

Weak-to-Strong Reasoning

Yuqing Yang^{2,4} Yan Ma^{2,3,4} Pengfei Liu^{1,3,4*}

¹Shanghai Jiao Tong University ²Fudan University

³Shanghai AI Laboratory ⁴Generative AI Research Lab (GAIR)

{yuqingyang21, yanma23}@m.fudan.edu.cn pengfei@sjtu.edu.cn

Abstract

When large language models (LLMs) surpass human capabilities, supervising them effectively becomes difficult. *Weak-to-strong learning*, where a less capable model enhances a stronger one, proves valuable in this context. Yet, the efficacy of this paradigm for complex reasoning tasks is still unexplored. In this paper, we introduce a progressive weak-to-strong reasoning framework that enables the strong model to autonomously refine its training data, maximizing the use of weak signals and unlocking its latent abilities. This framework begins with supervised fine-tuning on a selective small but high-quality dataset, followed by preference optimization on contrastive samples identified by the strong model itself. Experiments on the GSM8K and MATH datasets verify that our method can effectively improve the reasoning capabilities of Llama2-70b using three separate weak models. This work paves the way for a more scalable and sophisticated strategy to enhance AI reasoning powers. All relevant code and resources are available in <https://github.com/GAIR-NLP/weak-to-strong-reasoning>.

1 Introduction

As the pursuit of Artificial General Intelligence (AGI) advances, the creation of superintelligent systems—models that exceed human cognitive capabilities—remains a key ambition within the field (Robert, 2017; Altman et al., 2023; Puthumanailam et al., 2024). This quest introduces a host of challenges, especially concerning the *supervision* and *learning paradigms* for these advanced AI models. Conventional supervision methods, which typically depend on human oversight (Christiano et al., 2017; Ouyang et al., 2022; Sun et al., 2024) or guidance (i.e., distilled knowledge) from more advanced models (Bai et al., 2022; Lee et al., 2023;

*Corresponding author.

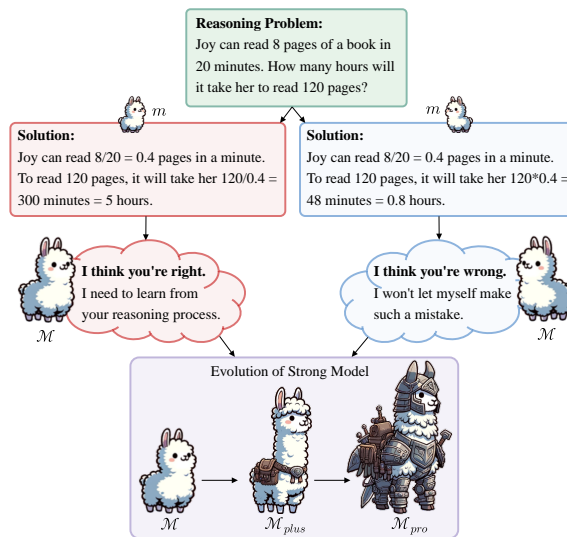


Figure 1: Illustration of weak-to-strong reasoning through the strong model self-refining its training data.

Peng et al., 2023), become inadequate as the capabilities of AI exceed those of their supervisors (Bowman et al., 2022; Sang et al., 2024). To address this issue, we focus on the *weak-to-strong learning* paradigm (Burns et al., 2023), which operates under a unique task setting where only a less capable model and a stronger¹ but not fully utilized model are available.

The central question of weak-to-strong learning is whether models with limited capabilities can effectively guide the development of more advanced, stronger models. Previous studies by Burns et al. (2023) have demonstrated the feasibility of it in classification, chess, and reward modeling tasks. However, the applicability of this setup to more complex reasoning tasks, which demand more than

¹Similar to Burns et al. (2023), we define “strong model” in the context of LLMs, taking into account their characteristics—that is, LLMs often contain the knowledge and capabilities needed to perform specific tasks, but these have not yet been fully elicited (Zhou et al., 2024). Typically, it refers to stronger and larger pre-trained language models whose capabilities have not been fully realized yet.

mere extrapolation or pattern recognition, remains an open question. Complex reasoning represents a key aspect of human cognition, crucial for assessing whether LLMs can emulate or surpass human-like capabilities in comprehending the world, making decisions, and solving problems (Qiao et al., 2023; Huang and Chang, 2023; Chang et al., 2023). Given the complexity and the critical nature of these tasks, applying the weak-to-strong learning framework to advanced reasoning challenges is essential, particularly within the broader context of achieving superintelligence.

Although Burns et al. (2023) suggest that naively fine-tuning strong models on the full set of noisy data produced by weak models, named *full weak fine-tuning*, can consistently improve their performance over the weaker counterparts, this approach is still far from recovering the full capabilities of strong models, and our experiments show that it loses effectiveness when facing more complex reasoning challenges. They also propose an auxiliary confidence loss to mitigate the issue of strong models imitating the errors of their supervisors. However, this method is tailored to classification tasks with a set of fixed labels and does not naturally extend to open-ended generation tasks including reasoning. Currently, there is a lack of effective methods beyond naive fine-tuning to prevent the overfit of weak errors and to further elicit the intrinsic reasoning abilities of strong models within the weak-to-strong reasoning framework.

To achieve the above goal, we introduce a progressive refinement learning framework, guided by the principle that a model can enhance its capabilities more effectively by initially focusing on smaller, more reliable subsets of data, and then iteratively expanding its learning scope, as illustrated in Fig. 1. In the first stage, we hypothesize that it is more advantageous to utilize smaller quantities of data that are likely to be more accurate. We achieve this by combining weak data, generated by the less capable model, with data self-generated by the more advanced model through in-context learning. This blend is then used to selectively curate datasets for subsequent supervised fine-tuning. In the second stage, upon having developed a strong model with improved reasoning capabilities, we utilize its ability to construct contrastive samples for preference optimization (Rafailov et al., 2023; Hong et al., 2024) and enable the model to learn effectively from the errors of the weaker model.

In implementation, we employ Llama2-70b

(Touvron et al., 2023) as the strong model, test three separate weak models: Llama2-7b, Gemma-2b (Mesnard et al., 2024), and Mistral-7b (Jiang et al., 2023), and conduct experiments on the commonly used math reasoning datasets GSM8K (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021). Experimental results reveal that:

(1) Full weak fine-tuning, while effective in classification tasks, falls short for complex reasoning tasks. (2) Our proposed method significantly outperforms full weak fine-tuning method, achieving a **26.99-point** improvement on GSM8K when supervised solely by the weak model (i.e., Gemma-2b) after the first stage of training ($\mathcal{M} \rightarrow \mathcal{M}_{\text{plus}}$), and further enhances performance by an additional **8.49** points through preference optimization without knowing the gold answer ($\mathcal{M}_{\text{plus}} \rightarrow \mathcal{M}_{\text{pro}}$). (3) Our proposed preference optimization phase enables the strong model to learn from errors made by the weak supervisor, **ultimately surpassing the strong model fine-tuned on gold-standard solutions** (i.e., strong ceiling) in challenging scenarios, such as level 4-5 MATH problems.

To more accurately approximate future scenarios, we conduct additional experiments on OlympicArena (Huang et al., 2024), an extremely challenging dataset with *no* definitive ground truth answers. Llama3-8b-instruct (AI@Meta, 2024), despite its smaller size, has been aligned and proven to effectively supervise the larger Llama3-70b, whose potential has not yet been fully realized. Moreover, our proposed two-stage training approach outperforms full weak fine-tuning by **3.19** points.

2 Preliminaries

2.1 Typical Learning Paradigms for LLMs

We outline common learning paradigms in large model training, primarily characterized by the need for ground truth answers and supervision from stronger models as shown in Tab. 1: **Generic-supervised learning** (Ouyang et al., 2022; Yuan et al., 2023), which relies on abundant labeled data, often proves impractical due to its high cost and complexity. Alternatives such as **distillation-based learning** (Lee et al., 2023; Peng et al., 2023; An et al., 2023; Agarwal et al., 2023; Chen et al., 2023) employ a stronger model to guide a weaker one, enhancing performance through knowledge distillation, but this approach only narrows performance gap marginally on tasks not well-represented in the imitation data (Gudibande et al., 2023). **Self-**

| | G.T. Answer | Stronger Model |
|--------------------|-------------|----------------|
| Generic-supervised | ✓ | — |
| Distillation-based | ✗ | ✓ |
| Self-improvement | ✓ | — |
| Semi-supervised | ✓ | — |
| Weak-to-strong | ✗ | ✗ |

Table 1: Typical Learning Paradigms for LLMs. “✓” and “✗” indicate whether supervision is required, and “—” indicates it is optional. “G.T.” represents Ground Truth.

improvement learning (Zelikman et al., 2022; Yuan et al., 2023; Singh et al., 2023; Hosseini et al., 2024) utilizes the model’s own outputs for updates, still requiring some external accurate feedback for effectiveness (Huang et al., 2023; Tyen et al., 2023), while **semi-supervised learning** (Tong et al., 2024; Hase et al., 2024; Sun et al., 2024) leverages a small set of high-quality data to refine the model’s capabilities. In contrast, **weak-to-strong learning** (Burns et al., 2023) uniquely utilizes weaker supervisory signals to enhance model performance when no superior model or ground truth is available, showing significant potential for tackling complex problems without stringent supervisory conditions.

