron, Louis Martin, Kevin Stone, Peter Al-	770
		bert, Amjad Almahairi, Yasmine Babaei, Nikolay	771
715	Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying	Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti	772
716	Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gon-	Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-	773
717	zalez, Hao Zhang, and Ion Stoica. 2023. Efficient	Ferrer, Moya Chen, Guillem Cucurull, David Esiobu,	774
718	memory management for large language model serv-	Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller,	775
719	ing with pagedattention. In <i>SOSP</i> , pages 611–626.	Cynthia Gao, Vedanuj Goswami, Naman Goyal, Antho-	776
720	ACM.	ny Hartshorn, Saghar Hosseini, Rui Hou, Hakan	777
		Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa,	778
721	Ming Li, Lichang Chen, Jiuhai Chen, Shwai He, Jiuxi-	Isabel Kloumann, Artem Korenev, Punit Singh Koura,	779
722	ang Gu, and Tianyi Zhou. 2024. Selective reflection-	Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Di-	780
723	tuning: Student-selected data recycling for LLM	ana Liskovich, Yinghai Lu, Yuning Mao, Xavier Mar-	781
724	instruction-tuning. <i>CoRR</i> , abs/2402.10110.	tinnet, Todor Mihaylov, Pushkar Mishra, Igor Moly-	782
		bog, Yixin Nie, Andrew Poulton, Jeremy Reizen-	783
725	Shiyang Li, Jianshu Chen, Yelong Shen, Zhiyu Chen,	stein, Rashi Rungta, Kalyan Saladi, Alan Schelten,	784
726	Xinlu Zhang, Zekun Li, Hong Wang, Jing Qian,		
727	Baolin Peng, Yi Mao, Wenhui Chen, and Xifeng		
728	Yan. 2022. Explanations from large language models		
729	make small reasoners better. <i>CoRR</i> , abs/2210.06726.		

785	Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open foundation and fine-tuned chat models. <i>CoRR</i> , abs/2307.09288.	
793	Peifeng Wang, Zhengyang Wang, Zheng Li, Yifan Gao, Bing Yin, and Xiang Ren. 2023a. SCOTT: self-consistent chain-of-thought distillation. In <i>ACL (1)</i> , pages 5546–5558. Association for Computational Linguistics.	
798	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 <i>ICLR</i> . OpenReview.net.	
803	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 <i>EMNLP</i> , pages 1948–1966. Association for Computational Linguistics.	
810	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. <i>Trans. Mach. Learn. Res.</i> , 2022.	
817	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022b. Chain-of-thought prompting elicits reasoning in large language models. In <i>NeurIPS</i> .	
822	Minghao Wu, Abdul Waheed, Chiyu Zhang, Muhammad Abdul-Mageed, and Alham Fikri Aji. 2023. Lamini-lm: A diverse herd of distilled models from large-scale instructions. <i>CoRR</i> , abs/2304.14402.	
826	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. <i>CoRR</i> , abs/2304.12244.	
831	Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. Hellaswag: Can a machine really finish your sentence? In <i>ACL (1)</i> , pages 4791–4800. Association for Computational Linguistics.	
836	Muru Zhang, Ofir Press, William Merrill, Alisa Liu, and Noah A. Smith. 2023. How language model hallucinations can snowball. <i>CoRR</i> , abs/2305.13534.	
839	Peiyuan Zhang, Guangtao Zeng, Tianduo Wang, and Wei Lu. 2024. Tinyllama: An open-source small language model. <i>CoRR</i> , abs/2401.02385.	
	Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. 2023. Agieval: A human-centric benchmark for evaluating foundation models. <i>CoRR</i> , abs/2304.06364.	842 843 844 845 846
	Yuhang Zhou and Wei Ai. 2024. Teaching-assistant-in-the-loop: Improving knowledge distillation from imperfect teacher models in low-budget scenarios.	847 848 849
	<b>A Additional Experimental Results</b>	850
	<b>A.1 Detailed Results on Mathematical Reasoning Tasks</b>	851 852
	Given the community’s focus on mathematical reasoning, we present detailed experimental results for math-related subtasks from our evaluation datasets in Table 5. It should be noted that the subtasks in the datasets we utilized involve mathematical reasoning tasks. We list as follows:	853 854 855 856 857 858
	1. BB-sub: includes subtasks like ‘elementary_math_qa’ and ‘identify_math_theorems’, as detailed in Table 9.	859 860 861
	2. BBH: includes subtasks such as ‘DateUnderstanding’ and ‘Multi-StepArithmetic’, as detailed in Table 8.	862 863 864
	3. AGIEval: includes subtasks like ‘AQuA-RAT’ and ‘SAT-Math’, as detailed in Table 6.	865 866
	While ARC does not specifically feature tasks for mathematical computation, we identify a number of scientific questions within this dataset that involve mathematical calculations. From the Table 5, We can observe that CasCoD has achieved performance improvements on mathematical reasoning tasks as well.	867 868 869 870 871 872 873
	<b>A.2 Ablation Study on In-domain Dataset</b>	874
	<b>A.2.1 W.R.T. Model Size</b>	875
	The results of the model size ablation study on IND datasets are presented in Figure 7. We observe that CasCoD outperforms the baselines on both the 7B and 13B model and significantly surpasses the teacher LLMs in the Zero-shot CoT setting.	876 877 878 879 880
	<b>A.2.2 W.R.T. Training Data Size</b>	881
	The results of the training data ablation study on IND datasets, as shown in Figure 8, indicate that CasCoD outperforms standard CoTs distillation across various sizes of training data. This demonstrates the efficiency of our proposed method.	882 883 884 885 886

