

Small Models Struggle to Learn from Strong Reasoners

Yuetai Li^{*} Xiang Yue[◇] Zhangchen Xu^{*} Fengqing Jiang^{*} Luyao Niu^{*}

Bill Yuchen Lin^{*} Bhaskar Ramasubramanian^{*} Radha Poovendran^{*}

^{*}University of Washington [◇]Carnegie Mellon University [♣]Western Washington University
{yuetaili, zxu9, fqjiang, luyaoni, byuchenl, rp3}@uw.edu,
xyue2@andrew.cmu.edu, ramasub@wwu.edu

Huggingface: <https://huggingface.co/UWNSL>
Project Page: <https://small-model-gap.github.io/>

Abstract

Large language models (LLMs) excel in complex reasoning tasks, and distilling their reasoning capabilities into smaller models has shown promise. However, we uncover an interesting phenomenon, which we term the *Small Model Learnability Gap*: small models (≤ 3 B parameters) do not consistently benefit from long chain-of-thought (CoT) reasoning or distillation from larger models. Instead, they perform better when fine-tuned on shorter, simpler reasoning chains that better align with their intrinsic learning capacity. To address this, we propose Mix Distillation, a simple yet effective strategy that balances reasoning complexity by combining long and short CoT examples or reasoning from both larger and smaller models. Our experiments demonstrate that Mix Distillation significantly improves small model reasoning performance compared to training on either data alone. These findings highlight the limitations of direct strong model distillation and underscore the importance of adapting reasoning complexity for effective reasoning capability transfer.

1 Introduction

Large language models (LLMs) (Anthropic, 2023; Brown et al., 2020; OpenAI, 2023; Touvron et al., 2023a) have demonstrated remarkable performance in complex reasoning tasks, enabling advancements in mathematical problem-solving, logical inference, and structured decision-making (Cobbe et al., 2021; Shao et al., 2024; Yang et al., 2024). A key advancement in improving LLM complex reasoning capability is the chain-of-thought (CoT) prompting. This technique decomposes complex problems into intermediate reasoning steps, enhancing both performance and interpretability. (Wei et al., 2023).

However, the high computational cost of LLMs hinders their deployment on resource-constrained devices, motivating the development of smaller models that offer similar capabilities at reduced

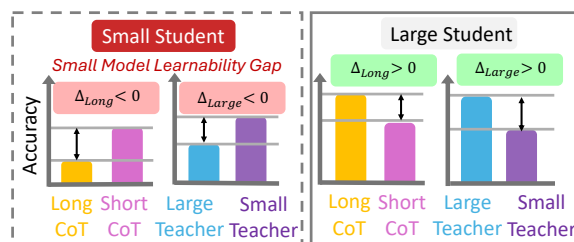


Figure 1: Small student models (≤ 3 B parameters) do not consistently benefit from long CoT reasoning or distillation from large teacher models. Instead, they perform better when fine-tuned on shorter CoT reasoning or distilled from smaller teachers, which better matches their intrinsic learning capacity. We term this phenomenon the *Small Model Learnability Gap*.

cost. A widely adopted strategy to achieve this is distillation (Agarwal et al., 2024; Hinton et al., 2015a; Kim et al., 2024a), where CoT sequences generated by a strong teacher model are used to fine-tune a weaker student model. Naturally, one might expect that distilling CoT sequences from stronger models would consistently improve small models’ complex reasoning capabilities (Agarwal et al., 2024; DeepSeek-AI et al., 2024; Min et al., 2024; Tunstall et al., 2023).

However, we reveal an interesting phenomenon, which we term the *Small Model Learnability Gap* (Fig. 1): small models do not consistently benefit from the complex reasoning sequences provided by strong teachers, such as long CoT reasoning or distillation from large models. In our experiments, we observe that when small models are exposed to long and intricate reasoning traces, they struggle to internalize the multi-step logic due to their constrained ability. Instead, small models perform better when fine-tuned on *shorter, simpler reasoning chains* that align more closely with their intrinsic learning capacity. This suggests that small models struggle to process overly elaborate reasoning traces or adapt to the distribution shifts introduced by stronger teachers, ultimately limiting their abil-

ity to generalize effectively.

To address the challenge described above, we propose *Mix Distillation*, a simple yet effective approach that balances reasoning complexity by blending different types of reasoning traces. Specifically, our method comprises two configurations: (1) *Mix-Long* – A combination of long and short CoT examples, ensuring that small models are exposed to both detailed and concise reasoning steps. (2) *Mix-Large* – A mixture of responses from both larger and smaller models, allowing small models to learn from reasoning chains that are better suited to their capacity.

Our experiments demonstrate that *Mix Distillation* consistently improves small model reasoning performance compared to standard distillation. For instance, Qwen2.5-3B-Instruct improves by more than 8 points on MATH and AMC using Mix-Long, compared to direct training on long CoT data. Qwen2.5-3B-Instruct gains more than 7 points on MATH, AIME and AMC using Mix-Large compared with training on large teacher CoT data.

These findings highlight a fundamental limitation of direct strong model distillation and emphasize the importance of *adapting reasoning complexity* for effective knowledge transfer. By carefully designing fine-tuning strategies, we provide new insights into overcoming the constraints of small model learning, making them more effective at reasoning-intensive tasks.

2 Preliminaries

2.1 Notation

Let $x = (x_1, x_2, \dots, x_n)$ represent an input sequence (e.g., a prompt), and $y = (y_1, y_2, \dots, y_m)$ be the corresponding output sequence. We consider a LLM parameterized by θ , which predicts the next token following a conditional distribution $\pi_\theta(y_t|x, y_{1:t-1})$. We denote by $\text{CoT}(y) \subseteq y$ the subset of tokens in the generated output that encodes a *chain-of-thought*, often serving as a reasoning trace or explanatory sequence.

Throughout this work, we use the term **short CoT**, to describe concise reasoning paths to arrive at solutions (Min et al., 2024; Yeo et al., 2025) and **long CoT** to describe an extended reasoning sequence that is not only longer but also demonstrates more complex reflective thoughts (Qwen, 2024b; Yeo et al., 2025). Additionally, we use the term **large teacher CoT** to refer to the reasoning trace generated by a larger teacher model, and the

term **small teacher CoT** for the reasoning steps produced by a smaller teacher model. Please see Appendix D for more examples.

2.2 Supervised Fine-Tuning (SFT)

Supervised fine-tuning (SFT) is widely adopted to enhance reasoning capabilities of LLMs on a dataset $\mathcal{D} = \{(x^i, y^i)\}_{i=1}^N$, where y^i can be short CoT, long CoT, strong model CoT or weak model CoT sequences. The SFT process updates the parameters θ of a language model by minimization the negative log-likelihood loss over the instruction dataset \mathcal{D} .

3 Small Model Learnability Gap

In this section, we fine-tune student models using different CoT data. We then reveal the small model learnability gap given the performance of fine-tuned models.

3.1 Experiment Setup

Datasets. We use the 7,500 prompt set of MATH (Hendrycks et al., 2021). This dataset encompasses seven math topics such as advanced calculus, geometry, and linear algebra.

Student models. Our study considers ten student models from the Qwen (Qwen, 2024a) and Llama (Meta, 2024a,b) model families of varying sizes. These models include the Instruct version of Qwen2.5-0.5B, Qwen2.5-1.5B, Qwen2.5-3B, Qwen2.5-7B, Qwen2.5-14B, and Qwen2.5-32B, and the Instruct version of Llama3.2-1B, Llama3.2-3B, Llama3.1-8B, and Llama3.3-70B. A comprehensive overview of the student models is presented in Table 4 of Appendix A.

Teacher models. To compare long CoT with short CoT, we use QwQ-32B-Preview (Qwen, 2024b) to generate long CoT sequences and Qwen2.5-32B-Instruct as the response generator for short CoT. Within each model family, we designate the larger scale model as the large teacher and the smaller scale model as the small teacher. This includes Qwen2.5-72B-Instruct vs Qwen2.5-3B-Instruct, Llama3.1-70B-Instruct vs Llama3.1-8B-Instruct, and Gemma2-27B-it vs Gemma2-9B-it.

Evaluation Benchmarks. We evaluate the reasoning capability of fine-tuned student models on a set of commonly used benchmarks, including MATH (Hendrycks et al., 2021), GSM8K (Cobbe

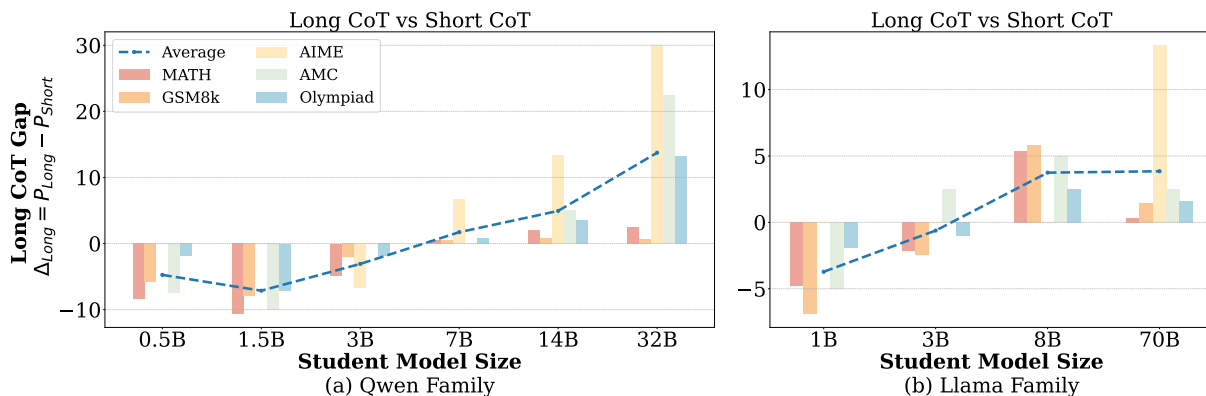


Figure 2: Long CoT Gap ($\Delta_{Long} = P_{Long} - P_{Short}$) of student models with different models sizes for (a) Qwen family (b) Llama family. For teacher models, QwQ-*preview*-32B is chosen to generate long CoT responses, while Qwen2.5-32B-*Instruct* is chosen to generate short CoT responses. Negative (Positive) Δ_{Long} indicates that long CoT is worse (better) than short CoT. Our results demonstrate that short CoT is better for smaller student models (indicated by $\Delta_{Long} < 0$), while long CoT is better for larger student models (indicated by $\Delta_{Long} > 0$).