2.2 Weak-to-Strong Reasoning Setup

In this paper, we address reasoning tasks in the weak-to-strong setting, as illustrated in Tab. 2. First, we examine mathematical reasoning tasks, such as those in GSM8k and MATH. These tasks require each step of the reasoning process to demonstrate fundamental mathematical problem-solving skills, including problem comprehension and algebraic operations, and build upon the previous steps. This imposes higher demands on the model’s learning and generalization capabilities. Unlike classification tasks, where models can rely on superficial pattern extrapolation or recognition, reasoning tasks offer minimal benefit from guessing. Then, we use a **weak model** (e.g., Llama2-7b) with a certain degree of mathematical problem-solving ability,² denoted as m . This model acts analogously to human supervisors with limited expertise in the era of superintelligence. Besides, we only have a **set of questions** $\mathcal{Q} = \{q_i\}$ without ground truth answers and the goal is to improve the reasoning capability of a **strong model** \mathcal{M} (e.g., Llama2-70b). To implement this, following Burns et al. (2023), we randomly divide the original training set into two equal parts, $\mathcal{D}_{\text{gold},1}$ and $\mathcal{D}_{\text{gold},2}$. The weak model

²Otherwise, the weak model can hardly provide useful supervision.

| Role | weak model | strong model | task question |
|-----------------|---|--------------|---|
| Analogue | Llama2-7b + SFT($\mathcal{D}_{\text{gold},1}$) | Llama2-70b | $\mathcal{Q} \in \text{GSM8K}$ $\in \text{MATH}$ |

Table 2: Weak-to-Strong Reasoning Setup.

is initially fine-tuned using $\mathcal{D}_{\text{gold},1}$ where the gold solutions are available, resulting in a weak model with some problem-solving capability, i.e. m . In contrast, the strong model can only access the questions from $\mathcal{D}_{\text{gold},2}$, without reasoning chains or final answers, i.e., \mathcal{Q} .

3 Methodology

In this section, we propose a weak-to-strong training method designed to maximize the use of weak data and to elicit the strong model’s innate talent. First, we identify potentially positive samples in the absence of ground truth and external signals. During Stage I, we exclusively utilize this subset of data for supervised fine-tuning. Then once the strong model has achieved a certain level of reasoning proficiency, we employ the full weak data, particularly the potentially negative samples in Stage II via preference learning-based approaches like DPO (Rafailov et al., 2023), encouraging the strong model to learn from mistakes made by the weaker model. The whole framework is depicted in Fig. 2.

3.1 Stage I: Learn from “Positive” Samples

Given a weak model m and a series of math problems \mathcal{Q} **without** ground truth, m generates weak data $\mathcal{D}_{\text{weak}} = \{q_i, c_{\text{weak},i}, a_{\text{weak},i}\}$, where $q_i \in \mathcal{Q}$, $c_{\text{weak},i}$ represents a reasoning chain, and $a_{\text{weak},i}$ represents the final answer. The correctness of $a_{\text{weak},i}$ is unknown. The central challenge is: **how can we maximize the use of m and $\mathcal{D}_{\text{weak}}$ to fully enhance and recover the mathematical reasoning capabilities of a stronger model \mathcal{M} ?**

3.1.1 Full Weak Fine-Tuning

Our initial strategy is to fine-tune the stronger model \mathcal{M} across the entirety of the weak dataset $\mathcal{D}_{\text{weak}}$. While prior research (Burns et al., 2023) has validated the effectiveness of this approach in text classification tasks, its efficacy in reasoning tasks remains unexplored. We have therefore embarked on an investigation to determine whether the phenomenon of *weak-to-strong generalization* can also enhance the reasoning capabilities of \mathcal{M} in this less examined domain.

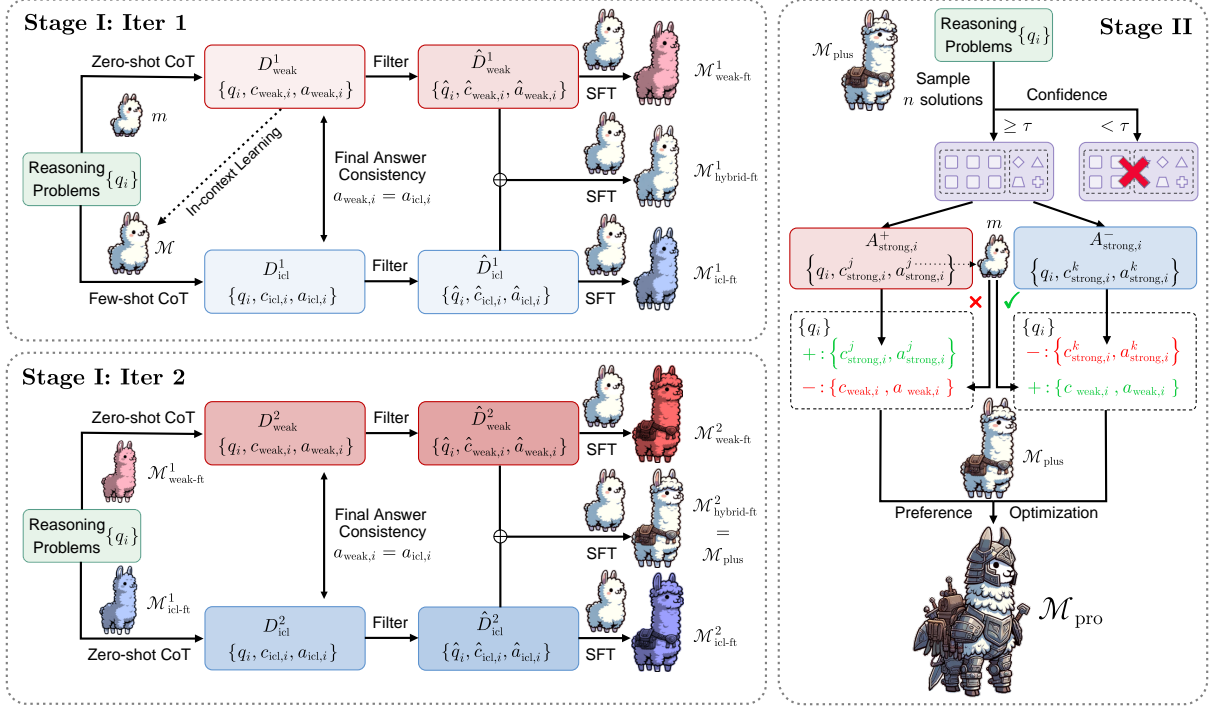


Figure 2: **Overview of our method evolving from $\mathcal{M} \rightarrow \mathcal{M}_{\text{plus}} \rightarrow \mathcal{M}_{\text{pro}}$.** **Left:** we utilize *final answer consistency* to selectively filter weak and icl data from diverse sources, which is used to fine-tune the strong model \mathcal{M} and obtain $\mathcal{M}_{\text{plus}}$ with enhanced mathematical reasoning capabilities. **Right:** we leverage the confidence of $\mathcal{M}_{\text{plus}}$ to identify contrastive samples for performance optimization, resulting in a more robust strong model \mathcal{M}_{pro} .

3.1.2 Weak In-Context Learning

Another straightforward approach is in-context learning (ICL, Dong et al. (2023b)), which requires only several training samples as demonstrations in the prompt. Specifically, we randomly select four samples from $\mathcal{D}_{\text{weak}}$ as demonstrations. Since we do not have access to the ground truth, these demonstrations cannot be provably correct.

3.1.3 Weak-ICL Fine-Tuning

Given that models can mimic weak errors through supervised fine-tuning (Charikar et al., 2024; Lang et al., 2024), we propose refining $\mathcal{D}_{\text{weak}}$ before use, instead of using all data blindly. Additionally, we seek to harness the innate abilities of the strong model activated via in-context learning. Building on these two ideas, we introduce *weak-icl fine-tuning*, employing both weak data $\mathcal{D}_{\text{weak}}$ and “icl data” $\mathcal{D}_{\text{icl}} = \{q_i, c_{\text{icl},i}, a_{\text{icl},i}\}$, where $q_i \in \mathcal{Q}$, $c_{\text{icl},i}$ and $a_{\text{icl},i}$ are generated by \mathcal{M} with few-shot demonstrations,³ as higher-quality supervision signals.

Note that, for both $\mathcal{D}_{\text{weak}}$ and \mathcal{D}_{icl} , we cannot determine whether a certain answer is correct or not.

³Experiments in §4.3 show that despite ICL being affected by demonstration selection, our method can achieve further improvements accordingly beyond ICL.

Nonetheless, when two models, employing distinct data representations, converge on the same answer in an open-ended task, it is indicative of a higher likelihood of accuracy. This phenomenon supports the reliability of the results when consistency is observed across different methodologies. We thus compare $\mathcal{D}_{\text{weak}}$ and \mathcal{D}_{icl} generated by the weak model and strong model, respectively, and select $\hat{\mathcal{D}}_{\text{weak}}$ and $\hat{\mathcal{D}}_{\text{icl}}$ if $a_{\text{weak},i} = a_{\text{icl},i}$, for subsequent supervised fine-tuning. We call this approach *final answer consistency*. Considering the combination of the two sets of data, we can obtain three versions of enhanced fine-tuned strong models:

- $\mathcal{M}_{\text{weak-ft}}$: \mathcal{M} fine-tuned on $\hat{\mathcal{D}}_{\text{weak}}$.
- $\mathcal{M}_{\text{icl-ft}}$: \mathcal{M} fine-tuned on $\hat{\mathcal{D}}_{\text{icl}}$.
- $\mathcal{M}_{\text{hybrid-ft}}$: \mathcal{M} fine-tuned on the union of $\hat{\mathcal{D}}_{\text{weak}}$ and $\hat{\mathcal{D}}_{\text{icl}}$.