Models & Methods	BBH-test	BB-sub	AGIEval	ARC-E	ARC-C	AVG
In-domain?	✓	✗	✗	✗	✗	
CodeLLaMA-7B + Std-CoT	<b>56.2</b>	<u>29.7</u>	19.2	<u>42.0</u>	<u>32.2</u>	<u>35.9</u>
CodeLLaMA-7B + Step-by-step	40.7	<u>29.0</u>	<u>23.9</u>	41.5	<b>32.8</b>	33.6
CodeLLaMA-7B + CasCoD	<u>54.8</u>	<b>35.4</b>	<b>25.8</b>	<b>42.9</b>	31.7	<b>38.1</b>
LLaMA3-8B + Std-CoT	<b>66.9</b>	33.9	32.7	69.8	60.2	52.7
LLaMA3-8B + Step-by-step	44.2	<u>35.5</u>	<u>38.8</u>	<u>83.7</u>	<u>70.7</u>	<u>54.5</u>
LLaMA3-8B + CasCoD	<u>65.2</u>	<b>42.9</b>	<b>40.1</b>	<b>87.2</b>	<b>74.0</b>	<b>61.9</b>
Mistral-7B-v0.2 + Std-CoT	<b>72.2</b>	37.6	<u>32.0</u>	68.8	57.9	<u>53.7</u>
Mistral-7B-v0.2 + Step-by-step	56.4	<u>38.9</u>	20.1	<u>76.4</u>	<u>62.3</u>	50.8
Mistral-7B-v0.2 + CasCoD	<u>71.7</u>	<b>42.5</b>	<b>40.1</b>	<b>83.9</b>	<b>74.2</b>	<b>62.5</b>

Table 3: Accuracy (%) on IND and OOD datasets with different student models distilled by different methods.




Figure 6: Ablation study on training data size for four OOD datasets. The dotted line indicates the performance of fine-tuning the student models by standard CoTs distillation using the full set (100% of) BBH-train dataset. The results in IND dataset can be found in Appendix A.2.2.