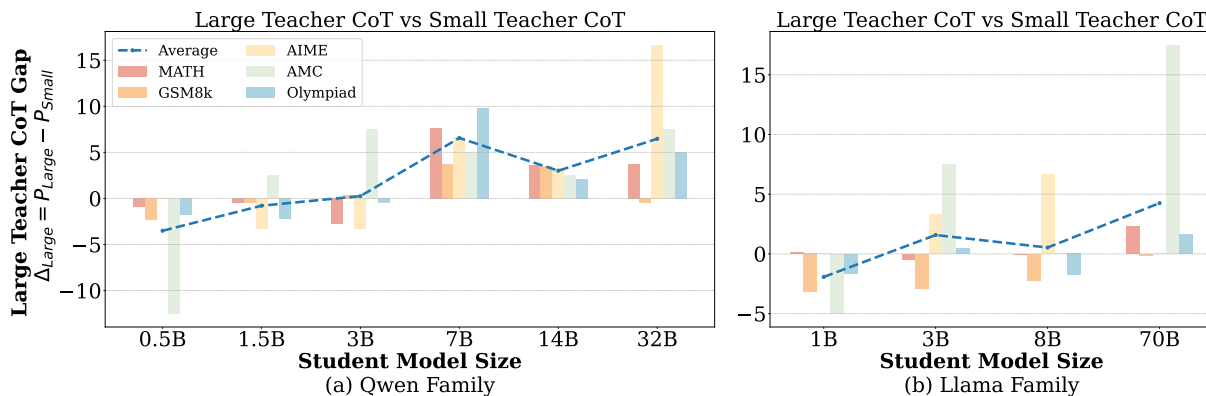


Figure 3: Large model CoT Gap ($\Delta_{Large} = P_{Large} - P_{Small}$) of student models with different models sizes for (a) Qwen family (b) Llama family. For teacher models, Qwen2.5-72B-*Instruct* is chosen as the large teacher to generate responses, while Qwen2.5-3B-*Instruct* is chosen as the small teacher to generate responses. Negative (positive) Δ_{Large} indicates that large teacher CoT is worse (better) than small teacher CoT. Our results demonstrate that small teacher CoT is better for smaller student models (indicated by $\Delta_{Large} < 0$), while large model CoT is better for larger student models (indicated by $\Delta_{Large} > 0$).

et al., 2021), AMC 2023, AIME 2024, and the English math subset of OlympiadBench (He et al., 2024). These benchmarks span a wide range of challenge levels, from elementary mathematics to advanced competition problems. We define the student model performance as the average score on five benchmarks. Unless otherwise specified, all fine-tuned models are evaluated in a zero-shot setting using greedy decoding. We set the maximum generation tokens as 16k. Please see Appendix A for detailed experimental setup.

We define the following performance scores:

- P_{Long} : Performance score of a student model fine-tuned on long CoT data.
- P_{Short} : Performance score of a student model

fine-tuned on short CoT data.

- P_{Large} : Performance score of a student model fine-tuned on CoT from a larger teacher.
- P_{Small} : Performance score of a student model fine-tuned on CoT from a smaller teacher.

Training Setup. Teacher models generate responses by rejection sampling (Dong et al., 2023; Gulcehre et al., 2023; Tong et al., 2024; Yuan et al., 2023; Yue et al., 2023; Zelikman et al., 2022) By default, teacher models employ greedy decoding. By combining the math problem instructions with corresponding solutions generated by teacher models, we construct problem-solution pairs to fine-tune student models. We train the models using the LLaMA-Factory framework (Zheng et al., 2024).

For student models of scale less than 14B, we use full-parameter SFT and implement a cosine learning rate schedule with a maximum learning rate of 10^{-5} to fine-tune student models for two epochs (Touvron et al., 2023b). For student models larger than 14B, we adopt LoRA fine-tuning with a learning rate of 10^{-4} for two epochs. Detailed hyperparameters and information about the experimental platform are provided in Appendix A.

3.2 Long CoT Gap

This section evaluates the reasoning capabilities of student models fine-tuned over long CoT data and short CoT data. We quantify the performance difference between long and short CoT data using *long CoT gap* Δ_{Long} , defined as:

$$\Delta_{Long} = P_{Long} - P_{Short}.$$

Figure 2 provides a comprehensive overview of the long CoT gap Δ_{Long} across different student models. The detailed benchmark scores on MATH, GSM8K, AIME, AMC, and OlympiadBench are deferred to Table 8 in Appendix B. We report the following key takeaways.

Takeaway 1: Long CoT Gap

Small student models tend to benefit more from short CoT, while large student models gain greater advantages from long CoT.

We observe that long CoT is more effective for larger models, consistently leading to improved performance across most math benchmarks. For example, the student model Qwen2.5-32B-Instruct improves about 15 points across all math metrics on average.

However, long CoT data is not effective for smaller models, yielding significantly less improvement compared to short CoT. On the MATH and AMC benchmarks, student model Qwen2.5-1.5B-Instruct performs over 10 points lower when fine-tuned with long CoT data. This shows that smaller models may not be able to effectively learn and utilize the long CoT paradigm. Please see more attribution analysis in Section 3.4.

3.3 Large Teacher CoT Gap

We investigate how effective small models may learn from large teacher and small teachers. We define a *large teacher CoT gap* as:

$$\Delta_{Large} = P_{Large} - P_{Small}.$$

Student Model	P_{Long}	P_{Short}	Δ_{Long}	Better?
Qwen2.5-0.5B	14.8	19.5	-4.7	Short
Qwen2.5-1.5B	27.0	34.2	-7.1	Short
Qwen2.5-3B	40.3	43.4	-3.1	Short
Qwen2.5-7B	48.9	47.2	1.7	Long
Qwen2.5-14B	59.2	54.3	4.9	Long
Qwen2.5-32B	73.0	59.3	13.7	Long
Llama-3.2-1B	15.8	19.5	-3.7	Short
Llama-3.2-3B	32.5	33.1	-0.6	Short
Llama-3.1-8B	35.2	31.5	3.7	Long
Llama-3.3-70B	58.2	54.3	3.8	Long

Table 1: Comparison of the average performance between fine-tuning with long CoT (P_{Long}) and short CoT (P_{Short}). We find that small student models may struggle to learn from long CoT data.

Student Model	P_{Large}	P_{Small}	Δ_{Large}	Better?
Qwen2.5-0.5B	16.9	20.4	-3.5	Weak
Qwen2.5-1.5B	32.2	33.0	-0.8	Weak
Qwen2.5-3B	39.7	39.4	0.3	Strong
Qwen2.5-7B	48.9	42.3	6.6	Strong
Qwen2.5-14B	52.9	49.9	3.0	Strong
Qwen2.5-32B	59.5	53.0	6.5	Strong
Llama-3.2-1B	16.5	18.5	-1.9	Weak
Llama-3.2-3B	32.8	31.2	1.6	Strong
Llama-3.2-8B	25.6	25.1	0.5	Strong
Llama-3.2-70B	57.6	53.3	4.3	Strong

Table 2: Comparison of average performance between fine-tuning with large teacher CoT (P_{Large}) and small teacher CoT (P_{Small}). We find that small student models may struggle to learn from large teacher CoT data.

Figure 3 provides a comprehensive comparison of the Δ_{Large} incurred by all student models. The detailed benchmark scores of MATH, GSM8K, AIME, AMC and OlympiadBench are deferred to Table 7 in Appendix B. More experimental results of different teacher models, including Llama3.1-70B vs Llama3.1-8B and Gemma2-27B vs Gemma2-9B are in Table 9 of Appendix B.

We observe that larger student models learn effectively from large teacher CoT. For example, Qwen2.5-7B-Instruct and Qwen2.5-32B-Instruct student models improve over 5 points on average, with Qwen2.5-32B-Instruct achieving more than a 15 point increase on the AIMC benchmark. However, smaller models do not learn effectively from large teacher models such as Qwen2.5-72B-Instruct. Instead, small teacher models such as Qwen2.5-3B-Instruct may serve as better teacher models for small student models. For instance, the performance of Qwen2.5-0.5B-Instruct degrades by more than 10 points on the AMC benchmark.

Student Model	Distillation Method	MATH	AMC	GSM8k	Olympiad Bench	AIME	Average
Qwen2.5-3B	Long CoT	56.2	37.5	80.0	24.4	<u>3.3</u>	40.3
	Short CoT	61.0	37.5	82.0	26.4	10.0	43.4
	Strong Model CoT	57.5	35.0	80.0	25.9	0.0	39.7
	Weak Model CoT	60.3	27.5	79.5	26.4	<u>3.3</u>	39.4
	Deepseek-R1-32B (Long CoT)	50.7	20.0	81.2	15.7	0.0	33.5
	<i>Ours</i>						
	Mix-Long	<u>64.7</u>	45.0	<u>81.4</u>	<u>28.6</u>	10.0	45.9
Mix-Large	65.8	<u>42.5</u>	81.7	29.0	10.0	<u>45.8</u>	
Llama3.2-3B	Long CoT	48.7	17.5	75.1	<u>17.6</u>	<u>3.3</u>	32.5
	Short CoT	50.9	15.0	77.5	18.7	<u>3.3</u>	33.1
	Strong Model CoT	47.4	25.0	71.2	16.9	<u>3.3</u>	32.8
	Weak Model CoT	47.9	17.5	74.1	16.4	<u>3.3</u>	31.2
	Deepseek-R1-32B (Long CoT)	48.5	17.5	<u>77.7</u>	16.1	6.7	33.3
	<i>Ours</i>						
	Mix-Long	53.0	<u>22.5</u>	79.4	17.2	<u>3.3</u>	35.1
Mix-Large	<u>51.8</u>	25.0	76.3	17.2	<u>3.3</u>	<u>34.7</u>	

Table 3: **Mix Distillation** outperforms the baseline models across most metrics. We use Llama3.2-3B-Instruct and Qwen2.5-3B-Instruct as the student model and 7.5k samples in MATH dataset as the training set. We distill different teacher models to generate responses as the baseline. Our proposed Mix-Long combines long CoT data and normal CoT data in a 1:4 ratio, while Mix-Large combines strong model response and weak model response with the same proportion. Experimental results demonstrate that both Mix-Long and Mix-Large surpass baselines in most evaluation metrics. The highest score is bolded, and the second highest score is underlined.

We remark that both larger teachers and small teachers generate short CoT data in this section to fine-tune student models, with no significant difference in average length. Specifically, the average token length is 432.98 for the 72B teacher and 440.70 for the 3B teacher. This helps eliminating CoT length as a confounding variable in our results of the large teacher CoT gap.