Iterative Training Upon closed examination of $\mathcal{M}_{\text{weak-ft}}$ and $\mathcal{M}_{\text{icl-ft}}$, we see that they still satisfy the condition of having different data representations, as they are trained on data from different sources. Hence, we can perform iterative training to bootstrap performance. We denote the initial round of supervised fine-tuning data as $\hat{\mathcal{D}}_{\text{weak}}^1$ and $\hat{\mathcal{D}}_{\text{icl}}^1$, resulting in models $\mathcal{M}_{\text{weak-ft}}^1$, $\mathcal{M}_{\text{icl-ft}}^1$, and

QUESTION (q_i): John has five more roommates than **twice** as many as Bob. If Bob has 10 roommates, how many roommates does John have?

WEAK RESPONSE ($\{c_{\text{weak},i}, a_{\text{weak},i}\}$): John has $10+5=15$ roommates. The answer is **15**.

SELF RESPONSE 1 ($\{c_{\text{strong},i}^1, a_{\text{strong},i}^1\} \in A_{\text{strong},i}^+$): Bob has 10 roommates. Twice as many as Bob is $2*10 = 20$ roommates. John has 5 more roommates than twice as many as Bob, so John has $20+5 = 25$ roommates. The answer is **25**.

SELF RESPONSE 2 ($\{c_{\text{strong},i}^2, a_{\text{strong},i}^2\} \in A_{\text{strong},i}^+$): Let x be the number of roommates Bob has. John has 5 more roommates than twice as many as Bob, so John has $2x+5$ roommates. Bob has 10 roommates, so $x=10$. John has $2*10+5 = 25$ roommates. The answer is **25**.

Table 3: A real case example. Given a math question, the **incorrect** “weak response” is generated by m , while the two **correct** “self responses” are sampled from $A_{\text{strong},i}^+$ self-generated by $\mathcal{M}_{\text{plus}}$. Benefiting from dual solutions in the training data during Stage I, $\mathcal{M}_{\text{plus}}$ is able to generate different reasoning paths that converge to the same final answer. Through Stage II, $\mathcal{M}_{\text{plus}}$ learns to avoid m ’s error of overlooking the key word “**twice**” in calculations.

$\mathcal{M}_{\text{hybrid-ft}}^1$. In the second iteration, we obtain zero-shot solutions from $\mathcal{M}_{\text{weak-ft}}^1$ applied to \mathcal{Q} to construct $\mathcal{D}_{\text{weak}}^2$, and those from $\mathcal{M}_{\text{icl-ft}}^1$ to construct $\mathcal{D}_{\text{icl}}^2$. Here, the subscripts “weak” and “icl” indicate the initial data source. Then we apply final answer consistency to obtain $\hat{\mathcal{D}}_{\text{weak}}^2$ and $\hat{\mathcal{D}}_{\text{icl}}^2$. Following another round of supervised fine-tuning, we have:

- $\mathcal{M}_{\text{weak-ft}}^2$: \mathcal{M} fine-tuned on $\hat{\mathcal{D}}_{\text{weak}}^2$.
- $\mathcal{M}_{\text{icl-ft}}^2$: \mathcal{M} fine-tuned on $\hat{\mathcal{D}}_{\text{icl}}^2$.
- $\mathcal{M}_{\text{hybrid-ft}}^2$: \mathcal{M} fine-tuned on the union of $\hat{\mathcal{D}}_{\text{weak}}^2$ and $\hat{\mathcal{D}}_{\text{icl}}^2$.

Note that the iterative training step is optional; it may lead to performance degradation when data quality is too low or the model overfits.

3.2 Stage II: Learn from “Negative” Samples

We denote the final iteration of $\mathcal{M}_{\text{hybrid-ft}}$ from Stage I as $\mathcal{M}_{\text{plus}}$, which has learned dual mathematical solutions and holds potential for further enhancement. Next, we apply preference optimization techniques to strategically utilize the potentially erroneous subset of the original weak dataset $\mathcal{D}_{\text{weak}}$ generated by m , which allows the strong model to identify and avoid similar errors in future reasoning processes. The key factor lies in how to construct contrastive samples for learning.

Without access to ground truth, the current strong model with enhanced reasoning capabilities identifies the most likely correct answers based

on its confidence. Specifically, for each question $q_i \in \mathcal{Q}$, we sample n responses from $\mathcal{M}_{\text{plus}}$, and define the probability of the answer that appears most frequently among these responses as *confidence*. When the confidence falls below a specified threshold τ , we consider the model’s judgment on this question unreliable and therefore discard it. Conversely, if the confidence is no less than τ , we regard the model as capable of solving the question and proceed to construct contrastive samples as follows.

- For a question q_i where $\mathcal{M}_{\text{plus}}$ is confident, we denote the most confident answer as $a_{\text{strong},i}^+$ and $P(a_{\text{strong},i}^+) \geq \tau$. It can be considered as the “correct” answer according to $\mathcal{M}_{\text{plus}}$. For instance, if we set $\tau = 0.6$ and 8 out of 10 sampled responses have the same final answer “4.2”, we say that $\mathcal{M}_{\text{plus}}$ considers “4.2” to be the correct answer to this question, i.e. $a_{\text{strong},i}^+ = 4.2$.
- Then we divide the sampled n responses of $\mathcal{M}_{\text{plus}}$ to q_i into two sets: $A_{\text{strong},i}^+ = \{c_{\text{strong},i}^j, a_{\text{strong},i}^j\}$ where $a_{\text{strong},i}^j = a_{\text{strong},i}^+$; $A_{\text{strong},i}^- = \{c_{\text{strong},i}^k, a_{\text{strong},i}^k\}$ where $a_{\text{strong},i}^k \neq a_{\text{strong},i}^+$. In the above example, $|A_{\text{strong},i}^+| = 8$ and $|A_{\text{strong},i}^-| = 2$.
- If the weak model holds an answer that the enhanced model considers “correct”, that is, $a_{\text{weak},i} = a_{\text{strong},i}^+$, we treat the weak model’s response $\{c_{\text{weak},i}, a_{\text{weak},i}\}$ as chosen response and randomly select a rejected response from $A_{\text{strong},i}^-$. Otherwise, if $a_{\text{weak},i} \neq a_{\text{strong},i}^+$, we treat $\{c_{\text{weak},i}, a_{\text{weak},i}\}$ as rejected response and randomly select a chosen response from $A_{\text{strong},i}^+$. Examples are shown in Tab. 3.

Further training $\mathcal{M}_{\text{plus}}$ on these samples enables it to distinguish between correct and incorrect solutions, leading to a stronger model \mathcal{M}_{pro} .

4 Experiments

4.1 Datasets

GSM8K (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021) are two widely used datasets for mathematical reasoning, and MATH comprises more challenging competition problems. The data statistics we use are presented in Tab. 4. Particularly, to ensure a sufficient amount of training data for developing preliminary mathematical skills in the weak model, we augment the GSM8K training set with the data constructed by Chern et al. (2023). Further details are available in §A.1.

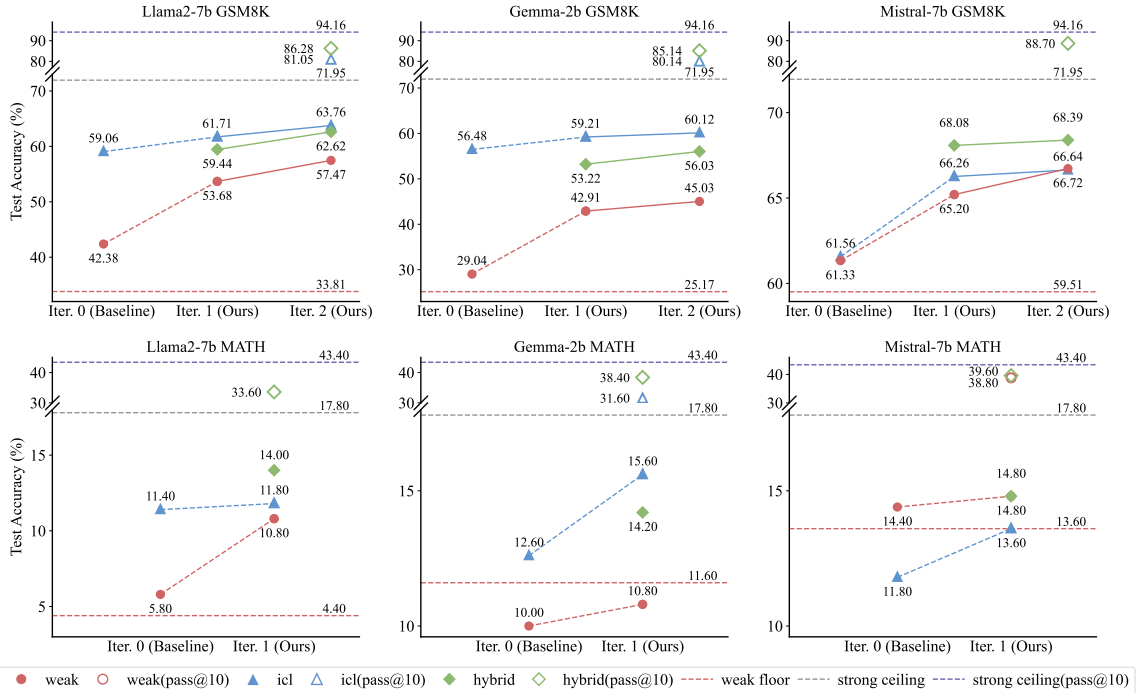


Figure 3: Main results of Stage I. “Iter. 0” presents the performance of two baselines, where “weak” indicates full weak fine-tuning, i.e., naively fine-tuning on the entire weak data, and “icl” refers to weak ICL without fine-tuning. Models connected by a line mean that they share the same training data sources. Results below “strong ceiling” present test accuracy via greedy decoding, while those above show pass@k scores ($k = 10$ and temperature = 1.0). For simplicity, we only present the pass@k scores of $\mathcal{M}_{\text{hybrid-ft}}$ and checkpoints that surpass it using greedy decoding, and full results are provided in §A.4.2.

| | # $\mathcal{D}_{\text{gold},1}$ | # $\mathcal{D}_{\text{gold},2}$ | # Test |
|-------|---------------------------------|---------------------------------|--------|
| GSM8K | 7,000 | 7,000 | 1,319 |
| MATH | 6,000 | 6,000 | 500 |

Table 4: Data Statistics. $\mathcal{D}_{\text{gold},1}$ and $\mathcal{D}_{\text{gold},2}$ are subsets of the training set. The weak model uses $\mathcal{D}_{\text{gold},1}$ to cultivate initial mathematical skills, while the strong model can only access questions from $\mathcal{D}_{\text{gold},2}$ without ground truths.