Method	ChatGPT	GPT4	AVG
Teacher	41.1	36.3	38.7
Std-CoT	40.8	29.8	35.3
SCOTT	36.2	29.4	32.8
MT-CoT	36.2	25.8	31
Step-by-step	6.6	-0.1	3.25
CasCoD (ours)	<b>40.8</b>	<b>31.6</b>	<b>36.2</b>

Table 4: Faithfulness (LAS, %) of the compared methods with different LLM evaluators on the IND dataset. The prompt templates can be found in Appendix C.2

## B Details of Experiment

### B.1 Dataset Statistics

Table 6, Table 7, Table 8 and Table 9 show the data statistics of AGIEval, ARC, BIG-Bench Hard (BBH) and BIG-Bench Sub (BB-sub)<sup>5</sup>, respectively.

<sup>5</sup>For detailed descriptions of the subtasks in BIG-Bench, please refer to <https://github.com/google/BIG-bench/>

### B.2 Hyperparameters Settings

In our study, we ensure consistency in the hyperparameter settings across all baselines, including our proposed CasCoD approach, to maintain the fairness of our comparative analysis. Here, we detail the hyperparameter configurations employed in our experiments.

**Training Steps and Batch Size.** The number of training steps is determined based on the size of the training dataset, the batch size, and the number of gradient accumulation steps required. We maintain a consistent batch size across all baselines to eliminate any performance discrepancies that could arise from varying batch sizes.

**Learning Rate.** Our exploratory experiments initially focus on the standard CoTs distillation method using the LLaMA-2 model, revealing that while the batch size had minimal impact on perfor-

[blob/main/bigbench/benchmark\\_tasks/README.md](blob/main/bigbench/benchmark_tasks/README.md).



Models & Tasks	elementary_math_qa	identify_math_theorems	DateUnderstanding	Multi-StepArithmetic	AVG
Std-CoT	11.0	20.7	82.0	8.0	30.4
SCOTT	12.0	34.0	58.0	4.0	27.0
MT-CoT	10.0	9.4	74.0	6.0	24.9
Step-by-step	17.0	32.1	68.0	0.0	29.3
CasCoD	<b>22.0</b>	<b>36.2</b>	<b>86.0</b>	<b>8.0</b>	<b>38.1</b>

Table 5: Accuracy (%) on mathematical reasoning tasks with different distillation methods.




Figure 7: Ablation study on model size in the IND (BBH-test). The dotted line indicates the performance of the teacher LLM under the Zero-shot-CoT setting.

No.	Task	Size	# Choices
1	AQuA-RAT	254	5
2	LogiQA-EN	651	4
3	LSAT-AR	230	5
4	LSAT-LR	510	5
5	LSAT-RC	269	5
6	SAT-Math	220	4
7	SAT-EN	206	4
8	SAT-EN (w/o Psg.)	206	4
	<b>Sum</b>	2546	-

Table 6: Statistics of AGIEval dataset.

911 mance, the learning rate was a critical factor. We  
912 test learning rates of  $1e-4$ ,  $2e-4$ , and  $3e-4$  and ob-  
913 serve optimal performance at  $2e-4$  across Std-CoT  
914 and other distillation baselines as well as our Cas-  
915 CoD. Therefore, we set the learning rate to  $2e-4$  for  
916 all methods involved in our study.

917 **Epochs and Evaluation Strategy.** Throughout  
918 our training process, we monitor the training loss  
919 curve and note that it generally plateaued by  
920 the 15th epoch, suggesting that the models have  
921 achieved convergence. Accordingly, we set the  
922 number of epochs to 15 for 7B models. The pro-  
923 cess of determining the number of epochs for other  
924 model sizes follows a similar pattern. To mitigate




Figure 8: Ablation study on training data size in the IND (BBH-test). The dotted line indicates the performance of fine-tuning the student models by standard CoT distillation using the full set (100% of) BBH-train dataset.

Task	Size	# Choices
ARC-E	2376	4-5
ARC-C	1172	4-5

Table 7: Statistics of ARC test dataset.

the potential risk of overfitting and to ensure that  
our evaluation reflects the most effective model  
configuration, we systematically select the check-  
points from the epoch that demonstrate the best  
performance on the IND task. These checkpoints  
are then used to evaluate performance on OOD  
tasks.