Note that prior studies (Kim et al., 2024b) also demonstrated that stronger models are not necessarily stronger teachers, emphasizing response generator and teacher-side factors. Our work differs in that we attribute this phenomenon primarily to the size of the student model.

Takeaway 2: Large Teacher CoT Gap

Small student models tend to learn better from small teachers, while large student models benefit more from large teachers.

Qwen2.5-1.5B-Instruct. Our findings show that the long CoT gap of the small math expert model is significantly smaller than that of general small models. Furthermore, the performance improvement of Qwen2.5-Math-1.5B when fine-tuned with large teacher CoT exceeds that of Qwen2.5-1.5B, suggesting that math expert models benefit more substantially from large teacher CoT. We conjecture that a key factor leading to the small model learnability gap is the *limited in-domain knowledge of small student models*. We summarize this observation in the following takeaway.

Takeaway 3: Effect of Domain Knowledge

Limited domain knowledge of small models may hinder their learning from strong reasoning teachers.

3.4 Analysis of Small Model Learnability Gap

Domain knowledge affects learnability gap.

We observe that math expert models, in spite of small model size, exhibit a smaller learnability gap for both long CoT and large teacher CoT data compared to general models in Figure 4. Specifically, we compare the learnability gaps between the student models Qwen2.5-Math-1.5B-Instruct and

Base models exhibit a more significant learnability gap.

We observe that base models generally exhibit a more significant learnability gap than Instruct models in Figure 5. This suggests that it is more challenging for small base models to effectively learn from long CoT data or large teacher CoT.

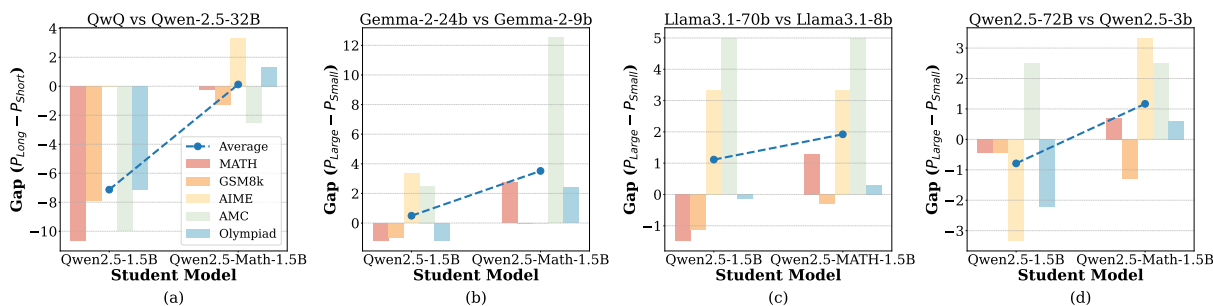


Figure 4: Math expert models usually have a less significant Learnability Gap than the general models. A positive Gap means long CoT or large teacher CoT is better while negative means worse. This indicates that the math expert model could more easily learn from long CoT data or large teacher CoT.

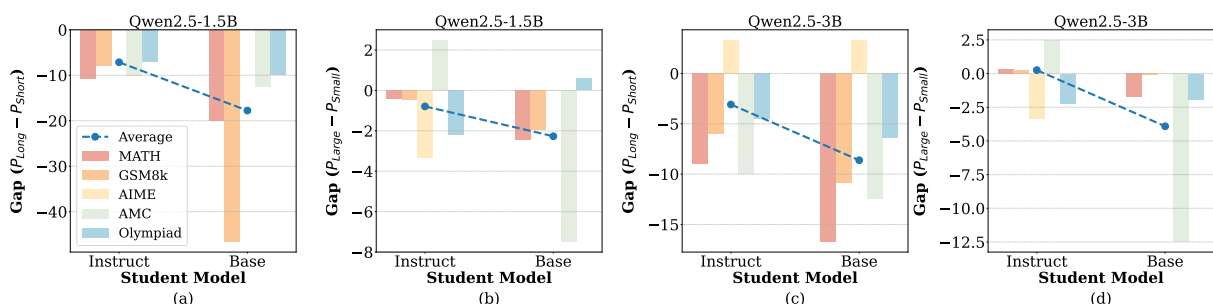


Figure 5: Base models generally exhibit a more significant learnability gap than Instruct models. A positive gap indicates that long CoT data or large teacher CoT enhance performance, whereas a negative gap suggests they have the opposite effect. This implies that it is more challenging for small base models to effectively learn from long CoT data or large teacher CoT.

Takeaway 4: Base vs Instruct

Small base models experience more significant learnability gap than Instruct models.

Speaking styles shift. We adopt the method from (Lin et al., 2023) to evaluate the rank shift of each token before and after fine-tuning on long CoT and Large teacher CoT data. This allows us to compare the token distribution shifts induced by the fine-tuning process. We then annotate the tokens that exhibit the largest rank shifts as the most shifted tokens. Our analysis reveals that these tokens are predominantly associated with expressive and stylistic elements, such as “wait”, “But”, and “Let”. Please see Appendix C for more details.

Takeaway 5: Speaking Styles Shift

Long CoT and large teacher CoT primarily shift the student model’s distribution of tokens associated with speaking styles.

Distribution Gap between student and teacher models. One contributing factor to Small Model Learnability Gap is that small models struggle to

process overly elaborate reasoning traces or adapt to the distribution shifts introduced by larger teachers, ultimately limiting their ability to generalize effectively. We present additional experimental results in Appendix B.3 to show the distribution gap between student and teacher models.

4 Mix Distillation: Bridge Small Model Learnability Gap

This section presents our Mix Distillation approach to bridge the small model learnability gap.

4.1 Mix Distillation

We propose *Mix Distillation* to address the learnability gap observed in small models. This approach blends easier-to-learn data with more challenging data for small models, thereby leveraging the strengths of both.

Our insight is that small models tend to perform better on data that closely matches their inherent distribution (such as short CoT or small teacher CoT), while they struggle with data that exhibits greater distribution shifts. The token distribution of the mixed long CoT and large teacher CoT data may become closer to that of small models’ in-

herent distribution, thereby enabling them to learn more effectively from challenging datasets.

We propose Mix-Long, which combines long and short CoT data with a weight of long CoT α and short CoT $1 - \alpha$. Similarly, we propose Mix-Large, which mixes large teacher CoT with a weight of α and small teacher CoT with a weight of $1 - \alpha$.

4.2 Experiment Results

We use Qwen2.5-3B-Instruct as the student model and MATH (7.5k) as the training set. We distill different teacher models to generate responses as the baseline. They include QwQ-32B (long CoT), Qwen2.5-32B (short CoT), Qwen2.5-72B (large teacher CoT), Qwen2.5-3B (small teacher CoT). We add Deepseek-R1-32B (DeepSeek-AI, 2025) as the teacher model to generate another set of long CoT data as baseline. We set $\alpha = 0.2$ in both configurations of Mix-Long and Mix-Large.

Experimental results demonstrate that both Mix-Long and Mix-Large surpass baselines in most evaluation metrics. We show that the small student model could achieve improved performance by Mix Distillation compared to training on a single dataset. For instance, Qwen2.5-3B-Instruct improves by more than 8 points on MATH and AMC using Mix-Long, compared to direct training on long CoT data. It also shows a more than 7-point gain on MATH, AIME and AMC for Qwen2.5-3B-Instruct by Mix-Large compared with training on large teacher CoT data. This implies that it is easier for small student models to learn from datasets generated by Mix Distillation.

Takeaway 6: Mix Distillation Bridges Gap

By mixing long CoT data (resp. large teacher CoTs) and short CoT data (resp. small teacher CoT), the small student model could achieve better performance compared to training on either data alone.

Figure 6 shows the average performance when taking different mix weight α of long CoT data or large teacher CoT. We choose Qwen2.5-3B-Instruct as the student model and find that a weight α of 0.2 achieves the highest average performance across five benchmarks for both Mix-Long and Mix-Large.

Interestingly, we find that after mixing long CoT and short CoT data, the small student model’s out-

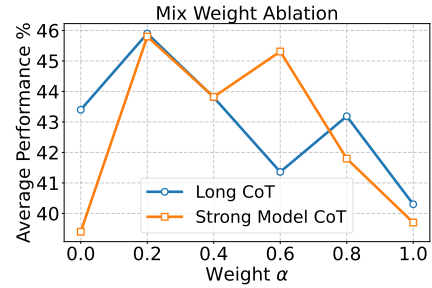


Figure 6: The average performance varies with the mix weight of long CoT or large teacher CoT data. Qwen2.5-3B-Instruct is chosen as the student model. At a weight of 0.2, mix distillation achieves the highest average performance.

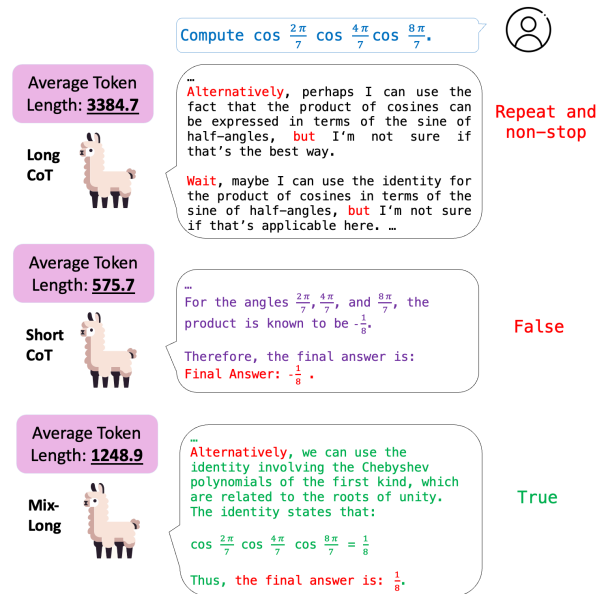


Figure 7: Case Study of Mix-Long. Models fine-tuned on long CoT tended to overthink, while those trained on short CoT produced incorrect answers. In contrast, Mix-Long, incorporating branching elements (e.g., “Alternatively”), achieved a balanced reasoning process and arrived at the correct answer.

put incorporates characteristics of long CoT, such as a branching process, while maintaining a reduced token length and avoiding overly elaborate thinking. This is illustrated in Figure 7. We observed that the small student model fine-tuned on long CoT data becomes overwhelmed by repeated thoughts and fails to stop, whereas the model fine-tuned on short CoT data produces incorrect answers. In contrast, our proposed Mix-Long, which incorporates branching elements (e.g., the use of “Alternatively”), delivers the correct answer. Additionally, the average token lengths of responses generated by long CoT, short CoT, and Mix-Long are 3384.7, 575.7, and 1248.9, respectively. We suggest that mixing long CoT and short CoT data is a practical approach to achieving a balanced CoT

length, thereby enhancing the reasoning capabilities of small student models.