4.2 Experiment Settings

We use Llama2-70b as the strong model and employ three weak models from different families: Llama2-7b, Gemma-2b, and Mistral-7b. We apply full parameter fine-tuning to the weak models on $\mathcal{D}_{\text{gold},1}$, and consistently adopt LoRA (Hu et al., 2022) for fine-tuning the strong model. In Stage I, we perform two rounds of iterations on GSM8K and one round on MATH according to the principles of iteration outlined in §3.1. In Stage II, we adopt two preference learning-based approaches, DPO (Rafailov et al., 2023) and ORPO (Hong et al., 2024). Details are provided in §A.2.

We evaluate the accuracy on the test set. The performance of the weak model m is defined as the “weak floor”. The performance of the strong model \mathcal{M} , fine-tuned with data containing gold solutions from $\mathcal{D}_{\text{gold},2}$, is termed the “strong ceiling”. It

represents the upper limit of the capabilities that the strong model can achieve with $\mathcal{D}_{\text{gold},2}$.

4.3 Results of Stage I

The main results of Stage I on both GSM8K and MATH datasets are depicted in Fig. 3.⁴ Notably, in the MATH experiments, we randomly sample additional data that is not chosen based on the final answer consistency, due to the small amount available. Please refer to §A.4.1 for details. According to Fig. 3, we have the following observations.

Weak-ICL fine-tuning demonstrates a notable enhancement. Using our proposed method, the performance of the strong model, supervised *only* by the weak Gemma-2b with 25.17 accuracy on GSM8K (without any gold answers), can be improved up to 60.12, surpassing naive full weak fine-tuning by 31.08, and $\mathcal{M}_{\text{plus}}$ (i.e., $\mathcal{M}_{\text{hybrid-ft}}^2$) outperforms it by 26.99. This verifies the effectiveness of data refining before supervised fine-tuning. Also, experimental results show that the mathematical reasoning capabilities of the strong model are increasingly recovered as the weak

⁴We do not incorporate the zero-shot performance of the strong model in the main results as it is significantly lower than that of weak ICL. See §A.4.5 for further details.




Figure 4: Results on GSM8K supervised by Gemma-2b. ■ and ■ are under original demonstrations, and ▨ and ▨ are under carefully selected demonstrations.

model improves, a conclusion verified by Liu and Alahi (2024) on classification tasks. In detail, the performance on GSM8K gradually improves for Gemma-2b, Llama-7b, and Mistral-7b (25.17 \rightarrow 33.81 \rightarrow 59.51). Hence, the maximum performance of the strong model, trained with data generated by these models, also progressively enhances (60.12 \rightarrow 63.76 \rightarrow 68.39).

$\mathcal{M}_{\text{hybrid-ft}}$ achieves the highest pass@k scores. As expected, $\mathcal{M}_{\text{hybrid-ft}}$ achieves the highest pass@k scores in the weak-to-strong setting, benefiting from its training data that incorporates two types of solutions—one from the weak model, and another from the strong model. This diversity enhances the robustness of the model by reducing the likelihood of overfitting. Additionally, the performance of $\mathcal{M}_{\text{icl-ft}}$ generally surpasses that of $\mathcal{M}_{\text{weak-ft}}$, which can be attributed to variations in process-level accuracy and possibly the solution format. Detailed analyses are conducted in §A.3.

Naive fine-tuning is inadequate for weak-to-strong reasoning. When using Gemma-2b as the weak model on the MATH dataset, full weak fine-tuning underperforms compared to the weak floor (10.0 v.s. 11.6). This indicates that naive fine-tuning, though successfully applied to classification, chess, and reward modeling tasks (Burns et al., 2023), falls short for intricate reasoning tasks, particularly those of substantial difficulty like questions in MATH. In contrast, our weak-icl fine-tuning method effectively bridges the gap, offering an effective and scalable solution for the weak-to-strong reasoning challenge.

Effect of ICL Performance Given that the efficacy of weak-icl fine-tuning partially depends on the effectiveness of weak ICL, we further explore

| Weak Model | Test Accuracy | | |
|--------------|---------------|---------------|---------------|
| | I | II. DPO | II. ORPO |
| GSM8K | | | |
| Llama2-7b | 62.62 | 66.19 (+3.57) | 68.16 (+5.54) |
| Gemma-2b | 56.03 | 64.52 (+8.49) | 63.91 (+7.88) |
| Mistral-7b | 68.39 | 70.96 (+2.57) | 72.18 (+3.79) |
| MATH | | | |
| Llama2-7b | 14.00 | 12.00 (-2.00) | 15.00 (+1.00) |
| Gemma-2b | 14.20 | 11.60 (-2.60) | 16.00 (+1.80) |
| Mistral-7b | 14.80 | 13.40 (-1.40) | 17.00 (+2.20) |

Table 5: Main results of Stage II.

how enhancing ICL performance through careful selection of demonstrations affects the performance of weak-icl fine-tuning. Fig. 4 shows the test accuracy on GSM8K using Gemma-2b as the weak model under a different set of demonstrations.

The results indicate that the performance of weak ICL with this particular group of demonstrations increases from the original 56.48 to 64.06. We then regenerate \mathcal{D}_{icl} with these demonstrations in the prompt and fine-tune the strong model on $\hat{\mathcal{D}}_{\text{icl}}$, which is selectively curated through final answer consistency. This further improves performance from 64.06 to 64.75, confirming the utility of self-directed data curation. It is worth noting that although weak ICL holds the potential for high performance, the selection of effective demonstrations in a weak-to-strong framework is a non-trivial thing, and is beyond the scope of this paper.

4.4 Results of Stage II

As discussed in §3.2, we employ the final iteration of $\mathcal{M}_{\text{hybrid-ft}}$ as $\mathcal{M}_{\text{plus}}$ for subsequent preference learning. The experimental results in §4.3 validate this checkpoint achieves higher pass@k and possesses significant potential for further refinement.

As shown in Tab. 5, our method for constructing positive and negative samples effectively enhance the strong model’s math reasoning capabilities. On GSM8K, both DPO and ORPO consistently achieve significant improvements using our constructed datasets, notably resulting in an increase of 8.49 points when supervised by Gemma-2b. Despite the inherently challenging nature of MATH problem, which compromises the strong model’s judgment and introduces inaccuracies in the training data, our method still achieves improvements on MATH through ORPO by at least 1 point.⁵

⁵Pang et al. (2024); Xu et al. (2024); Yuan et al. (2024) demonstrate that DPO can cause performance degradation on MATH due to the lack of regularization in its loss.




Figure 5: Accuracy across varying difficulty levels on the MATH test set. We use ORPO to obtain \mathcal{M}_{pro} .

Data Construction Recipe When constructing preference data, we always use weak responses generated by the weak model as one of the chosen/rejected responses, instead of relying exclusively on self-generated data. We also test the self-generated setting on GSM8K using Llama2-7b as the weak model, where both chosen and rejected responses are generated by the strong model itself. The DPO test accuracy in this setting is 62.40 (-0.22), indicating a slight performance degradation. Without ground truth, the constructed positive and negative samples actually correspond to the more frequently and less frequently occurring answers, respectively, and are related to the answers the model tends to choose. Since preference optimization essentially performs ranking, the potential benefit of this self-generated setting is minimal. Therefore, incorporating weak data signals in the preference data construction process proves to be a better approach.

4.5 Analysis

For further analysis, we examine the accuracy across different difficulty levels in the MATH test set (See §A.1.2 for data statistics).

As shown in Fig. 5, the strong model exhibits better generalization on easier problems. Specifically, even though Llama2-7b achieves only 6.98 points accuracy on level 1 problems, Llama2-70b can achieve an accuracy exceeding 30 points after training using this weak supervision. For more challenging problems (levels 4-5), \mathcal{M}_{pro} , enhanced with ORPO, even **surpasses the strong ceiling obtained by supervised fine-tuning solely on gold solutions**. This phenomenon serves to validate the effectiveness of learning from incorrect data.

4.6 Experiments Closer to Future Scenarios

In preliminary tests with Llama3-70b (AI@Meta, 2024), we observe that on GSM8K and MATH,

| | Test Accuracy |
|--|---------------|
| Weak Floor | 11.82 |
| Full Weak FT | 12.46 |
| Weak ICL | 8.63 |
| $\mathcal{M}_{\text{weak-ft}}^1$ | 12.78 |
| $\mathcal{M}_{\text{icl-ft}}^1$ | 9.58 |
| $\mathcal{M}_{\text{hybrid-ft}}^1$ | 11.18 |
| $\mathcal{M}_{\text{weak-ft}}^2$ | <u>13.10</u> |
| $\mathcal{M}_{\text{icl-ft}}^2$ | 11.50 |
| $\mathcal{M}_{\text{hybrid-ft}}^2$ ($\mathcal{M}_{\text{plus}}$) | 11.82 |
| \mathcal{M}_{pro} | 15.65 |

Table 6: Results on *OlympicArena* using *Llama3 family*. The best result is in **bold**, and the best result of supervised fine-tuning is underlined.