Finally, the detailed hyperparameters in training  
and inference can be found in Table 10 and Table  
11, respectively.

## C Prompts

### C.1 Prompts of Generating CoTs for ChatGPT

We use the prompt template shown in Table 12 to  
call the ChatGPT API to generate the CoTs for the  
BBH-train datasets.

No.	Task	Size	# Choices
1	Boolean Expressions	250	2
2	Causal Judgement	187	2
3	Date Understanding	250	6
4	Disambiguation QA	250	4
5	Dyck Languages	250	-
6	Formal Fallacies Syllogisms Negation	250	2
7	Geometric Shapes	250	11
8	Hyperbaton (Adjective Ordering)	250	2
9	Logical Deduction (3 objects)	250	3
10	Logical Deduction (5 objects)	250	5
11	Logical Deduction (7 objects)	250	7
12	Movie Recommendation	250	5
13	Multi-Step Arithmetic	250	-
14	Navigate	250	2
15	Object Counting	250	-
16	Penguins in a Table	146	5
17	Reasoning about Colored Objects	250	18
18	Ruin Names	250	11
19	Salient Translation Error Detection	250	6
20	Snarks	178	2
21	Sports Understanding	250	2
22	Temporal Sequences	250	4
23	Tracking Shuffled Objects (3 objects)	250	3
24	Tracking Shuffled Objects (5 objects)	250	5
25	Tracking Shuffled Objects (7 objects)	250	7
26	Web of Lies	250	2
27	Word Sorting	250	-
	<b>Sum</b>	6511	-

Table 8: Statistics of BIG-Bench Hard dataset.

## C.2 Prompts of Simulators

We use the prompt templates shown in Table 13 and Table 14 to call the ChatGPT and GPT4 API to predict the answers given a question or with an additional rationale, respectively.

## D Case Study

Here we show 4 cases in Table 15, 16, 17 and 18 to compare the CoT generated by CasCoD with the teacher LLM and the standard CoTs distillation method (Std-CoT). We utilize ✓ and ✗ to denote whether the CoT is correct or incorrect, respectively.

Table 15 and 16 show that while Std-CoT correctly predicts the final answer on in-domain tasks, it generates incorrect intermediate reasoning steps, indicating that Std-CoT causes student models to capture spurious correlations between questions and answers, without learning to reason with diversity. This is evidenced by the fact that its generated rationales are almost identical to the teacher’s CoTs. In contrast, CasCoD produces more distinct CoTs, differing from the teacher’s reasoning format, which indicates an enhancement in the student’s reasoning generalizability. Tables 17 and 18 illustrate that on out-of-domain tasks, Std-CoT fails to reason correctly, including logical and factual mistakes, whereas CasCoD can clearly understand the question and provides concise, logically rich