5 Related Work

5.1 Chain-of-Thought

Early research on CoT primarily focused on short CoT, where models produce succinct reasoning paths to reach a solution (Lambert et al., 2025; Longpre et al., 2023; Wei et al., 2023; Yu et al., 2024). Recently, researchers have turned to long CoT prompting, which encourages the generation of extended and detailed reasoning chains (DeepSeek-AI, 2025; Hou et al., 2025; Kimi Team, 2025; NovaSky, 2025; OpenAI, 2024; Pan et al., 2025; Zeng et al., 2025). The model systematically explores multiple paths (branching) and reverts to earlier points if a particular path proves wrong (backtracking). Although several studies have investigated methods such as distillation and reinforcement learning to integrate long CoT capabilities into LLMs, these efforts have predominantly concentrated on large models. In contrast, our work specifically targets the challenges associated with training smaller models.

5.2 Synthetic Reasoning Data

Although human-crafted reasoning datasets have been used to enhance LLM reasoning capabilities (Hendrycks et al., 2021; LI et al., 2024), their development is both time-consuming and labor-intensive. Recent advancements have streamlined this process by generating instructions or responses directly from LLMs (Hui et al., 2024; Toshniwal et al., 2024; Xu et al., 2024a; Yue et al., 2023; Zhang et al., 2025) or extracting data directly from web (Paster et al., 2023; Yue et al., 2024), yielding more detailed and diverse chain-of-thought reasoning pathways. Recent study has investigated the impact of various response generators (Kim et al., 2024b), suggesting that in the domains of instruction following and reasoning, responses from stronger teacher models do not necessarily produce the most effective learning effects for student models. However, these investigations have not recognized student model size as a critical factor influencing this phenomenon, nor have they performed the more attribution and mitigation analyses as in this paper.

5.3 Distillation

Knowledge distillation has been extensively employed to transfer knowledge from large teacher

models to smaller student models (Hinton et al., 2015b; Phan et al., 2024; Gou et al., 2021; Xu et al., 2024b). Recent research in LLMs has increasingly adopted token-level distillation as an alternative to traditional logit-level distillation approaches (Phan et al., 2024; Ho et al., 2023; Agarwal et al., 2024). In conventional classification tasks, several studies have investigated the capacity gap phenomenon, where excessive differences in capacity between teacher and student models can compromise distillation effectiveness (Mirzadeh et al., 2019; Cho and Hariharan, 2019; Zhang et al., 2023). However, these work focused primarily on classification tasks, our work investigates reasoning generation tasks, where student models must internalize complex CoT reasoning traces.

6 Conclusion and Future Work

In this paper, we showed that long CoT data and large model responses were not uniformly beneficial for small student models. We found that small models may perform better when fine-tuned with short CoT and small model CoT. We termed this challenge as the Small Model Learnability Gap. The reason behind it may be that small student models excel on data that closely match their inherent distribution but struggle with significant distribution shifts. To bridge the gap, we introduced Mix Distillation, including Mix-Long, which combined long CoT and short CoT data in a ratio, and Mix-Large, which integrated large and small teacher CoT. Experimental results showed that both Mix-Long and Mix-Large outperform baselines across most evaluation metrics, which implied mix distillation outperforms training on a single data distribution. This paper provided practical insights for optimizing post-training strategies to enhance small language model reasoning capability.

We will explore several promising directions as future work. First, we will refine mix distillation by optimally combining diverse data sources and proposing more fine-grained mixing algorithms to boost reasoning capabilities. Second, we propose to study how strong reasoning teachers can generate data that is better suited for tuning small student models, thereby facilitating more effective knowledge transfer. Third, we will conduct further theoretical and model interpolability studies on the small model learnability gap. Lastly, we will investigate which SFT methods yield the best initial policies for subsequent RL procedure.

Limitations

While our study provides valuable insights into the understanding of small model learnability gap in math reasoning, we acknowledge that our research primarily focuses on this specific domain and does not evaluate other crucial skills such as instruction following, code generation, or multi-modal understanding. We also did not investigate the impact of fine-grained variations in pre-training data composition on the small model learnability gap. A more detailed analysis of how different pre-training data sources and their proportions affect learning outcomes could offer valuable insights into optimizing data selection strategies for mitigating this gap.

Ethical Statement

This paper focuses on the evaluation and enhancement of reasoning capabilities in small language models through distillation techniques. The dataset and benchmarks used in our experiments are publicly available. We do not introduce or endorse any applications that could cause harm or be misused. This paper does not present any ethical concerns.

Acknowledgment

This work is partially supported by the Air Force Office of Scientific Research (AFOSR) under grant FA9550-23-1-0208, the Office of Naval Research (ONR) under grant N0014-23-1-2386, and the National Science Foundation (NSF) AI Institute for Agent-based Cyber Threat Intelligence and Operation (ACTION) under grant IIS 2229876. Results presented in this paper were partially obtained using the Chameleon testbed (Keahey et al., 2020) supported by the National Science Foundation.

This work is supported in part by funds provided by the National Science Foundation, Department of Homeland Security, and IBM. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF or its federal agency and industry partners.

References

Rishabh Agarwal, Nino Vieillard, Yongchao Zhou, Piotr Stanczyk, Sabela Ramos, Matthieu Geist, and Olivier Bachem. 2024. [On-policy distillation of language models: Learning from self-generated mistakes](#). *Preprint*, arXiv:2306.13649.

Anthropic. 2023. [Introducing claude](#).

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](#). *Preprint*, arXiv:2005.14165.

Jang Hyun Cho and Bharath Hariharan. 2019. [On the efficacy of knowledge distillation](#). In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4794–4802.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.

DeepSeek-AI. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *Preprint*, arXiv:2501.12948.

DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, and et al. 2024. [Deepseek-v3 technical report](#). *Preprint*, arXiv:2412.19437.

Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao, Jipeng Zhang, Kashun Shum, and Tong Zhang. 2023. [Raft: Reward ranked finetuning for generative foundation model alignment](#). *Preprint*, arXiv:2304.06767.

Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. 2021. [Knowledge distillation: A survey](#). *International Journal of Computer Vision*, 129(6):1789–1819.

Caglar Gulcehre, Tom Le Paine, Srivatsan Srinivasan, Ksenia Konyushkova, Lotte Weerts, Abhishek Sharma, Aditya Siddhant, Alex Ahern, Miaosun Wang, Chenjie Gu, Wolfgang Macherey, Arnaud Doucet, Orhan Firat, and Nando de Freitas. 2023. [Reinforced self-training \(rest\) for language modeling](#). *Preprint*, arXiv:2308.08998.

Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun. 2024. [Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems](#). *Preprint*, arXiv:2402.14008.

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](#). *Preprint*, arXiv:2103.03874.

- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015a. [Distilling the knowledge in a neural network](#). *Preprint*, arXiv:1503.02531.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015b. [Distilling the knowledge in a neural network](#). *arXiv preprint arXiv:1503.02531*.
- Namgyu Ho, Laura Schmid, and Se-Young Yun. 2023. [Large language models are reasoning teachers](#). *Preprint*, arXiv:2212.10071.
- Zhenyu Hou, Xin Lv, Rui Lu, Jiajie Zhang, Yujiang Li, Zijun Yao, Juanzi Li, Jie Tang, and Yuxiao Dong. 2025. [Advancing language model reasoning through reinforcement learning and inference scaling](#). *Preprint*, arXiv:2501.11651.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *Preprint*, arXiv:2106.09685.
- Tingfeng Hui, Lulu Zhao, Guanting Dong, Yaqi Zhang, Hua Zhou, and Sen Su. 2024. [Smaller language models are better instruction evolvers](#). *Preprint*, arXiv:2412.11231.
- Kate Keahey, Jason Anderson, Zhuo Zhen, Pierre Riteau, Paul Ruth, Dan Stanzione, Mert Cevik, Jacob Colleran, Haryadi S. Gunawi, Cody Hammock, Joe Mambretti, Alexander Barnes, François Halbach, Alex Rocha, and Joe Stubbs. 2020. Lessons learned from the chameleon testbed. In *Proceedings of the 2020 USENIX Annual Technical Conference (USENIX ATC '20)*. USENIX Association.
- Gyeongman Kim, Doohyuk Jang, and Eunho Yang. 2024a. [Promptkd: Distilling student-friendly knowledge for generative language models via prompt tuning](#). *Preprint*, arXiv:2402.12842.
- Seungone Kim, Juyoung Suk, Xiang Yue, Vijay Viswanathan, Seongyun Lee, Yizhong Wang, Kiril Gashteovski, Carolin Lawrence, Sean Welleck, and Graham Neubig. 2024b. [Evaluating language models as synthetic data generators](#). *Preprint*, arXiv:2412.03679.
- Kimi Team. 2025. [Kimi k1.5: Scaling reinforcement learning with llms](#). *Preprint*, arXiv:2501.12599.
- Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafford, Chris Wilhelm, Luca Soldaini, Noah A. Smith, Yizhong Wang, Pradeep Dasigi, and Hannaneh Hajishirzi. 2025. [Tulu 3: Pushing frontiers in open language model post-training](#). *Preprint*, arXiv:2411.15124.
- Jia LI, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Costa Huang, Kashif Rasul, Longhui Yu, Albert Jiang, Ziju Shen, Zihan Qin, Bin Dong, Li Zhou, Yann Fleureau, Guillaume Lample, and Stanislas Polu. 2024. [Numina-math](#). [<https://github.com/project-numina/aimo-progress-prize>](https://github.com/project-numina/aimo-progress-prize/blob/main/report/numina_dataset.pdf).
- Bill Yuchen Lin, Abhilasha Ravichander, Ximing Lu, Nouha Dziri, Melanie Sclar, Khyathi Chandu, Chandra Bhagavatula, and Yejin Choi. 2023. [The unlocking spell on base llms: Rethinking alignment via in-context learning](#). *Preprint*, arXiv:2312.01552.
- Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V. Le, Barret Zoph, Jason Wei, and Adam Roberts. 2023. [The flan collection: Designing data and methods for effective instruction tuning](#). *Preprint*, arXiv:2301.13688.
- Meta. 2024a. Llama-3.2-3b. <https://huggingface.co/meta-llama/Llama-3.2-3B>.
- Meta. 2024b. Meet llama 3.1. <https://llama.meta.com>.
- Yingqian Min, Zhipeng Chen, Jinhao Jiang, Jie Chen, Jia Deng, Yiwen Hu, Yiru Tang, Jiapeng Wang, Xiaoxue Cheng, Huatong Song, Wayne Xin Zhao, Zheng Liu, Zhongyuan Wang, and Ji-Rong Wen. 2024. [Imitate, explore, and self-improve: A reproduction report on slow-thinking reasoning systems](#). *Preprint*, arXiv:2412.09413.
- Seyed-Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, Nir Levine, Akihiro Matsukawa, and Hassan Ghasemzadeh. 2019. [Improved knowledge distillation via teacher assistant](#). *Preprint*, arXiv:1902.03393.
- NovaSky. 2025. [Sky-T1: Train your own o1 preview model within \\$450](#). Accessed: 2025-01-09.
- OpenAI. 2023. [Gpt-4 technical report](#).
- OpenAI. 2024. [Learning to reason with llms](#).
- Jiayi Pan, Junjie Zhang, Xingyao Wang, Lifan Yuan, Hao Peng, and Alane Suhr. 2025. [Tinyzero](https://github.com/Jiayi-Pan/TinyZero). <https://github.com/Jiayi-Pan/TinyZero>. Accessed: 2025-01-24.
- Keiran Paster, Marco Dos Santos, Zhangir Azerbayev, and Jimmy Ba. 2023. [Openwebmath: An open dataset of high-quality mathematical web text](#). *Preprint*, arXiv:2310.06786.
- Phuc Phan, Hieu Tran, and Long Phan. 2024. [Distillation contrastive decoding: Improving llms reasoning with contrastive decoding and distillation](#). *arXiv preprint arXiv:2402.14874*.
- Qwen. 2024a. [Qwen2.5: A party of foundation models](#).
- Qwen. 2024b. [Qwq: Reflect deeply on the boundaries of the unknown](#). 2024b.

- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#). *Preprint*, arXiv:2402.03300.
- Avi Singh, John D. Co-Reyes, Rishabh Agarwal, Ankesh Anand, Piyush Patil, Xavier Garcia, Peter J. Liu, James Harrison, Jaehoon Lee, Kelvin Xu, Aaron Parisi, Abhishek Kumar, Alex Alemi, Alex Rizkowsky, Azade Nova, Ben Adlam, Bernd Bohnet, Gamaleldin Elsayed, Hanie Sedghi, Igor Mordatch, Isabelle Simpson, Izzeddin Gur, Jasper Snoek, Jeffrey Pennington, Jiri Hron, Kathleen Keanealy, Kevin Swersky, Kshiteej Mahajan, Laura Culp, Lechao Xiao, Maxwell L. Bileschi, Noah Constant, Roman Novak, Rosanne Liu, Tris Warkentin, Yundi Qian, Yamini Bansal, Ethan Dyer, Behnam Neyshabur, Jascha Sohl-Dickstein, and Noah Fiedel. 2024. [Beyond human data: Scaling self-training for problem-solving with language models](#). *Preprint*, arXiv:2312.06585.
- Yuxuan Tong, Xiwen Zhang, Rui Wang, Ruidong Wu, and Junxian He. 2024. [Dart-math: Difficulty-aware rejection tuning for mathematical problem-solving](#). *Preprint*, arXiv:2407.13690.
- Shubham Toshniwal, Wei Du, Ivan Moshkov, Branislav Kisanin, Alexan Ayrapetyan, and Igor Gitman. 2024. [Openmathinstruct-2: Accelerating ai for math with massive open-source instruction data](#). *Preprint*, arXiv:2410.01560.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. [Llama: Open and efficient foundation language models](#). *Preprint*, arXiv:2302.13971.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Lewis Tunstall, Edward Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro von Werra, Clémentine Fourrier, Nathan Habib, Nathan Sarrazin, Omar Sanseviero, Alexander M. Rush, and Thomas Wolf. 2023. [Zephyr: Direct distillation of lm alignment](#). *Preprint*, arXiv:2310.16944.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-thought prompting elicits reasoning in large language models](#). *Preprint*, arXiv:2201.11903.
- Zhangchen Xu, Fengqing Jiang, Luyao Niu, Yuntian Deng, Radha Poovendran, Yejin Choi, and Bill Yuchen Lin. 2024a. [Magpie: Alignment data synthesis from scratch by prompting aligned llms with nothing](#). *Preprint*, arXiv:2406.08464.
- Zhangchen Xu, Fengqing Jiang, Luyao Niu, Bill Yuchen Lin, and Radha Poovendran. 2024b. [Stronger models are not stronger teachers for instruction tuning](#). *Preprint*, arXiv:2411.07133.
- An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, Keming Lu, Mingfeng Xue, Runji Lin, Tianyu Liu, Xingzhang Ren, and Zhenru Zhang. 2024. [Qwen2.5-math technical report: Toward mathematical expert model via self-improvement](#). *Preprint*, arXiv:2409.12122.
- Edward Yeo, Yuxuan Tong, Morry Niu, Graham Neubig, and Xiang Yue. 2025. [Demystifying long chain-of-thought reasoning in llms](#). *Preprint*, arXiv:2502.03373.
- Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T. Kwok, Zhengguo Li, Adrian Weller, and Weiyang Liu. 2024. [Metamath: Bootstrap your own mathematical questions for large language models](#). *Preprint*, arXiv:2309.12284.
- Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting Dong, Keming Lu, Chuanqi Tan, Chang Zhou, and Jingren Zhou. 2023. [Scaling relationship on learning mathematical reasoning with large language models](#). *Preprint*, arXiv:2308.01825.
- Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhui Chen. 2023. [Mammoth: Building math generalist models through hybrid instruction tuning](#). *Preprint*, arXiv:2309.05653.
- Xiang Yue, Tuney Zheng, Ge Zhang, and Wenhui Chen. 2024. [Mammoth2: Scaling instructions from the web](#). *Preprint*, arXiv:2405.03548.
- Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. 2022. [Star: Bootstrapping reasoning with reasoning](#). *Preprint*, arXiv:2203.14465.
- Weihao Zeng, Yuzhen Huang, Wei Liu, Keqing He, Qian Liu, Zejun Ma, and Junxian He. 2025. [7b model and 8k examples: Emerging reasoning with reinforcement learning is both effective and efficient](#). <https://hkust-nlp.notion.site/simpler1-reason>. Notion Blog.
- Chen Zhang, Yang Yang, Jiahao Liu, Jingang Wang, Yunsen Xian, Benyou Wang, and Dawei Song. 2023. [Lifting the curse of capacity gap in distilling language models](#). *Preprint*, arXiv:2305.12129.
- Dylan Zhang, Qirun Dai, and Hao Peng. 2025. [The best instruction-tuning data are those that fit](#). *Preprint*, arXiv:2502.04194.

Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, Zhangchi Feng, and Yongqiang Ma. 2024. [Llamafactory: Unified efficient fine-tuning of 100+ language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, Bangkok, Thailand. Association for Computational Linguistics.

A Detailed Experimental Setups

Category	Models
Teacher Models	
Long CoT vs ShortCoT	QwQ-32B-Preview vs Qwen2.5-32B-Instruct
Large Teacher vs Small Teacher	
<i>Qwen Family</i>	Qwen2.5-72B-Instruct vs Qwen2.5-3B-Instruct
<i>Llama Family</i>	Llama3.1-70B-Instruct vs Llama3.1-8B-Instruct
<i>Gemma Family</i>	Gemma2-27B-it vs Gemma2-9B-it
Student Models	
<i>Qwen Family</i>	Qwen2.5-0.5B-Instruct, Qwen2.5-1.5B-Instruct, Qwen2.5-3B-Instruct, Qwen2.5-7B-Instruct, Qwen2.5-14B-Instruct, Qwen2.5-32B-Instruct
<i>Llama Family</i>	Llama3.2-1B-Instruct, Llama3.2-3B-Instruct, Llama3.1-8B-Instruct, Llama3.3-70B-Instruct

Table 4: Overview of Teacher and Student Models

A.1 Models

Table 4 presents a comprehensive overview of student and teacher models used in our paper.

A.2 Training Setup

Our model training is conducted using LLaMA-Factory (Zheng et al., 2024), on a server with four NVIDIA A100-SXM4-80GB GPUs, an AMD EPYC 7763 64-Core Processor, and 512 GB of RAM. We use full parameter fine-tuning on student models less than 14B parameters. When the student model is larger than 14B, we use LoRA fine-tuning (Hu et al., 2021). Table 5 and Table 6 list hyper-parameters for full parameter fine-tuning and LoRA fine-tuning respectively.

Hyper-parameter	Value
Learning Rate	1×10^{-5}
Number of Epochs	2
Number of Devices	4
Per-device Batch Size	2
Optimizer	Adamw
Learning Rate Scheduler	cosine
Max Sequence Length	16384

Table 5: This table shows the hyper-parameters for full parameter fine-tuning.

Hyper-parameter	Value
Learning Rate	1×10^{-4}
Number of Epochs	2
Number of Devices	4
Per-device Batch Size	1
Lora Target	full
Learning Rate Scheduler	cosine
Warmup Ratio	0.03
Max Sequence Length	16384

Table 6: This table shows the hyper-parameters for LoRA fine-tuning.

Teacher models generate responses by rejection sampling (Zelikman et al., 2022; Tong et al., 2024; Yue et al., 2023; Singh et al., 2024; Gulcehre et al., 2023; Yuan et al., 2023; Dong et al., 2023). By default, teacher models employ greedy decoding. By combining the math problem instructions with corresponding solutions generated by teacher models, we construct problem-solution pairs to fine-tune student models. We perform pairwise comparisons of solutions generated by different teacher models and filter out problem-solution pairs that are correct for both models to fine-tune student models.

A.3 Evaluation Setup

We evaluate the reasoning capability of fine-tuned student models on a set of commonly used benchmarks, including MATH (Hendrycks et al., 2021), GSM8K (Cobbe et al., 2021), AMC 2023, AIME 2024, and the English math subset of Olympiad-Bench (He et al., 2024).

Unless otherwise specified, all fine-tuned models are evaluated in a zero-shot setting using greedy decoding. We set the maximum generation tokens as 16k. The evaluation prompt is shown below.

Prompt
Solve the following math problem and present the final answer in the format: Final Answer: {your answer} Problem: {problem} Answer:

After extracting the final answer of the evaluated model, we first employ exact matching to determine the correctness of the answer. If the answer is incorrect, we use Qwen-32B-Instruct as a judge to compare the extracted final answers against that of the ground truth. The prompt is shown below.