Llama3-70b can largely unlock its potential through in-context learning, with marginal or even adverse impacts from parameter updates due to training instabilities. Consequently, we focus on a more challenging dataset developed after the release of Llama3-70b, *OlympicArena* (Huang et al., 2024), to simulate a more realistic future scenario.

We only consider English questions in *OlympicArena*, excluding the CODE (Code Generation) and OT (Others) problem types that require case-based or expert evaluation. This results in 6,020 training data without solutions and final answers, and 313 test data with final answers to assess the performance of different methods. We use Llama3-8b-instruct (without initial fine-tuning on a subset of training data) as the weak model and Llama3-70b as the strong model to be improved. This configuration more closely resembles future real-world weak-to-strong scenarios.

Experimental results are displayed in Tab. 6. “Weak Floor” represents the zero-shot performance of Llama3-8b-instruct, “Full Weak FT” denotes the performance of Llama3-70b after supervised fine-tuning on the full set (i.e., 6,020) of weak solutions generated by Llama3-8b-instruct on the training set, and “Weak ICL” indicates the performance

of Llama3-70b under 4-shot weak demonstrations generated by Llama3-8b-instruct. Despite having more parameters, Llama3-70b under in-context learning still performs lower than the zero-shot performance of Llama3-8b-instruct due to insufficient mining capabilities.

$\mathcal{M}_{\text{weak-ft}}^1$, obtained by our proposed weak-icl fine-tuning method, achieves higher performance than Full Weak FT with fewer training data (i.e., 746), outperforming it by 0.32 points. After the second stage of preference optimization, which further exploits the weak model and training questions without answers, the strong model’s performance is improved by an additional 3.19 points over Full Weak FT. This demonstrates the robustness and generalizability of our method in scenarios closer to future conditions.

5 Related Work

LLM Training LLMs can improve their task-solving abilities and alignment with human instructions through supervised fine-tuning (Zhang et al., 2023; Dong et al., 2023a; Lv et al., 2023b,a), with the process heavily influenced by the quality of training data (Zhou et al., 2023a; Wang et al., 2023b). In this paper, we investigate the potential of learning from weak supervision. To better utilize both positive and negative feedback, there are additional training methods such as reinforcement learning from human feedback (RLHF, Ouyang et al. (2022); Bai et al. (2022)) and DPO (Rafailov et al., 2023)). DPO modifies reward functions in RLHF and has led to variants like ORPO (Hong et al., 2024) and SimPO (Meng et al., 2024), which improve stability and performance. This paper explores the capabilities of DPO and ORPO using our constructed contrastive samples in a weak-to-strong setting.

Mathematical Reasoning The exploration of mathematical reasoning within LLMs has been a focal point for evaluating their cognitive capabilities akin to human reasoning (Qiao et al., 2023; Lu et al., 2023). Researchers have developed various methods to enhance the mathematical reasoning capabilities of LLMs after pre-training, which can be broadly classified into two categories: (1) Prompting: Some works (Kojima et al., 2022; Wei et al., 2022; Zhou et al., 2023b; Liu et al., 2023) aim to elicit the intrinsic reasoning abilities of LLMs by specific prompting engineering, without updating the model parameters; (2) Fine-tuning: Another

line of studies focuses on generating a more extensive and higher-quality collection of question-answer pairs (Yu et al., 2023; Wang et al., 2023c,a). Through supervised fine-tuning and preference optimization (Luo et al., 2023; Azerbayev et al., 2023; Mitra et al., 2024; Xu et al., 2024), the models can achieve significant improvements in their mathematical problem-solving capabilities.

6 Conclusion

In this paper, we explore the efficacy of weak-to-strong framework in complex reasoning tasks. We introduce a new method that elicits strong capabilities using weak supervision, without relying on annotations from humans or more advanced models. This method focuses on the strong model’s ability to autonomously refine its training data, even if it has not learned the task before. By iteratively expanding its learning scope, the strong model continuously broadens its reasoning skills. This self-directed data curation is crucial for scaling up the enhancement of AI reasoning capabilities, making the model more independent and effective in its developmental trajectory. Through this work, we seek to illuminate new pathways for AI development, emphasizing the critical role of innovative model supervision in advancing AGI and beyond.

Limitations

In our experiments, we use Llama2-70b as a proxy for hypothetical superintelligent models of the future. We acknowledge that there might be performance discrepancies compared to a genuine future advanced model. Nonetheless, our efforts lay the groundwork for investigating methodologies in weak-to-strong reasoning. Additionally, this paper does not explore supervision at the process level, such as the model’s ability to learn from partially correct data (Ni et al., 2023; Lightman et al., 2023). In the weak-to-strong scenario, the presence of non-negligible errors and noise at the process level yields only limited performance improvements in our early experiments, thereby posing challenges for future research.

Acknowledgments

We thank Xuefeng Li, Haoyang Zou, and Ting Wu for their valuable insights during discussions. This work was supported by Shanghai Artificial Intelligence Laboratory, SJTU SEIEE - ByteDance Large Language Model Joint Laboratory.

References

- Rishabh Agarwal, Nino Vieillard, Piotr Stanczyk, Sabela Ramos, Matthieu Geist, and Olivier Bachem. 2023. **GKD: generalized knowledge distillation for auto-regressive sequence models.** *CoRR*, abs/2306.13649.
- AI@Meta. 2024. **Llama 3 model card.**
- Sam Altman, Greg Brockman, and Ilya Sutskever. 2023. **Governance of superintelligence.** <https://openai.com/index/governance-of-superintelligence/>.
- Shengnan An, Zexiong Ma, Zeqi Lin, Nanning Zheng, Jian-Guang Lou, and Weizhu Chen. 2023. **Learning from mistakes makes LLM better reasoner.** *CoRR*, abs/2310.20689.
- Zhangir Azerbayev, Hailey Schoelkopf, Keiran Paster, Marco Dos Santos, Stephen McAleer, Albert Q. Jiang, Jia Deng, Stella Biderman, and Sean Welleck. 2023. **Llemma: An open language model for mathematics.** *CoRR*, abs/2310.10631.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosiute, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemí Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. 2022. **Constitutional AI: harmfulness from AI feedback.** *CoRR*, abs/2212.08073.
- Samuel R. Bowman, Jeeyoon Hyun, Ethan Perez, Edwin Chen, Craig Pettit, Scott Heiner, Kamile Lukosiute, Amanda Askell, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Christopher Olah, Daniela Amodei, Dario Amodei, Dawn Drain, Dustin Li, Eli Tran-Johnson, Jackson Kernion, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Liane Lovitt, Nelson Elhage, Nicholas Schiefer, Nicholas Joseph, Noemí Mercado, Nova DasSarma, Robin Larson, Sam McCandlish, Sandipan Kundu, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Timothy Telleen-Lawton, Tom Brown, Tom Henighan, Tristan Hume, Yuntao Bai, Zac Hatfield-Dodds, Ben Mann, and Jared Kaplan. 2022. **Measuring progress on scalable oversight for large language models.** *CoRR*, abs/2211.03540.
- Collin Burns, Pavel Izmailov, Jan Hendrik Kirchner, Bowen Baker, Leo Gao, Leopold Aschenbrenner, Yining Chen, Adrien Ecoffet, Manas Joglekar, Jan Leike, Ilya Sutskever, and Jeff Wu. 2023. **Weak-to-strong generalization: Eliciting strong capabilities with weak supervision.** *CoRR*, abs/2312.09390.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Kaijie Zhu, Hao Chen, Linyi Yang, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S. Yu, Qiang Yang, and Xing Xie. 2023. **A survey on evaluation of large language models.** *CoRR*, abs/2307.03109.
- Moses Charikar, Chirag Pabbaraju, and Kirankumar Shiragur. 2024. **Quantifying the gain in weak-to-strong generalization.** *Preprint*, arXiv:2405.15116.
- Kai Chen, Chunwei Wang, Kuo Yang, Jianhua Han, Lanqing Hong, Fei Mi, Hang Xu, Zhengying Liu, Wenyong Huang, Zhenguo Li, Dit-Yan Yeung, Lifeng Shang, Xin Jiang, and Qun Liu. 2023. **Gaining wisdom from setbacks: Aligning large language models via mistake analysis.** *CoRR*, abs/2310.10477.
- Ethan Chern, Haoyang Zou, Xuefeng Li, Jiewen Hu, Kehua Feng, Junlong Li, and Pengfei Liu. 2023. **Generative ai for math: Abel.** <https://github.com/GAIR-NLP/abel>.
- Paul F. Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. **Deep reinforcement learning from human preferences.** In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 4299–4307.
- 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.** *CoRR*, abs/2110.14168.
- Guanting Dong, Hongyi Yuan, Keming Lu, Chengpeng Li, Mingfeng Xue, Dayiheng Liu, Wei Wang, Zheng Yuan, Chang Zhou, and Jingren Zhou. 2023a. **How abilities in large language models are affected by supervised fine-tuning data composition.** *CoRR*, abs/2310.05492.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, Lei Li, and Zhifang Sui. 2023b. **A survey for in-context learning.** *CoRR*, abs/2301.00234.
- Run-Ze Fan, Xuefeng Li, Haoyang Zou, Junlong Li, Shwai He, Ethan Chern, Jiewen Hu, and Pengfei Liu. 2024. **Reformatted alignment.** *CoRR*, abs/2402.12219.
- Arnav Gudibande, Eric Wallace, Charlie Snell, Xinyang Geng, Hao Liu, Pieter Abbeel, Sergey Levine, and Dawn Song. 2023. **The false promise of imitating proprietary llms.** *CoRR*, abs/2305.15717.
- Peter Hase, Mohit Bansal, Peter Clark, and Sarah Wierffe. 2024. **The unreasonable effectiveness of easy training data for hard tasks.** *CoRR*, abs/2401.06751.

- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. [Measuring mathematical problem solving with the MATH dataset](#). In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*.
- Jiwoo Hong, Noah Lee, and James Thorne. 2024. [ORPO: monolithic preference optimization without reference model](#). *CoRR*, abs/2403.07691.
- Arian Hosseini, Xingdi Yuan, Nikolay Malkin, Aaron C. Courville, Alessandro Sordoni, and Rishabh Agarwal. 2024. [V-star: Training verifiers for self-taught reasoners](#). *CoRR*, abs/2402.06457.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [Lora: Low-rank adaptation of large language models](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Jie Huang and Kevin Chen-Chuan Chang. 2023. [Towards reasoning in large language models: A survey](#). In *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 1049–1065. Association for Computational Linguistics.
- Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. 2023. [Large language models cannot self-correct reasoning yet](#). *CoRR*, abs/2310.01798.
- Zhen Huang, Zengzhi Wang, Shijie Xia, Xuefeng Li, Haoyang Zou, Ruijie Xu, Run-Ze Fan, Lyumanshan Ye, Ethan Chern, Yixin Ye, Yikai Zhang, Yuqing Yang, Ting Wu, Binjie Wang, Shichao Sun, Yang Xiao, Yiyuan Li, Fan Zhou, Steffi Chern, Yiwei Qin, Yan Ma, Jiadi Su, Yixiu Liu, Yuxiang Zheng, Shaoting Zhang, Dahua Lin, Yu Qiao, and Pengfei Liu. 2024. [Olympicarena: Benchmarking multi-discipline cognitive reasoning for superintelligent ai](#). *Preprint*, arXiv:2406.12753.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *CoRR*, abs/2310.06825.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Hunter Lang, David Sontag, and Aravindan Vijayaraghavan. 2024. [Theoretical analysis of weak-to-strong generalization](#). *Preprint*, arXiv:2405.16043.
- Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie Lu, Thomas Mesnard, Colton Bishop, Victor Carbone, and Abhinav Rastogi. 2023. [RLAIF: scaling reinforcement learning from human feedback with AI feedback](#). *CoRR*, abs/2309.00267.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. [Let’s verify step by step](#). *CoRR*, abs/2305.20050.
- Tengxiao Liu, Qipeng Guo, Yuqing Yang, Xiangkun Hu, Yue Zhang, Xipeng Qiu, and Zheng Zhang. 2023. [Plan, verify and switch: Integrated reasoning with diverse x-of-thoughts](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 2807–2822. Association for Computational Linguistics.
- Yuejiang Liu and Alexandre Alahi. 2024. [Co-supervised learning: Improving weak-to-strong generalization with hierarchical mixture of experts](#). *CoRR*, abs/2402.15505.
- Pan Lu, Liang Qiu, Wenhao Yu, Sean Welleck, and Kai-Wei Chang. 2023. [A survey of deep learning for mathematical reasoning](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 14605–14631. Association for Computational Linguistics.
- Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng, Qingwei Lin, Shifeng Chen, and Dongmei Zhang. 2023. [Wizardmath: Empowering mathematical reasoning for large language models via reinforced evol-instruct](#). *CoRR*, abs/2308.09583.
- Kai Lv, Hang Yan, Qipeng Guo, Haijun Lv, and Xipeng Qiu. 2023a. [Adalomo: Low-memory optimization with adaptive learning rate](#). *CoRR*, abs/2310.10195.
- Kai Lv, Yuqing Yang, Tengxiao Liu, Qinghui Gao, Qipeng Guo, and Xipeng Qiu. 2023b. [Full parameter fine-tuning for large language models with limited resources](#). *CoRR*, abs/2306.09782.
- Yu Meng, Mengzhou Xia, and Danqi Chen. 2024. [Simpo: Simple preference optimization with a reference-free reward](#). *Preprint*, arXiv:2405.14734.
- Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivi  re, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, L  onard Hussenot, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, Am  lie H  liou, Andrea Tacchetti, Anna Bulanova, Antonia Paterson, Beth

- Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo, Clément Crepy, Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng Yan, George Tucker, George-Cristian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski, Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, and et al. 2024. [Gemma: Open models based on gemini research and technology](#). *CoRR*, abs/2403.08295.
- Arindam Mitra, Hamed Khanpour, Corby Rosset, and Ahmed Awadallah. 2024. [Orca-math: Unlocking the potential of slms in grade school math](#). *CoRR*, abs/2402.14830.
- Ansong Ni, Jeevana Priya Inala, Chenglong Wang, Alex Polozov, Christopher Meek, Dragomir Radev, and Jianfeng Gao. 2023. [Learning math reasoning from self-sampled correct and partially-correct solutions](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Richard Yuanzhe Pang, Weizhe Yuan, Kyunghyun Cho, He He, Sainbayar Sukhbaatar, and Jason Weston. 2024. [Iterative reasoning preference optimization](#). *CoRR*, abs/2404.19733.
- Arjun Panickssery, Samuel R. Bowman, and Shi Feng. 2024. [LLM evaluators recognize and favor their own generations](#). *CoRR*, abs/2404.13076.
- Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. [Instruction tuning with GPT-4](#). *CoRR*, abs/2304.03277.
- Gokul Puthumanai, Manav Vora, Pranay Thangeda, and Melkior Ornik. 2024. [A moral imperative: The need for continual superalignment of large language models](#). *CoRR*, abs/2403.14683.
- Shuofei Qiao, Yixin Ou, Ningyu Zhang, Xiang Chen, Yunzhi Yao, Shumin Deng, Chuanqi Tan, Fei Huang, and Huajun Chen. 2023. [Reasoning with language model prompting: A survey](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 5368–5393. Association for Computational Linguistics.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Xuan Ren, Biao Wu, and Lingqiao Liu. 2024. [I learn better if you speak my language: Enhancing large language model fine-tuning with style-aligned response adjustments](#). *CoRR*, abs/2402.11192.
- Christian P. Robert. 2017. [Superintelligence: Paths, dangers, strategies](#). *CHANCE*, 30:42 – 43.
- Jitao Sang, Yuhang Wang, Jing Zhang, Yanxu Zhu, Chao Kong, Junhong Ye, Shuyu Wei, and Jinlin Xiao. 2024. [Improving weak-to-strong generalization with scalable oversight and ensemble learning](#). *CoRR*, abs/2402.00667.
- 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 F. Elsayed, Hanie Sedghi, Igor Mordatch, Isabelle Simpson, Izzeddin Gur, Jasper Snoek, Jeffrey Pennington, Jiri Hron, Kathleen Keenally, 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. 2023. [Beyond human data: Scaling self-training for problem-solving with language models](#). *CoRR*, abs/2312.06585.
- Zhiqing Sun, Longhui Yu, Yikang Shen, Weiyang Liu, Yiming Yang, Sean Welleck, and Chuang Gan. 2024. [Easy-to-hard generalization: Scalable alignment beyond human supervision](#). *CoRR*, abs/2403.09472.
- Yongqi Tong, Sizhe Wang, Dawei Li, Yifan Wang, Simeng Han, Zi Lin, Chengsong Huang, Jiaxin Huang, and Jingbo Shang. 2024. [Optimizing language model’s reasoning abilities with weak supervision](#). *Preprint*, arXiv:2405.04086.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutli Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten,

- 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](#). *CoRR*, abs/2307.09288.
- Gladys Tyen, Hassan Mansoor, Peter Chen, Tony Mak, and Victor Carbune. 2023. [LLMs cannot find reasoning errors, but can correct them!](#) *CoRR*, abs/2311.08516.
- Peiyi Wang, Lei Li, Zhihong Shao, R. X. Xu, Damai Dai, Yifei Li, Deli Chen, Y. Wu, and Zhifang Sui. 2023a. [Math-shepherd: Verify and reinforce llms step-by-step without human annotations](#). *CoRR*, abs/2312.08935.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023b. [Self-instruct: Aligning language models with self-generated instructions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 13484–13508. Association for Computational Linguistics.
- Zengzhi Wang, Rui Xia, and Pengfei Liu. 2023c. [Generative AI for math: Part I - mathpile: A billion-token-scale pretraining corpus for math](#). *CoRR*, abs/2312.17120.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Ting Wu, Xuefeng Li, and Pengfei Liu. 2024. [Progress or regress? self-improvement reversal in post-training](#). *Preprint*, arXiv:2407.05013.
- Shijie Xia, Xuefeng Li, Yixin Liu, Tongshuang Wu, and Pengfei Liu. 2024. [Evaluating mathematical reasoning beyond accuracy](#). *CoRR*, abs/2404.05692.
- Yifan Xu, Xiao Liu, Xinghan Liu, Zhenyu Hou, Yueyan Li, Xiaohan Zhang, Zihan Wang, Aohan Zeng, Zhengxiao Du, Wenyi Zhao, Jie Tang, and Yuxiao Dong. 2024. [Chatglm-math: Improving math problem-solving in large language models with a self-critique pipeline](#). *CoRR*, abs/2404.02893.
- Fangxu Yu, Lai Jiang, Haoqiang Kang, Shibo Hao, and Lianhui Qin. 2024. [Flow of reasoning: Efficient training of llm policy with divergent thinking](#). *Preprint*, arXiv:2406.05673.
- Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T. Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. 2023. [Meta-math: Bootstrap your own mathematical questions for large language models](#). *CoRR*, abs/2309.12284.
- Lifan Yuan, Ganqu Cui, Hanbin Wang, Ning Ding, Xingyao Wang, Jia Deng, Boji Shan, Huimin Chen, Ruobing Xie, Yankai Lin, Zhenghao Liu, Bowen Zhou, Hao Peng, Zhiyuan Liu, and Maosong Sun. 2024. [Advancing LLM reasoning generalists with preference trees](#). *CoRR*, abs/2404.02078.
- Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting Dong, Chuanqi Tan, and Chang Zhou. 2023. [Scaling relationship on learning mathematical reasoning with large language models](#). *CoRR*, abs/2308.01825.
- Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. 2022. [Star: Bootstrapping reasoning with reasoning](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Zhongshen Zeng, Pengguang Chen, Haiyun Jiang, and Jiaya Jia. 2023. [Challenge llms to reason about reasoning: A benchmark to unveil cognitive depth in llms](#). *CoRR*, abs/2312.17080.
- Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, and Guoyin Wang. 2023. [Instruction tuning for large language models: A survey](#). *CoRR*, abs/2308.10792.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. 2023a. [LIMA: less is more for alignment](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. 2024. [Lima: Less is more for alignment](#). *Advances in Neural Information Processing Systems*, 36.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V. Le, and Ed H. Chi. 2023b. [Least-to-most prompting enables complex reasoning in large language models](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

A Appendix

A.1 Dataset Details

A.1.1 Dataset Construction

For GSM8K, we evenly divide the original training dataset of 7,473 samples into two subsets, $\mathcal{D}_{\text{gold},1}$ and $\mathcal{D}_{\text{gold},2}$. Additionally, we supplement both $\mathcal{D}_{\text{gold},1}$ and $\mathcal{D}_{\text{gold},2}$ with the data of the same distribution developed by (Chern et al., 2023), until each contains 7,000 samples. Thus, the weak model uses $\mathcal{D}_{\text{gold},1}$, which includes both questions and gold solutions, to obtain basic problem-solving capabilities. Meanwhile, the strong model can only access a training dataset $\mathcal{Q} = \{q_i\}$, where $q_i \in \mathcal{D}_{\text{gold},2}$, consisting of 7,000 mathematical problems without ground truth answers. GSM8K also includes 1,319 test samples.

For MATH, we employ the same subset of 500 representative problems as the test set, identical to that used in (Lightman et al., 2023). We then split the remaining 12,000 samples evenly between $\mathcal{D}_{\text{gold},1}$ and $\mathcal{D}_{\text{gold},2}$, each containing 6,000 samples.

A.1.2 Statistics of MATH test set

The distribution of difficulty levels across the 500 test data samples in MATH is listed in Tab. 7.

| # L1 | # L2 | # L3 | # L4 | # L5 | # Total |
|------|------|------|------|------|---------|
| 43 | 90 | 105 | 128 | 134 | 500 |

Table 7: Data statistics of the MATH test set.

A.2 Training Details

For supervised fine-tuning in Stage I, we adopt LoRA to fine-tune the strong model \mathcal{M} with a learning rate of 1×10^{-4} and search for weight decay in the set $\{0, 0.01\}$. We run 2 epochs on GSM8K and 3 epochs on MATH, with a batch size of 8. In Stage II, we employ two preference optimization methods. For DPO, we train the enhanced strong model $\mathcal{M}_{\text{plus}}$ with a learning rate of 1×10^{-5} and run 1 epoch. For ORPO, we search for β in the set $\{0.1, 0.5, 1.0\}$ with a learning rate of 3×10^{-5} and run 1 epoch. For experiments on OlympicArena using Llama3 family, the hyperparameters are consistent with those used for GSM8K. All experiments are conducted using A100 GPUs.

When constructing contrastive samples in Stage II, we sample $n = 10$ responses at temperature = 1.0, and use a confidence threshold of $\tau = 0.6$. Normally, we evaluate using greedy decoding.

| | | Final Answer | Process-Level |
|--------------|-------------------------------------|--------------|---------------|
| GSM8K | | | |
| Llama2-7b | $\hat{\mathcal{D}}_{\text{weak}}^1$ | 89.82 | 72.50 |
| | $\hat{\mathcal{D}}_{\text{icl}}^1$ | 89.82 | 76.50 |
| Gemma-2b | $\hat{\mathcal{D}}_{\text{weak}}^1$ | 87.97 | 73.10 |
| | $\hat{\mathcal{D}}_{\text{icl}}^1$ | 87.97 | 73.80 |
| Mistral-7b | $\hat{\mathcal{D}}_{\text{weak}}^1$ | 92.38 | 80.10 |
| | $\hat{\mathcal{D}}_{\text{icl}}^1$ | 92.38 | 77.90 |
| MATH | | | |
| Llama2-7b | $\hat{\mathcal{D}}_{\text{weak}}$ | 46.11 | 32.04 |
| | $\hat{\mathcal{D}}_{\text{icl}}$ | 46.11 | 39.22 |
| Gemma-2b | $\hat{\mathcal{D}}_{\text{weak}}$ | 30.40 | 26.30 |
| | $\hat{\mathcal{D}}_{\text{icl}}$ | 31.90 | 29.90 |
| Mistral-7b | $\hat{\mathcal{D}}_{\text{weak}}$ | 24.75 | 21.50 |
| | $\hat{\mathcal{D}}_{\text{icl}}$ | 25.25 | 25.60 |

Table 8: Training accuracy of Stage I.

For calculating pass@k, we set $k = 10$ and temperature = 1.0.

A.3 Additional Analysis

A.3.1 Diversity Analysis

To investigate why $\mathcal{M}_{\text{hybrid-ft}}$ achieves high pass@k scores despite lower greedy decoding results, we explore the diversity of responses generated by $\mathcal{M}_{\text{hybrid-ft}}$ and $\mathcal{M}_{\text{icl-ft}}$. We specifically examine the frequency distribution of the number of distinct solutions for each question across the two strong model checkpoints. Given a question from $\mathcal{D}_{\text{gold},2}$, we sample $n = 10$ responses at temperature = 1.0 for each checkpoint. We consider two responses distinct if their ROUGE-L similarity is less than 0.7. We then compute the number of clusters formed by these distinct responses and plot their frequency distribution in Fig. 6.

As shown in Fig. 6, $\mathcal{M}_{\text{icl-ft}}^2$ tends to produce nearly the same sampled responses for each question in more than 36% of the instances. This indicates a limited exploration of problem-solving paths and difficulty in generating diverse, correct solutions during the sampling process. In contrast, $\mathcal{M}_{\text{hybrid-ft}}^2$ generates a variety of responses, increasing its hit rate with multiple sampling and thus achieving higher pass@k scores. Additionally, diverse solutions are crucial for robust outcomes and model generalization (Yu et al., 2024; Wu et al., 2024). In Stage II, diverse solutions also ensure the distinction between positive and negative samples, demonstrating the rationale for selecting $\mathcal{M}_{\text{hybrid-ft}}^2$ for preference optimization in Stage II.




Figure 6: Frequency distribution of the number of distinct solutions on GSM8K supervised by Llama2-7b.

A.3.2 Training Accuracy of Stage I

Tab. 8 presents the final answer accuracy and process-level accuracy for both weak data and icl data utilized in the initial round.⁶ To compute process-level accuracy, we randomly sample a maximum of 1,000 training sample from each of weak data and icl data, and evaluate them using GPT-4o following Xia et al. (2024); Zeng et al. (2023), the prompt we use is illustrated in Tab. 14. Accuracy at this level is determined strictly on the basis that there are no errors throughout the intermediate reasoning steps.

From the results we can see that despite having consistent final answer accuracy (with the exceptions of Gemma-2b and Mistral-7b on MATH using augmented training data), there are noticeable differences in process-level performance, leading to variations in the effectiveness of $\mathcal{M}_{\text{weak-ft}}$ and $\mathcal{M}_{\text{icl-ft}}$. Moreover, it is counterintuitive that models trained on icl data with relatively low process-level accuracy achieve higher performance. This might be because the models prefer self-generated solutions and can more effectively learn those that better align with their inherent distribution (Panickssery et al., 2024; Ren et al., 2024; Fan et al., 2024).

A.4 Additional Experiments

A.4.1 Details of Stage I on MATH

In the Stage I experiment conducted on the MATH dataset, it is found that the amount of training data selected via final answer consistency is so limited that the strong model can hardly learn the effective features through supervised fine-tuning. To address this, we randomly sample additional inconsistent

⁶The relatively low accuracy observed in MATH explains why we choose to perform one round of iteration.

| | Test Accuracy | # Training Data |
|------------------------------------|---------------|-----------------|
| Gemma-2b | | |
| SFT on Full Weak | 10.00 | 6,000 |
| SFT on Gold Weak | 15.60 | 644 |
| $\mathcal{M}_{\text{weak-ft}}^1$ | 11.00 | 448 |
| $\mathcal{M}_{\text{icl-ft}}^1$ | 11.40 | 448 |
| $\mathcal{M}_{\text{hybrid-ft}}^1$ | 13.20 | 448×2 |
| Mistral-7b | | |
| SFT on Full Weak | 14.40 | 6,000 |
| SFT on Gold Weak | 16.60 | 861 |
| $\mathcal{M}_{\text{weak-ft}}^1$ | 12.40 | 584 |
| $\mathcal{M}_{\text{icl-ft}}^1$ | 15.60 | 584 |
| $\mathcal{M}_{\text{hybrid-ft}}^1$ | 14.20 | 584×2 |

Table 9: Stage I results on MATH without augmenting training data.

data. Based on the weak model’s performance (Llama-7b < Gemma-2b < Mistral-7b on MATH), we supplement the data (both \hat{D}_{weak} and \hat{D}_{icl}) to 1,000 instances for Gemma-2b and 2,000 instances for Mistral-7b, and present the results in Fig. 3. The original amount of training data and test accuracy for these two weak models are shown in Tab. 9.

A.4.2 Pass@k Results

Tab. 10 summarizes the greedy decoding and pass@k results for the three variants of enhanced strong models obtained through weak-icl fine-tuning. Notably, $\mathcal{M}_{\text{hybrid-ft}}$ utilizes a training set that combines those used by $\mathcal{M}_{\text{weak-ft}}$ and $\mathcal{M}_{\text{icl-ft}}$. The results indicate that $\mathcal{M}_{\text{hybrid-ft}}$ outperforms its counterparts in terms of pass@k, achieving superior pass@k scores with margins of up to 5.23 points. The only exception occurs in the MATH dataset supervised by Llama2-7b, where the under-performance is likely due to limited training data.

The superior performance of $\mathcal{M}_{\text{hybrid-ft}}$ can be attributed to the diversity of solutions in its training set (verified in §A.3.1), validating our approach

| | | Greedy Decoding | Pass@k |
|--------------|------------------------------------|-----------------|--------------|
| GSM8K | | | |
| Llama2-7b | $\mathcal{M}_{\text{weak-ft}}^2$ | 57.47 | 77.26 |
| | $\mathcal{M}_{\text{icl-ft}}^2$ | 63.76 | 81.05 |
| | $\mathcal{M}_{\text{hybrid-ft}}^2$ | 62.62 | 86.28 |
| Gemma-2b | $\mathcal{M}_{\text{weak-ft}}^2$ | 45.03 | 71.49 |
| | $\mathcal{M}_{\text{icl-ft}}^2$ | 60.12 | 80.14 |
| | $\mathcal{M}_{\text{hybrid-ft}}^2$ | 56.03 | 85.14 |
| Mistral-7b | $\mathcal{M}_{\text{weak-ft}}^2$ | 66.72 | 85.67 |
| | $\mathcal{M}_{\text{icl-ft}}^2$ | 66.64 | 84.08 |
| | $\mathcal{M}_{\text{hybrid-ft}}^2$ | 68.39 | 88.70 |
| MATH | | | |
| Llama2-7b | $\mathcal{M}_{\text{weak-ft}}^1$ | 10.80 | 34.80 |
| | $\mathcal{M}_{\text{icl-ft}}^1$ | 11.80 | 35.00 |
| | $\mathcal{M}_{\text{hybrid-ft}}^1$ | 14.00 | 33.60 |
| Gemma-2b | $\mathcal{M}_{\text{weak-ft}}^1$ | 14.80 | 38.80 |
| | $\mathcal{M}_{\text{icl-ft}}^1$ | 13.60 | 33.60 |
| | $\mathcal{M}_{\text{hybrid-ft}}^1$ | 14.80 | 39.60 |
| Mistral-7b | $\mathcal{M}_{\text{weak-ft}}^1$ | 10.80 | 34.20 |
| | $\mathcal{M}_{\text{icl-ft}}^1$ | 15.60 | 31.60 |
| | $\mathcal{M}_{\text{hybrid-ft}}^1$ | 14.20 | 38.40 |

Table 10: Greedy decoding and pass@k results ($k = 10$ and temperature = 1.0) for the three variants of enhanced strong models obtained through weak-icl fine-tuning. The best results are in **bold**.

of adopting the final iteration of $\mathcal{M}_{\text{hybrid-ft}}$ from Stage I for preference optimization in Stage II. It is important to note that while higher pass@k scores suggest greater potential, the true challenge lies in effectively harnessing this potential, particularly in the weak-to-strong setting where no ground truths are available. Our proposed weak-to-strong preference optimization in Stage II successfully addresses this challenge, transforming theoretical potential into tangible performance gains in greedy decoding, as proved in §4.4.

A.4.3 PGR of Stage I

Burns et al. (2023) propose a new metric called performance gap recovered (PGR) to measure the fraction of the performance gap that can be recovered through weak supervision, as illustrated in Eq. 1. Tab. 11 displays the results of the naive full weak fine-tuning (i.e., Full Weak FT) and our best weak-icl fine-tuning (i.e., Weak-ICL FT) in terms of PGR, which also demonstrate that our method can outperform the simple competitor. However, the variations in PGR across different weak models do not provide meaningful insights. In the experiments described in the main text, we use test accuracy instead to provide a more detailed depiction of model performance.

| Weak Model | Full Weak FT | Weak-ICL FT |
|--------------|---------------|-------------|
| GSM8K | | |
| Llama2-7b | 22.47 | 78.53 |
| Gemma2-2b | 8.27 | 75.71 |
| Mistral-7b | 14.63 | 71.38 |
| MATH | | |
| Llama2-7b | 10.45 | 71.64 |
| Gemma-2b | -25.81 | 64.52 |
| Mistral-7b | 19.05 | 28.57 |

Table 11: Performance Gap Recovered (PGR) in Stage I.

| Weak Model | SFT Data | Test Accuracy |
|------------|-----------|----------------|
| Llama2-7b | Full Weak | 42.38 |
| | Gold Weak | 54.21 (+11.83) |
| | Our Weak | 53.68 (+11.30) |
| | Full ICL | 59.14 |
| | Gold ICL | 64.29 (+5.15) |
| | Our ICL | 61.71 (+2.57) |
| Gemma-2b | Full Weak | 29.04 |
| | Gold Weak | 46.40 (+17.36) |
| | Our Weak | 42.91 (+13.87) |
| | Full ICL | 58.61 |
| | Gold ICL | 63.86 (+5.25) |
| | Our ICL | 59.21 (+0.60) |
| Mistral-7b | Full Weak | 61.33 |
| | Gold Weak | 67.55 (+6.22) |
| | Our Weak | 65.96 (+4.63) |
| | Full ICL | 62.32 |
| | Gold ICL | 66.64 (+4.32) |
| | Our ICL | 65.43 (+3.11) |

Table 12: Detailed results of Stage I.

$$\text{PGR} = \frac{\text{weak-to-strong} - \text{weak floor}}{\text{strong ceiling} - \text{weak floor}}. \quad (1)$$

A.4.4 Effect of SFT Data

Tab. 12 presents more detailed comparative experimental results of Stage I on GSM8K. “Full Weak” denotes full weak fine-tuning, “Our Weak” is equivalent to $\mathcal{M}_{\text{weak-ft}}^1$, and “Our ICL” is equivalent to $\mathcal{M}_{\text{icl-ft}}^1$. “Gold Weak” refers to the scenario where weak data with correct final answers are filtered and used for supervised fine-tuning, which is impossible in the weak-to-strong setting and just used for experimental analysis. Similarly, “Gold ICL” refers to the scenario where solutions with correct final answers, generated by the strong model via weak ICL, are filtered.

Compared to using a large volume of noisy data (i.e., Full Weak and Full ICL), reducing the data quantity while enhancing data quality can significantly improve the accuracy of the trained model,

| Test Accuracy | |
|-------------------|-------|
| GSM8K | |
| Llama2-70b wo CoT | 12.36 |
| Llama2-70b w/ CoT | 18.35 |
| MATH | |
| Llama2-70b wo CoT | 6.40 |
| Llama2-70b w/ CoT | 7.20 |

Table 13: Zero-Shot Results of Llama2-70b on GSM8K and MATH.

with potential gains over 17 points. Although our method performs slightly lower than the gold results, it proves highly effective and stable in scenarios where obtaining the ground truth is impossible.

A.4.5 Zero-Shot Results

To obtain zero-shot performance, we follow [Kojima et al. \(2022\)](#) using a two-stage prompting approach. Specifically, we use the first prompt to extract a full reasoning path, where “wo CoT” denotes the standard prompt “Question: {question}\nAnswer:”, while “w/ CoT” denotes the CoT prompt “Question: {question}\nLet’s think step by step.\nAnswer:”. Then we use the second prompt, which concatenates “The answer is” with the generated reasoning path, to extract the answer in the correct format. The zero-shot results of Llama2-70b on the two reasoning datasets are presented in Tab. 13. We can observe that these results are significantly lower than those achieved with weak ICL. This notably poor zero-shot performance aligns with our hypothesis about the strong model: before any fine-tuning with weak supervision, the strong model’s capabilities have not been fully realized.

Question:
{question}

Student Solution:
{solution}

Your task involves three parts:

1. **Step-by-step Evaluation:** Go through the student solution carefully and identify key errors and potential misunderstandings that led to the incorrect solution.
2. **Final Judgement:** Provide an overall judgement on the correctness of the student's solution.
3. **First Error Step:** If the solution is incorrect, generate the step number where the first error occurs, otherwise generate N/A here.

Here's the format I want:

Step-by-step Evaluation: [Provide a step by step examination of the student solution and identify key errors and misunderstandings here.]

Final Judgement: [Insert only **correct** or **wrong** here]

First Error Step: [Insert either N/A or the step number where the first error occurs]

Please follow this format without any additional introductory or concluding statements.

Table 14: Prompt used to evaluate process-level accuracy.