No.	Task	Size	# Choices
1	abstract_narrative_understanding	100	5
2	anachronisms	100	2
3	analogical_similarity	100	7
4	analytic_ entailment	70	2
5	cause_and_effect	100	2
6	checkmate_in_one	100	26
7	cifar10_classification	100	10
8	code_line_description	60	4
9	conceptual_combinations	100	4
10	crass_ai	44	4
11	elementary_math_qa	100	5
12	emoji_movie	100	5
13	empirical_judgments	99	3
14	english_russian_proverbs	80	4
15	entailed_polarity	100	2
16	entailed_polarity_hindi	100	2
17	epistemic_reasoning	100	2
18	evaluating_information_essentiality	68	5
19	fantasy_reasoning	100	2
20	figure_of_speech_detection	59	10
21	goal_step_wikihow	100	4
22	gre_reading_comprehension	31	5
23	human_organs_senses	42	4
24	identify_math_theorems	53	4
25	identify_odd_metaphor	47	5
26	implicatures	100	2
27	implicit_relations	82	25
28	indic_cause_and_effect	100	2
29	intersect_geometry	100	26
30	kanji_ascii	100	5
31	kannada	100	4
32	key_value_maps	100	2
33	logic_grid_puzzle	100	3
34	logical_args	32	5
35	logical_fallacy_detection	100	2
36	metaphor_boolean	100	2
37	metaphor_understanding	100	4
38	minute_mysteries_qa	100	4
39	mnist_ascii	100	10
40	moral_permissibility	100	2
41	movie_dialog_same_or_different	100	2
42	nonsense_words_grammar	50	4
43	odd_one_out	86	5
44	parsinlu_qa	100	4
45	physical_intuition	81	4
46	play_dialog_same_or_different	100	2
47	presuppositions_as_nli	100	3
48	riddle_sense	49	5
49	similarities_abstraction	76	4
50	simple_ethical_questions	100	4
51	social_iqa	100	3
52	strange_stories	100	2
53	strategyqa	100	2
54	swahili_english_proverbs	100	4
55	swedish_to_german_proverbs	72	4
56	symbol_interpretation	100	5
57	timedial	100	3
58	undo_permutation	100	5
59	unit_interpretation	100	5
60	vitaminc_fact_verification	100	3
61	winowhy	100	2
	<b>Sum</b>	5384	-

Table 9: Statistics of BIG-Bench sub dataset. We filter the original dataset by retrieving tasks with keywords "multiple choice" and randomly sample up to 100 examples per task. Note, the task in BBH will not be involved in BB-sub.

reasoning.

Hyperparameter	TinyLLaMA-1.1B	LLaMA2-7B	LLaMA2-13B
gradient accumulation steps	4	4	8
per device batch size	16	16	8
learning rate	2e-4	2e-4	2e-4
epoches	20	15	10
max length	1024	1024	1024
$\beta$ of AdamW	(0.9,0.999)	(0.9,0.999)	(0.9,0.999)
$\epsilon$ of AdamW	1e-8	1e-8	1e-8
$\gamma$ of Scheduler	0.95	0.95	0.95
weight decay	0	0	0
warmup ratio	0	0	0
rank of LoRA	64	64	64
$\alpha$ of LoRA	32	32	32
target modules	q_proj, v_proj	q_proj, v_proj	q_proj, v_proj
drop out of LoRA	0.05	0.05	0.05

Table 10: Training hyperparameters.

Arguments	Student	Teacher
do sample	False	True
temperature	-	0.2
top-p	1.0	1.0
top-k	-	-
max new tokens	1024	2048
# return sequences	1	1

Table 11: Generation configs of students and teachers.

---

{Task Description}. Your response should conclude with the format "Therefore, the answer is".

Q: {Task Example Question No.1}  
A: Let's think step by step. {Human-Curated-CoTs No.1}.

Q: {Task Example Question No.2}  
A: Let's think step by step. {Human-Curated-CoTs No.2}.

Q: {Task Example Question No.3}  
A: Let's think step by step. {Human-Curated-CoTs No.3}.

Q: {QUESTION}  
A: Let's think step by step.

---

Table 12: Prompt template of gpt-3.5-turbo for generating the CoTs data with 3 shots.

system content	You are a helpful and precise assistant for following the given instruction.
user content	[Instruction] {Please read the question and then give your answer based on the question without any explanations.} Task Description: {TASK_DESCRIPTION} Question: {QUESTION} Your Answer:

Table 13: Prompt template of simulators for predicting the answers when given the question.

system content	You are a helpful and precise assistant for following the given instruction.
user content	[Instruction] {Please read the question and the rationale, and then give your answer based on the question and the rationale without any explanations.} Task Description: {TASK_DESCRIPTION} Question: {QUESTION} Rationale: {RATIONALE} Your Answer:

Table 14: Prompt template of simulators for predicting the answers when given the question and rationale.



<b>Question:</b>	Evaluate the result of a random Boolean expression: not False and False or True or False.
<b>Teacher</b>	<p>Remember that (i) expressions inside brackets are always evaluated first and that (ii) the order of operations from highest priority to lowest priority is "not", "and", "or", respectively.</p> <p>We first simplify this expression "Z" as follows: "Z = not False and False or True or False = A and B or C or D" where "A = not False", "B = False", "C = True", and "D = False".</p> <p>Let's evaluate A: A = not False = not (False) = True.</p> <p>Plugging in A, B, C, and D, we get: Z = A and B or C or D = True and False or True or False = False or True or False = True.</p> <p>Therefore, the answer is True. ✓</p>
<b>Std-CoT</b>	<p>Remember that (i) expressions inside brackets are always evaluated first and that (ii) the order of operations from highest priority to lowest priority is "not", "and", "or", respectively.</p> <p>We first simplify this expression "Z" as follows: "Z = not False and False or True or False = A and B or C or D" where "A = not False", "B = False", "C = True", and "D = False".</p> <p>Let's evaluate A: A = not False = not (False) = True.</p> <p>Plugging in A, B, C, and D, we get: Z = A and B or C or D = True and False or True or False = False or False or False = False.</p> <p>Therefore, the answer is True. ✗</p>
<b>CasCoD</b>	<p>Removing the brackets first: (not False and False) or (True or False) or False.</p> <p>Let's evaluate these expressions one by one:</p> <ul style="list-style-type: none"> <li>- (not False and False): The boolean expression "not False and False" evaluates to True.</li> <li>- (True or False): The boolean expression "True or False" evaluates to True.</li> <li>- False: The boolean expression False evaluates to False.</li> </ul> <p>When we evaluate these expressions, we get True.</p> <p>Therefore, the answer is True. ✓</p>

Table 15: A boolean-expression case from BIG-Bench Hard.

<b>Question:</b>	<p>Answer extremely simple questions about the colors of objects on a surface. Q: On the table, you see several objects arranged in a row: a burgundy paperclip, a purple cup, a mauve pencil, a yellow keychain, and a silver teddy bear. What is the color of the object furthest from the teddy bear?</p> <p>Options: (A) red (B) orange (C) yellow (D) green (E) blue (F) brown (G) magenta (H) fuchsia (I) mauve (J) teal (K) turquoise (L) burgundy (M) silver (N) gold (O) black (P) grey (Q) purple (R) pink</p>
<b>Teacher</b>	<p>According to this question, the objects are arranged in a row, from left to right, as follows: (1) a burgundy paperclip, (2) a purple cup, (3) a mauve pencil, (4) a yellow keychain, (5) a silver teddy bear.</p> <p>The teddy bear is the last item on the list, namely (5). The object furthest from the teddy bear corresponds to (1), which is a burgundy paperclip.</p> <p>The color of the paperclip is burgundy. Therefore, the answer is (L). ✓</p>
<b>Std-CoT</b>	<p>According to this question, the objects are arranged in a row, from left to right, as follows: (1) a burgundy paperclip, (2) a purple cup, (3) a mauve pencil, (4) a yellow keychain, (5) a silver teddy bear. The object furthest from the teddy bear is the mauve pencil, which is the third item on the list.</p> <p>The color of the mauve pencil is mauve. Therefore, the answer is (L). ✗</p>
<b>CasCoD</b>	<p>Rationale: According to this question, the objects are arranged in a row, from left to right, as follows: (1) a burgundy paperclip, (2) a purple cup, (3) a mauve pencil, (4) a yellow keychain, (5) a silver teddy bear.</p> <p>The object furthest from the teddy bear corresponds to (1), which is a burgundy paperclip.</p> <p>The color of the paperclip is burgundy.</p> <p>Therefore, the answer is (L). ✓</p>

Table 16: A reasoning about colored objects case from BIG-Bench Hard.

<b>Question:</b>	$h = 3a + 28.6$ A pediatrician uses the model above to estimate the height $h$ of a boy, in inches, in terms of the boy's age $a$ , in years, between the ages of 2 and 5. Based on the model, what is the estimated increase, in inches, of a boy's height each year? Options: (A)3 (B)5.7 (C)9.5 (D)14.3
<b>Teacher</b>	<p>Step 1: We are given the model <math>h = 3a + 28.6</math> which represents the estimated height <math>h</math> of a boy in terms of his age <math>a</math>.</p> <p>Step 2: We need to find the estimated increase in height each year. This means we need to find the change in height for a change in age of 1 year.</p> <p>Step 3: To find the change in height, we can subtract the height at age <math>a + 1</math> from the height at age <math>a</math>. So, the change in height is <math>(3(a + 1) + 28.6) - (3a + 28.6)</math>.</p> <p>Step 4: Simplifying the expression, we get <math>3a + 3 + 28.6 - 3a - 28.6</math>.</p> <p>Step 5: The terms <math>3a</math> and <math>-3a</math> cancel out, leaving us with <math>3 + 3</math>.</p> <p>Step 6: The change in height is 6 inches.</p> <p>Therefore, the answer is (A) 3. <b>X</b></p>
<b>Std-CoT</b>	This question focuses on the increase in height each year. According to the model, the height of a boy between the ages of 2 and 5 is approximately 36 inches. Therefore, the estimated increase in height each year is 3 inches. Therefore, the answer is (A). <b>X</b>
<b>CasCoD</b>	<p>Rationale: This question asks about the estimated increase, in inches, of a boy's height each year. According to the model, the height of a boy increases by 3 inches each year between the ages of 2 and 5.</p> <p>Therefore, the answer is (A) 3 inches. <b>✓</b></p>

Table 17: A sat-math case from AGIEval.

<b>Question:</b>	Which gas is the most abundant in Earth's atmosphere? Options: (A) oxygen (B) nitrogen (C) water vapor (D) carbon dioxide										
<b>Teacher</b>	<p>First, we know that Earth's atmosphere is primarily composed of gases.</p> <p>Second, we need to determine which gas is the most abundant.</p> <p>Third, we can eliminate options (A) oxygen and (C) water vapor because they are present in smaller quantities compared to the other gases.</p> <p>Finally, we are left with options (B) nitrogen and (D) carbon dioxide.</p> <p>Considering that nitrogen makes up about 78% of Earth's atmosphere, while carbon dioxide only makes up about 0.04%, it is clear that nitrogen is the most abundant gas in Earth's atmosphere.</p> <p>Therefore, the answer is (B) nitrogen. <b>✓</b></p>										
<b>Std-CoT</b>	<p>The percentage of each gas in the atmosphere is given by the following table:</p> <table border="1"> <thead> <tr> <th>Gas</th> <th>Percentage</th> </tr> </thead> <tbody> <tr> <td>oxygen</td> <td>20.95</td> </tr> <tr> <td>nitrogen</td> <td>78.09</td> </tr> <tr> <td>water vapor</td> <td>0.4</td> </tr> <tr> <td>carbon dioxide</td> <td>0.04</td> </tr> </tbody> </table> <p>According to this table, oxygen is the most abundant gas in Earth's atmosphere. Therefore, the answer is (A). <b>X</b></p>	Gas	Percentage	oxygen	20.95	nitrogen	78.09	water vapor	0.4	carbon dioxide	0.04
Gas	Percentage										
oxygen	20.95										
nitrogen	78.09										
water vapor	0.4										
carbon dioxide	0.04										
<b>CasCoD</b>	<p>Rationale: According to the World Bank, the composition of Earth's atmosphere is approximately 78% nitrogen, 21% oxygen, 0.93% argon, and small amounts of other gases. Nitrogen is the most abundant gas in Earth's atmosphere.</p> <p>Therefore, the answer is (B). <b>✓</b></p>										

Table 18: A case from AI2 Reasoning Challenge.