Model	MATH			GSM8k			AIME			AMC			Olympiad			Average
	P_{Strong}	P_{Weak}	Δ_{Strong}	P_{Strong}	P_{Weak}	Δ_{Strong}	P_{Strong}	P_{Weak}	Δ_{Strong}	P_{Strong}	P_{Weak}	Δ_{Strong}	P_{Strong}	P_{Weak}	Δ_{Strong}	Δ_{Strong}
Llama-3.2-1B	29.8	29.6	0.160	44.4	47.5	-3.18	0.00	0.00	0.00	2.50	7.50	-5.00	6.07	7.70	-1.63	-1.93
Llama-3.2-3B	47.4	47.9	-0.500	71.2	74.1	-2.88	3.33	0.00	3.33	25.0	17.5	7.50	16.9	16.4	0.445	1.58
Llama-3.2-8B	37.6	37.6	-0.040	67.0	69.2	-2.20	6.67	0.00	6.67	7.50	7.50	0.00	9.19	11.0	-1.78	0.530
Llama-3.2-70B	74.5	72.2	2.28	92.0	92.2	-0.152	16.7	16.7	0.00	67.5	50.0	17.5	37.3	35.7	1.63	4.25
Qwen2.5-0.5B	30.0	31.0	-0.920	43.1	45.4	-2.35	0.00	0.00	0.00	5.00	17.5	-12.5	6.52	8.30	-1.78	-3.51
Qwen2.5-1.5B	50.3	50.7	-0.440	70.6	71.0	-0.455	0.00	3.33	-3.33	22.5	20.0	2.50	17.8	20.0	-2.22	-0.790
Qwen2.5-3B	57.5	60.3	-2.82	79.9	79.5	0.379	0.00	3.33	-3.33	35.0	27.5	7.50	25.9	26.4	-0.444	0.256
Qwen2.5-7B	71.3	63.6	7.66	87.8	84.1	3.72	6.67	0.00	6.67	40.0	35.0	5.00	38.8	29.0	9.78	6.56
Qwen2.5-14B	76.4	72.8	3.66	93.1	89.6	3.49	6.67	3.33	3.33	47.5	45.0	2.50	41.0	39.0	2.07	3.01
Qwen2.5-32B	80.5	76.8	3.72	92.2	92.7	-0.531	20.0	3.33	16.7	57.5	50.0	7.50	47.4	42.4	5.04	6.48

Table 7: This table summarizes the performance of models in Llama and Qwen families fine-tuned with large teacher CoT and small teacher CoT when evaluated on MATH, GSM8K, AIME, AMC, and OlympiadBench. Qwen-2.5-72B-Instruct is chosen as the large teacher while Qwen-2.5-3B-Instruct is chosen as the small teacher. We observe that small student models may experience degraded performance when distilled from a large teacher compared to a small teacher, whereas larger student models benefit more from the distilling a large teacher.

Prompt

Given a math problem, its correct final answer, and the model’s generated final answer, determine if the model’s answer is correct. Respond with ‘True’ if the it is correct and ‘False’ if it is incorrect.

Problem: {problem}

Correct Final Answer: {ground truth}

Model’s Generated Final Answer: {resp answer}

Your Judgement:

B More Experiments Results

In this section we present additional experiment results of long CoT gap and large teacher CoT gap.

B.1 Long CoT Gap: Additional Results

Table 8 shows the detailed performance scores and gap of each benchmark for different student models fine-tuned on long CoT and short CoT. QwQ-32B-Preview is chosen to generate long CoT and awhile Qwen-2.5-32B-Instruct is chosen to generate short CoT. We observe that small student models tend to benefit more from short CoT, while large student models gain greater advantages from long CoT.

B.2 Large Teacher CoT Gap: Additional Results

Table 7 shows the detailed performance scores and gap of each benchmark for different student models distilled from large teacher and small teacher. We summarize the performance of 10 student models from the Llama and Qwen families across various

model sizes. Qwen-2.5-72B-Instruct is chosen as the large teacher while Qwen-2.5-3B-Instruct is chosen as the small teacher. The results are shown in Table 7. Our findings indicate that small student models may experience degraded performance when distilled from a large teacher compared to a small teacher, whereas larger student models benefit more from distilling a large teacher.

Table 9 shows more experiment results for teacher models in different model families, including Gemma-27B-it vs Gemma-9B-it and Llama3.1-72B-Instruct vs Llama3.1-8B-Instruct.

B.3 Empirical Evidence for Distribution Gap Between Student and Teacher Models

We suggest that the distribution gap between student models and teacher models may be a key factor leading to the small model learnability gap. We provide empirical evidence through perplexity measurements and text similarity comparisons.

Training Data Perplexity Analysis. To quantify the distribution gap between student and teacher models, we measured perplexity (PPL) of teacher-generated training data on different student models in Table 10. Lower PPL indicates better alignment between the student model’s distribution and the training data distribution. Our findings reveal several key patterns:

1. Small students struggle with complex sequences: Small student models assign significantly higher PPL to large teacher CoT or long CoT sequences, indicating difficulty in modeling such complex reasoning traces.
2. Aligned teacher-student pairs show better dis-

Model	MATH			GSM8K			AIME			AMC			Olympiad			Average Δ_{Long}
	P_{Long}	P_{Short}	Δ_{Long}	P_{Long}	P_{Short}	Δ_{Long}	P_{Long}	P_{Short}	Δ_{Long}	P_{Long}	P_{Short}	Δ_{Long}	P_{Long}	P_{Short}	Δ_{Long}	
Llama-3.2-1B	28.6	33.4	-4.78	42.3	49.2	-6.90	0.00	0.00	0.00	2.50	7.50	-5.00	5.48	7.40	-1.92	-3.72
Llama-3.2-3B	48.7	50.9	-2.14	75.1	77.5	-2.42	3.33	3.33	0.00	17.5	15.0	2.50	17.6	18.7	-1.04	-0.619
Llama-3.1-8B	50.0	44.6	5.36	81.4	75.5	5.84	0.00	0.00	0.00	27.5	22.5	5.00	17.3	14.8	2.52	3.74
Llama-3.3-70B	75.3	74.9	0.340	92.7	91.2	1.44	26.7	13.3	13.3	55.0	52.5	2.50	41.3	39.7	1.63	3.85
Qwen2.5-0.5B	23.0	31.5	-8.44	39.5	45.3	-5.84	0.00	0.00	0.00	7.50	15.0	-7.50	4.00	5.93	-1.93	-4.74
Qwen2.5-1.5B	41.6	52.3	-10.7	63.8	71.7	-7.89	0.00	0.00	0.00	17.5	27.5	-10.0	12.3	19.4	-7.11	-7.13
Qwen2.5-3B	56.2	61.0	-4.84	80.0	82.0	-1.98	3.33	10.0	-6.67	37.5	37.5	0.00	24.4	26.4	-1.93	-3.08
Qwen2.5-7B	68.2	67.8	0.460	86.2	85.7	0.560	13.3	6.67	6.67	40.0	40.0	0.00	36.6	35.7	0.889	1.72
Qwen2.5-14B	78.3	76.2	2.04	93.3	92.5	0.760	20.0	6.67	13.3	60.0	55.0	5.00	44.4	40.9	3.56	4.94
Qwen2.5-32B	84.8	82.3	2.44	94.9	94.3	0.610	40.0	10.0	30.0	85.0	62.5	22.5	60.4	47.3	13.2	13.7

Table 8: This table summarizes the performance of models in Llama and Qwen families fine-tuned with long CoT and short CoT data. They are evaluated on MATH, GSM8K, AIME, AMC, and OlympiadBench. QwQ-32B-Preview is chosen to generate long CoT and awhile Qwen-2.5-32B-Instruct is chosen to generate short CoT. We observe that small student models tend to benefit more from short CoT, while large student models gain greater advantages from long CoT.

Model	Gemma2-9B vs Gemma2-27B						Llama3.1-8B vs Llama3.1-70B					
	MATH	AMC	Olympiad	AIME	GSM8k	Average	MATH	AMC	Olympiad	AIME	GSM8k	Average
Llama3.2-1B	-1.42	-7.50	0.00	0.00	-0.227	-1.83	-1.42	-5.00	-0.296	3.33	0.152	-0.646
Llama3.2-3B	2.08	-7.50	-0.888	0.00	1.67	-0.928	-0.14	10.0	-0.593	3.33	1.06	2.73
Llama3.1-8B	0.56	0.00	0.078	0.00	-0.516	0.0243	-2.18	7.50	2.67	0.00	-1.29	1.34
Llama3.1-70B	0.02	7.50	-0.741	10.0	0.152	3.39	2.72	17.5	5.48	6.67	0.986	6.67
Qwen2.5-0.5B	-4.56	0.00	0.741	0.00	0.592	-0.645	-1.88	0.00	0.185	0.00	-1.74	-0.688
Qwen2.5-1.5B	-1.20	2.50	-1.19	0.00	-0.986	-0.174	-1.48	5.00	-0.148	3.33	-1.14	1.11
Qwen2.5-3B	0.44	5.00	1.78	0.00	-0.758	1.29	-1.26	5.00	-0.741	-3.33	-1.29	-0.325
Qwen2.5-7B	0.22	5.00	1.04	-3.33	3.94	1.37	3.68	20.0	4.15	3.33	2.81	6.79
Qwen2.5-14B	1.32	2.50	-0.148	0.00	-0.986	0.537	2.18	0.00	0.445	3.33	-0.303	1.13
Qwen2.5-32B	0.10	2.50	1.48	3.44	1.36	1.78	2.72	-2.50	5.63	3.33	0.834	2.00

Table 9: This table presents the performance of student models distilled from different teacher models, including Gemma-27B-it vs Gemma-9B-it and Llama3.1-72B-Instruct vs Llama3.1-8B-Instruct. We observe that small student models may experience degraded performance when distilled from a large teacher compared to a small teacher, whereas larger student models benefit more from the distilling a large teacher.

tribution matching. Small teacher CoT yields lower PPL in small students, suggesting reduced distribution gap when teacher and student capacities are more aligned.

3. As student model size increases, the PPL gap between long and short CoT (and between large and small teacher CoT) shrinks, indicating that larger students can more easily adapt to complex reasoning distributions.

Text Similarity Analysis. We conducted additional analyses comparing responses generated by student models with those from small and large teachers using two text similarity metrics:

- **TF-IDF cosine similarity:** Measures lexical similarity between texts
- **Embedding similarity:** Uses all-mpnet-base-v2 to capture semantic similarity

The text similarity analysis provides interpretable evidence of the distribution gap in Table 11. We found that student responses are consistently more similar to small teacher CoT than to large teacher CoT across both lexical and semantic similarity metrics. The confidence intervals are tight and do not overlap, indicating that the observed differences are statistically significant and not due to outlier effects.

C Examples of Speaking Style Shift

We adopt the method from (Lin et al., 2023) to evaluate the most shifted tokens after fine-tuning on long CoT and Large teacher CoT data. Figure 8 shows the calculation process. This allows us to compare the token distribution shifts induced by the fine-tuning process. We annotate the tokens that exhibit the largest rank shifts as the most shifted tokens. We choose Qwen2.5-3B-Instruct as the student model. We put the results of most shifted tokens after fine-tuning on long CoT data in Figure 9 and 10. The results of most shifted tokens after fine-tuning on large teacher CoT data are shown in Figure 11. Our analysis reveals that these tokens are predominantly associated with expressive and stylistic elements, such as “wait”, “But”, and “Let”.

D Examples of Various CoT Data

This section demonstrates examples of different CoT data including long CoT, short CoT, large teacher CoT and small teacher CoT. See the following examples.

Table 10: Perplexity Analysis Results

Student Model	Long CoT	Short CoT	Δ (L-S)	Large Teacher	Small Teacher	Δ (Lg-Sm)
Qwen-0.5B	2.237	1.278	0.959	1.246	1.217	0.028
Qwen-1.5B	2.969	1.226	0.743	1.204	1.178	0.026
Qwen-3B	1.963	1.246	0.716	1.225	1.155	0.069
Qwen-7B	1.923	1.222	0.700	1.197	1.180	0.016
Qwen-14B	1.902	1.218	0.683	1.198	1.189	0.009
Qwen-32B	1.265	1.050	0.215	1.053	1.051	0.002

Table 11: Text Similarity Analysis Results

Student Model	Metric	Small Teacher	Large Teacher
Qwen2.5-1.5B	TF-IDF Similarity	0.8329 \pm 0.004	0.8235 \pm 0.004
	Embedding Similarity	0.9461 \pm 0.002	0.9413 \pm 0.002
Qwen2.5-0.5B	TF-IDF Similarity	0.7928 \pm 0.003	0.7854 \pm 0.003
	Embedding Similarity	0.9372 \pm 0.001	0.9297 \pm 0.002

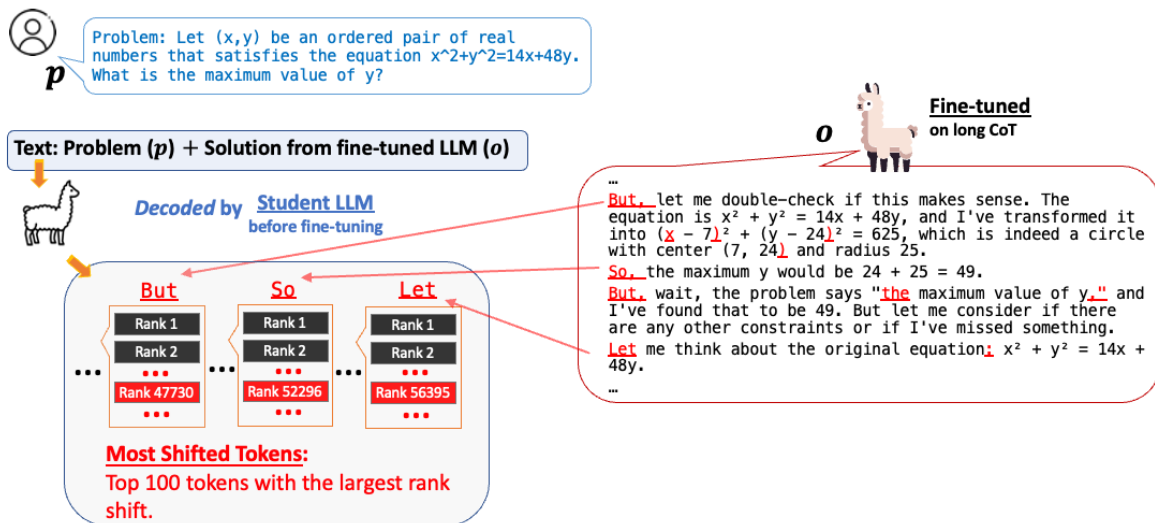


Figure 8: The process of calculating most shifted tokens. We decode each token generated by the fine-tuned LLM in the student model before fine-tuning. Then we calculate the rank shift in the student model for each token generated by the fine-tuned model. We annotate the tokens that exhibit the largest rank shifts as the most shifted tokens. We found that these tokens are predominantly associated with expressive and stylistic elements, such as “But” and “Let”.

Problem

Let (x, y) be an ordered pair of real numbers that satisfies the equation $x^2 + y^2 = 14x + 48y$. What is the maximum value of y ?

Example of **most shifted tokens** after fine-tuning on long CoT data. (1/2)

So I've got this math problem here. It says: "Let (x, y) be an ordered pair of real numbers that satisfies the equation $x^2 + y^2 = 14x + 48y$. What is the maximum value of y ? Okay, let's see how to approach this.

First, I need to understand what the equation is telling me. It's an equation involving x and y , and it's set equal to zero:

$$x^2 + y^2 - 14x - 48y = 0.$$

I need to find the maximum value of y for real numbers x and y that satisfy this equation.

Hmm, maybe I can **rearrange** this equation to make it easier to work with. Let's see:

$$x^2 - 14x + y^2 - 48y = 0.$$

I can complete the square for both x and y terms.

For x : $x^2 - 14x$, to complete the square, I take half of -14 , which is -7 , and square it to get 49 . So, $x^2 - 14x$ becomes $(x - 7)^2 - 49$.

For y : $y^2 - 48y$, half of -48 is -24 , squared is 576 . So, $y^2 - 48y$ becomes $(y - 24)^2 - 576$.

Now, **plugging** these back into the equation:

$$(x - 7)^2 - 49 + (y - 24)^2 - 576 = 0.$$

Simplify:

$$(x - 7)^2 + (y - 24)^2 - 625 = 0,$$

$$\text{So, } (x - 7)^2 + (y - 24)^2 = 625.$$

This looks like the equation of a circle with center at $(7, 24)$ and radius 25 .

Wait, but the problem is asking for the maximum value of y , and I have a circle equation. **So**, the maximum y would be the y -coordinate of the highest point on the circle, which is the center plus the radius.

Wait, but the center is at $(7, 24)$, and the radius is 25 , so the highest point would be at $(7, 24 + 25) = (7, 49)$.

But, let me double-check if this makes sense. The equation is

$$x^2 + y^2 = 14x + 48y,$$

and I've transformed it into

$$(x - 7)^2 + (y - 24)^2 = 625,$$

which is indeed a circle with center $(7, 24)$ and radius 25 .

So, the maximum y would be $24 + 25 = 49$.

But, wait, the problem says "the maximum value of y ," and I've found that to be 49 . But let me consider if there are any other constraints or if I've missed something.

Let me think about the original equation:

$$x^2 + y^2 = 14x + 48y.$$

I can **rearrange** it to:

$$x^2 - 14x + y^2 - 48y = 0.$$

Figure 9: Example of **most shifted tokens** after fine-tuning on long CoT data. All the most shifted tokens are labeled as red.

Example of **most shifted tokens** after fine-tuning on long CoT data (2/2)

Then, completing the square:

$$(x - 7)^2 - 49 + (y - 24)^2 - 576 = 0,$$

$$(x - 7)^2 + (y - 24)^2 = 625.$$

Yes, that's a circle with center $(7, 24)$ and radius 25.

So, the maximum y is indeed $24 + 25 = 49$.

But, just to be thorough, let's consider if there are any points on the circle where y is greater than 49. The center is at $(7, 24)$, and the radius is 25, so the highest point is at $(7, 49)$, and the lowest point would be at $(7, -1)$.

But, the problem is asking for the maximum value of y , and I've found that to be 49. Therefore, the final answer is 49.

Final Answer:

49

Figure 10: **Most shifted tokens** after fine-tuning on long CoT data (2/2). All the most shifted tokens are labeled as red.

Example of **Most Shifted Tokens** after fine-tuning on large model CoT

To solve the problem, we start with the given equation:

$$x^2 + y^2 = 14x + 48y$$

First, we rearrange the equation to bring all terms to one side:

$$x^2 - 14x + y^2 - 48y = 0$$

Next, we complete the square for both x and y .

For x :

$$x^2 - 14x$$

To complete the square, we add and subtract $(\frac{14}{2})^2 = 49$:

$$x^2 - 14x + 49 - 49$$

$$(x - 7)^2 - 49$$

For y :

$$y^2 - 48y$$

To complete the square, we add and subtract $(\frac{48}{2})^2 = 576$:

$$y^2 - 48y + 576 - 576$$

$$(y - 24)^2 - 576$$

Now, substitute these completed squares back into the equation:

$$(x - 7)^2 - 49 + (y - 24)^2 - 576 = 0$$

$$(x - 7)^2 + (y - 24)^2 - 625 = 0$$

$$(x - 7)^2 + (y - 24)^2 = 625$$

This is the equation of a circle with center $(7, 24)$ and radius 25.

To find the maximum value of y , we need to find the highest point on the circle. The highest point on the circle is directly above the center, which is at $y = 24 + 25 = 49$.

Thus, the maximum value of y is:

$$y = 49$$

Final Answer:

Figure 11: Example of **most shifted tokens** after fine-tuning on large teacher CoT. All the most shifted tokens are labeled as red.

Problem

What is the sum of the positive odd divisors of 60?

Long CoT Example

To find the sum of the positive odd divisors of 60, I first need to understand what the divisors of 60 are and then identify which of those are odd.

First, let's find all the positive divisors of 60. To do this, I can start by factoring 60 into its prime factors.

$$60 = 2^2 \times 3^1 \times 5^1$$

Now, to find all the positive divisors, I can take all possible products of these prime factors raised to their possible exponents. For 2, the exponents can be 0, 1, or 2; for 3, 0 or 1; and for 5, 0 or 1. Listing them out:

$$\begin{aligned}1 &= 2^0 \times 3^0 \times 5^0, \\2 &= 2^1 \times 3^0 \times 5^0, \\4 &= 2^2 \times 3^0 \times 5^0, \\3 &= 2^0 \times 3^1 \times 5^0, \\6 &= 2^1 \times 3^1 \times 5^0, \\12 &= 2^2 \times 3^1 \times 5^0, \\5 &= 2^0 \times 3^0 \times 5^1, \\10 &= 2^1 \times 3^0 \times 5^1, \\20 &= 2^2 \times 3^0 \times 5^1, \\15 &= 2^0 \times 3^1 \times 5^1, \\30 &= 2^1 \times 3^1 \times 5^1, \\60 &= 2^2 \times 3^1 \times 5^1.\end{aligned}$$

So, the positive divisors of 60 are:

$$1, 2, 3, 4, 5, 6, 10, 12, 15, 20, 30, 60.$$

Now, I need to identify which of these are odd. Odd numbers are not divisible by 2, so I can eliminate any divisors that are even.

The odd divisors are:

$$1, 3, 5, 15.$$

Now, I need to find the sum of these odd divisors: $1 + 3 + 5 + 15 = 24$. Alternatively, using the prime factorization approach: Since the odd divisors are those that do not include the factor 2, I can ignore the 2's in the prime factorization. The remaining odd part is:

$$3^1 \times 5^1.$$

The sum of the divisors of this odd part is:

$$(1 + 3)(1 + 5) = 4 \times 6 = 24.$$

This matches the previous result.

Therefore, for $60 = 2^2 \times 3^1 \times 5^1$, the sum of the odd divisors is:

$$\boxed{24}.$$

Figure 12: Long CoT Example

Short CoT Example

To find the sum of the positive odd divisors of 60, we first need to factorize 60 into its prime factors. The prime factorization of 60 is:

$$60 = 2^2 \times 3 \times 5$$

Since we are only interested in the odd divisors, we can ignore the factor of 2. Therefore, we only consider the factors 3 and 5. The odd divisors of 60 are the divisors of $3 \times 5 = 15$. The divisors of 15 are 1, 3, 5, and 15.

Next, we sum these divisors:

$$1 + 3 + 5 + 15 = 24$$

Therefore, the sum of the positive odd divisors of 60 is:

$$\boxed{24}$$

Figure 13: Short CoT example

Problem

Find the greatest a such that

$$\frac{7\sqrt{(2a)^2 + (1)^2} - 4a^2 - 1}{\sqrt{1 + 4a^2} + 3} = 2.$$

Large Teacher CoT Example (1/2)

To solve the equation $\frac{7\sqrt{(2a)^2+1^2-4a^2-1}}{\sqrt{1+4a^2+3}} = 2$, we start by simplifying the expression inside the square root in the numerator. The equation becomes:

$$\frac{7\sqrt{4a^2+1}-4a^2-1}{\sqrt{1+4a^2+3}} = 2$$

Let's denote $x = \sqrt{1+4a^2}$. Then the equation transforms into:

$$\frac{7x-4a^2-1}{x+3} = 2$$

Next, we clear the fraction by multiplying both sides by $x+3$:

$$7x-4a^2-1 = 2(x+3)$$

Expanding and simplifying the right side, we get:

$$7x-4a^2-1 = 2x+6$$

Rearranging terms to isolate x and a^2 , we have:

$$7x-2x-4a^2-1-6 = 0 \implies 5x-4a^2-7 = 0 \implies 5x = 4a^2+7 \implies x = \frac{4a^2+7}{5}$$

Recall that $x = \sqrt{1+4a^2}$. Therefore, we have:

$$\sqrt{1+4a^2} = \frac{4a^2+7}{5}$$

To eliminate the square root, we square both sides:

$$1+4a^2 = \left(\frac{4a^2+7}{5}\right)^2$$

Expanding the right side, we get:

$$1+4a^2 = \frac{(4a^2+7)^2}{25} = \frac{16a^4+56a^2+49}{25}$$

Multiplying both sides by 25 to clear the denominator, we obtain:

$$25+100a^2 = 16a^4+56a^2+49$$

Large Teacher CoT Example (2/2)

Rearranging all terms to one side, we get a standard polynomial equation:

$$16a^4 + 56a^2 + 49 - 100a^2 - 25 = 0 \implies 16a^4 - 44a^2 + 24 = 0$$

Let $y = a^2$. Then the equation becomes a quadratic in y :

$$16y^2 - 44y + 24 = 0$$

We solve this quadratic equation using the quadratic formula $y = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$, where $a = 16$, $b = -44$, and $c = 24$:

$$y = \frac{44 \pm \sqrt{(-44)^2 - 4 \cdot 16 \cdot 24}}{2 \cdot 16} = \frac{44 \pm \sqrt{1936 - 1536}}{32} = \frac{44 \pm \sqrt{400}}{32} = \frac{44 \pm 20}{32}$$

This gives us two solutions for y :

$$y = \frac{64}{32} = 2 \quad \text{and} \quad y = \frac{24}{32} = \frac{3}{4}$$

Since $y = a^2$, we have $a^2 = 2$ or $a^2 = \frac{3}{4}$. The corresponding values for a are $a = \pm\sqrt{2}$ and $a = \pm\frac{\sqrt{3}}{2}$. The greatest value of a is $\sqrt{2}$.

Thus, the greatest a that satisfies the equation is:

$$\boxed{\sqrt{2}}$$

Small Teacher CoT Example (1/2)

To solve the equation $\frac{7\sqrt{(2a)^2+(1)^2-4a^2-1}}{\sqrt{1+4a^2}+3} = 2$, we start by simplifying the expression inside the square root and the equation itself.

First, simplify the expression under the square root in the numerator:

$$(2a)^2 + 1^2 = 4a^2 + 1$$

So the equation becomes:

$$\frac{7\sqrt{4a^2 + 1} - 4a^2 - 1}{\sqrt{4a^2 + 1} + 3} = 2$$

Next, let $x = \sqrt{4a^2 + 1}$. Then the equation transforms into:

$$\frac{7x - 4a^2 - 1}{x + 3} = 2$$

Since $x = \sqrt{4a^2 + 1}$, we can express $4a^2$ in terms of x :

$$4a^2 = x^2 - 1$$

Substitute $4a^2 = x^2 - 1$ into the equation:

$$\frac{7x - (x^2 - 1) - 1}{x + 3} = 2$$

Simplify the numerator:

$$7x - x^2 + 1 - 1 = 7x - x^2$$

So the equation becomes:

$$\frac{7x - x^2}{x + 3} = 2$$

Multiply both sides by $x + 3$ to clear the denominator:

$$7x - x^2 = 2(x + 3)$$

Expand and simplify:

$$7x - x^2 = 2x + 6$$

Rearrange all terms to one side:

$$-x^2 + 7x - 2x - 6 = 0$$

$$-x^2 + 5x - 6 = 0$$

Multiply through by -1 to make the leading coefficient positive:

$$x^2 - 5x + 6 = 0$$

Small Teacher CoT Example (2/2)

Factor the quadratic equation:

$$(x - 2)(x - 3) = 0$$

So, the solutions are:

$$x = 2 \quad \text{or} \quad x = 3$$

Recall that $x = \sqrt{4a^2 + 1}$. We need to check both solutions:

1. For $x = 2$:

$$\sqrt{4a^2 + 1} = 2$$

Square both sides:

$$4a^2 + 1 = 4$$

$$4a^2 = 3$$

$$a^2 = \frac{3}{4}$$

$$a = \pm \frac{\sqrt{3}}{2}$$

2. For $x = 3$:

$$\sqrt{4a^2 + 1} = 3$$

Square both sides:

$$4a^2 + 1 = 9$$

$$4a^2 = 8$$

$$a^2 = 2$$

$$a = \pm\sqrt{2}$$

The greatest value of a among the solutions is $\sqrt{2}$.

Thus, the final answer is:

$$\boxed{\sqrt{2}}$$

Problem

What is the greatest common divisor of 128, 144 and 480?

Large Teacher CoT Example

To find the greatest common divisor (GCD) of 128, 144, and 480, we can use the prime factorization method.

1. Prime Factorization:

$$128 = 2^7,$$

$$144 = 2^4 \times 3^2,$$

$$480 = 2^5 \times 3 \times 5.$$

2. Identify the common prime factors:

The common prime factor among 128, 144, and 480 is 2.

3. Determine the lowest power of the common prime factor:

The lowest power of 2 in the factorizations is 2^4 .

4. Calculate the GCD:

The GCD is:

$$2^4 = 16.$$

Therefore, the greatest common divisor of 128, 144, and 480 is:

$$\boxed{16}$$

Small Teacher CoT Example

To find the greatest common divisor (GCD) of 128, 144, and 480, we start by finding the prime factorization of each number.

The prime factorization of 128 is:

$$128 \div 2 = 64$$

$$64 \div 2 = 32$$

$$32 \div 2 = 16$$

$$16 \div 2 = 8$$

$$8 \div 2 = 4$$

$$4 \div 2 = 2$$

$$2 \div 2 = 1$$

Thus,

$$128 = 2^7.$$

The prime factorization of 144 is:

$$144 \div 2 = 72$$

$$72 \div 2 = 36$$

$$36 \div 2 = 18$$

$$18 \div 2 = 9$$

$$9 \div 3 = 3$$

$$3 \div 3 = 1$$

Thus,

$$144 = 2^4 \times 3^2.$$

The prime factorization of 480 is:

$$480 \div 2 = 240$$

$$240 \div 2 = 120$$

$$120 \div 2 = 60$$

$$60 \div 2 = 30$$

$$30 \div 2 = 15$$

$$15 \div 3 = 5$$

$$5 \div 5 = 1$$

Thus,

$$480 = 2^5 \times 3 \times 5.$$

Now, we identify the common prime factors. The factorizations are:

$$128 = 2^7, \quad 144 = 2^4 \times 3^2, \quad 480 = 2^5 \times 3 \times 5.$$

The only common prime factor among all three numbers is 2.

Next, we determine the smallest power of the common prime factor. The lowest power of 2 among these factorizations is 2^4 .

Therefore, the GCD of 128, 144, and 480 is:

Problem

Rounded to 2 decimal places, what is $\frac{7}{9}$?

Large Teacher CoT Example

To solve the problem, we need to divide 7 by 9 and then round the result to two decimal places.

$$\frac{7}{9} = 0.777\dots$$

Rounded to two decimal places, this becomes 0.78.

Final Answer:

Small Teacher CoT Example

To solve the problem of finding the value of $\frac{7}{9}$ rounded to 2 decimal places, we can follow these steps:

1. **Perform the division:** Calculate $\frac{7}{9}$.
2. **Convert to decimal:** Perform the division to get the decimal equivalent.
3. **Round to 2 decimal places:** Round the decimal to 2 decimal places.

Let's start with the division:

$$\frac{7}{9} \approx 0.7777\dots$$

Next, we need to round this decimal to 2 decimal places. The digit in the third decimal place is 7, which is greater than or equal to 5. Therefore, we round the second decimal place up from 7 to 8.

So, $0.7777\dots$ rounded to 2 decimal places is 0.78.

The final answer is: