

Debate, Reflect, and Distill: Multi-Agent Feedback with Tree-Structured Preference Optimization for Efficient Language Model Enhancement

Xiaofeng Zhou^{1*}, Heyan Huang¹, Lizi Liao²

¹Beijing Institute of Technology, ²Singapore Management University
{zhouxiaofeng, hhy63}@bit.edu.cn, lzliao@smu.edu.sg

Abstract

Large Language Models (LLMs) continue to set new standards in knowledge-intensive and complex reasoning tasks, yet their high computational demands limit widespread adoption. While distilling large models into smaller ones offers a sustainable solution, current techniques—such as static knowledge distillation, resource-intensive reinforcement learning from human feedback, or limited self-reflection—struggle to yield substantial and lasting performance gains. In this paper, we present a novel **Debate and Reflect (D&R)** framework that orchestrates multi-turn debates between smaller models and stronger teacher models, eliciting actionable feedback (*e.g.*, error analysis, corrective strategies) to guide student models. Further, we introduce Tree-structured Direct Preference Optimization (T-DPO) to efficiently leverage these debate logs, organizing interactions into a hierarchical format for effective training. Empirical evaluations across diverse NLP benchmarks demonstrate that our approach significantly improves smaller-model accuracy, robustness, and generalization, outperforming conventional baselines by a large margin.

1 Introduction

LLMs have emerged as powerful tools for tackling complex reasoning tasks (Zelikman et al., 2024; Qwen et al., 2025; He et al., 2024). Their ability to process and generate coherent, context-aware outputs has set a new standard in artificial intelligence (AI). However, the computational demands of training and deploying these models make them inaccessible to many users, highlighting the need for more efficient alternatives. Developing smaller models capable of achieving similar knowledge and reasoning performance is an essential step toward making advanced AI technologies more accessible and sustainable.

*This work was done during an internship at SMU.

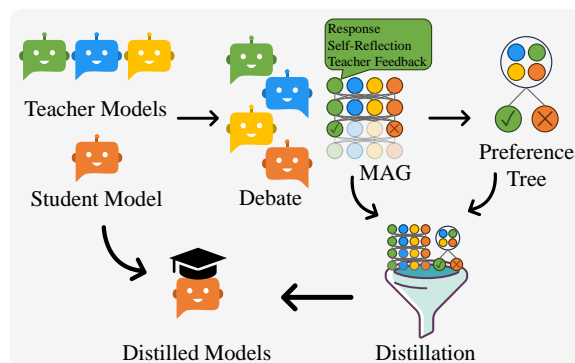


Figure 1: Overview of D&R. The student model debates with teacher models, receiving self-reflection and teacher feedback. A Multi-Agent Interaction Graph (MAG) records debate content, from which preference trees are extracted. All information is then distilled into the student, enhancing its knowledge and reasoning.

Existing strategies for enhancing smaller models generally fall into three categories. First, **knowledge distillation** transfers knowledge from a larger teacher to a smaller student model (Hinton et al., 2015; Shridhar et al., 2023). While effective for reducing model size, traditional distillation procedures are typically one-off processes that lack targeted guidance from the teacher. That is, the teacher’s rich insights into how the student can improve are not fully leveraged, and the student often receives only generic training signals. Second, **reinforcement learning from human feedback** (RLHF) aligns model behavior with human preferences and can substantially boost reasoning (Stiennon et al., 2020; Liang et al., 2024c). However, it requires extensive human intervention, raising concerns about scalability and cost. Third, **multi-agent systems and self-reflection** frameworks enable models to exchange critiques or engage in introspective refinement during inference (Liang et al., 2024b; Madaan et al., 2023; Wang et al., 2024a). Although these methods show promise in improving model reasoning, they often apply feedback only transiently—without permanently

updating model weights—or use simplistic fine-tuning schemes that do not fully exploit iterative feedback.

In this paper, we introduce a novel **Debate and Reflect (D&R)** framework that combines the strengths of multi-agent debate and actionable teacher feedback to robustly enhance smaller models. Concretely, a student model engages in multi-turn debates with one or more stronger teacher models, collecting detailed feedback in the form of error analyses, corrective strategies, and best practices. Crucially, the feedback is not limited to a single inference; rather, we incorporate both **self-reflection** and **teacher feedback** mechanisms that help the student to internalize these insights for lasting, parameter-level improvements (Yuan et al., 2024; Ying et al., 2024; Lan et al., 2024). To maximize the efficiency of this training signal, we further propose Tree-structured Direct Preference Optimization (T-DPO) to transform debate logs into hierarchical preference trees (Yuan et al., 2025). By systematically organizing intermediate debate turns and teacher guidance, T-DPO enables the student to learn not just what the best response is, but why certain answers are preferred. Through extensive empirical evaluations, we show that D&R with T-DPO significantly outperforms conventional methods in accuracy, robustness, and generalization, offering a promising route for achieving high performance in smaller, more accessible language models.

In summary, our contributions are threefold:

- We introduce a Debate and Reflect paradigm that combines the strengths of multi-agent interaction and targeted teacher feedback to continuously improve smaller models.
- We propose a Tree-structured Direct Preference Optimization (T-DPO) to better leverage debate data by structuring it into preference trees for efficient and robust training.
- Extensive experiments across diverse NLP tasks, demonstrating that our method significantly outperforms standard knowledge distillation and supervised fine-tuning.

2 Related Work

2.1 Knowledge Distillation

Knowledge distillation (Hinton et al., 2015) is a widely used technique for transferring knowledge

from large and high-performing teacher models to smaller student models, enhancing performance while reducing inference costs. Given the varying parameter sizes, capabilities, and computational requirements of LLMs (Touvron et al., 2023; Jiang et al., 2023; Li et al., 2024; Zhang et al., 2024b; Wang et al., 2023; Ren et al., 2022; Zhou et al., 2024; Bai et al., 2024), distilling knowledge from powerful teacher models is a common strategy for optimizing the trade-off between performance and efficiency. Typically, knowledge distillation involves querying teacher models with a set of seed data to extract relevant knowledge, which is then used to train student models through various distillation algorithms (Xu et al., 2024; Gu et al., 2024). In the context of reasoning tasks, Chain of Thought (CoT) reasoning serves as a crucial learning signal, allowing student models to learn structured reasoning steps from their teachers (Mukherjee et al., 2023). Beyond CoT, feedback has also been employed to provide process supervision, further refining student models’ reasoning capabilities (Luo et al., 2023). However, conventional knowledge distillation methods are typically static, meaning they lack iterative interactions between teacher and student models. To address this, our work utilizes CoT and feedback from debate-based interactions, providing more targeted and dynamic guidance for student model improvement.

2.2 Multi-Agent Debate

Multi-agent debate refines LLM reasoning and decision-making based on LLMs’ powerful conversational ability (Chu et al., 2024; Liang et al., 2024a; Yang et al., 2024) and agents’ collaborative ability (Li et al., 2023; Tu et al., 2024), by exploring diverse reasoning paths and cross-verifying claims, thereby reducing hallucinations and errors (Du et al., 2024). It also serves as an evaluation mechanism for tasks like question answering and summarization (Chan et al., 2024). Debate formats vary—from collaborative (Xiong et al., 2023) to adversarial (Cohen et al., 2023)—with different communication protocols (Chan et al., 2024). Recent studies show that such frameworks can even drive self-improvement without human supervision (Fu et al., 2023). For instance, MAGDI (Chen et al., 2024) uses multi-teacher debates to build interaction graphs that yield richer learning signals than single-teacher setups. In our work, we leverage multi-agent debates with stronger models to generate richer reasoning and actionable feedback,

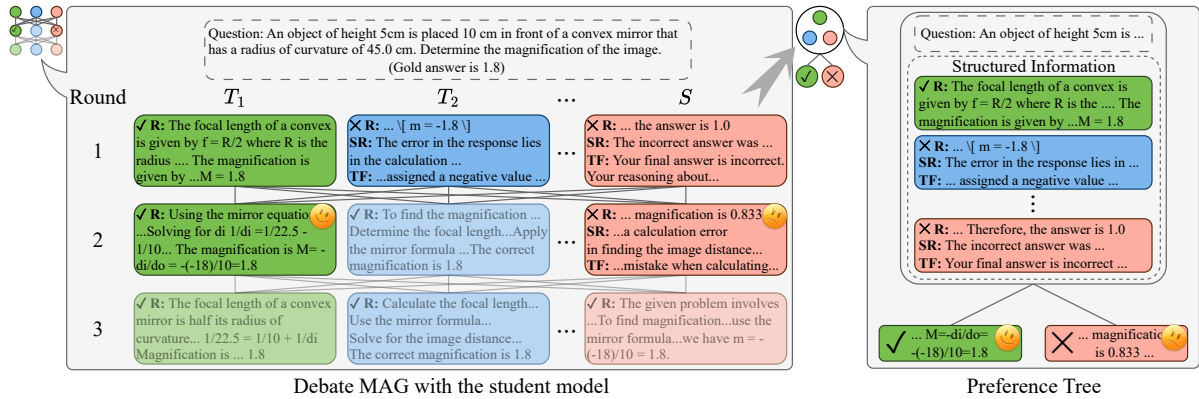


Figure 2: A debate example of our proposed D&R. The left part demonstrates the MAG of a debate involving the teacher models T_1, T_2, \dots and the student model S . Each node corresponds to an agent in a round, including responses (R), self-reflection (SR), and teacher feedback (TF). The highlighted nodes are extracted to construct a preference tree for T-DPO. In its root node, there is the structured information representing the previous debate.

integrating these insights into our Tree-structured Direct Preference Optimization pipeline.

2.3 DPO and Reasoning

Direct Preference Optimization (DPO) (Rafailov et al., 2023) offers a stable and computationally efficient way to align LLMs with specific objectives, making it valuable for refining both behavior and reasoning quality. Several extensions illustrate how DPO can be harnessed to improve reasoning. For instance, RPO (Pang et al., 2024) iteratively applies DPO alongside negative log-likelihood losses to fine-tune models on complex tasks. CPO (Zhang et al., 2024a) and Eurus (Yuan et al., 2025) take a tree-structured approach, building hierarchical reasoning paths to guide DPO-based alignment. Further refinements like Step-DPO (Lai et al., 2024) and Step-Controlled DPO (Lu et al., 2024) apply alignment at each individual reasoning step, ensuring that the model learns correct transitions between steps. In addition, M-DPO (Xiong et al., 2025) shows how DPO can be extended to multi-turn dialogues, incorporating tool usage for solving math problems. In our work, we extend DPO with a Tree-structured approach (T-DPO) that turns debate logs into hierarchical preference trees, enabling the student model to learn from both final answers and the intermediate reasoning in multi-turn debates.

3 Debate & Reflect Method

To improve small models using **Debate & Reflect (D&R)**, we construct debates where the small student model engages in discussions with their teacher models, receiving feedback from them (Section 3.1). From debates, we construct preference trees in Multi-Agent Interaction Graphs (MAGs)

including debate context for T-DPO (Section 3.2). Finally, we distill these structured insights into the student model with T-DPO, enabling more effective learning (Section 3.3).

3.1 Small Models Engage in Debates

To enable iterative interaction and targeted feedback, we design a debate framework where a student model engages in discussions with stronger teacher models. Figure 2 shows a debate example. We use training sets as our data sources to obtain gold answers. During the debates, multiple agents engage in several rounds of discussion on the same problem, exchanging their insights, and ultimately achieving a consensus answer. In each debate round, the teacher models $\mathcal{T} = \{T_i\}_{i=1}^n$ and the student model S analyze the structured information from the previous debate round. This information includes prior responses, along with the corresponding self-reflection and feedback. Based on this analysis, each agent generates a new response that includes both a reasoning process and a final answer. To evaluate performance, we compare each agent’s response against the gold answer. Agents who get an incorrect answer (such as the student model S) are prompted to self-reflect, while agents with correct answers (such as a teacher model T_i) provide constructive feedback to incorrect agents. Teacher models actively highlight student errors and suggest corrections, while the student model attempts to recognize and address its own mistakes. Then the newly generated responses, including reasoning processes, self-reflections, and feedback, are passed to the next round as updated structured information. In the next round, the models also analyze the strengths and weaknesses of their pre-

vious responses again, identifying potential errors and incorporating corrective feedback to generate next-round responses. This iterative process enables the student model to progressively refine its reasoning and responses. Debates continue until all agents reach a consensus or the maximum number of rounds is reached. Following [Chen et al. \(2024\)](#), we construct MAGs for debates, which record debate processes as structured graphs for distillation.

3.2 Construct Preference Trees for T-DPO

To effectively train models on reasoning trajectories, we construct preference trees as structured representations of debate interactions. Inspired by previous studies ([Yuan et al., 2025](#); [Zhang et al., 2024a](#)), we propose Tree-structured DPO (T-DPO) for debate distillation, which constructs preference trees from graphs for preference optimization.

The right hand side of [Figure 2](#) shows an example of our preference tree. Preference trees capture small, specific reasoning cases from debates. In the debates mentioned above, agents receive structured information from the previous round and generate new responses. Each agent is given the same input prompt in the same round, containing the problem and the previous structured information. However, due to varying capabilities, agents produce different responses; some are correct, and others are incorrect. Therefore, we choose an input prompt as the root node of a preference tree. A correct response and an incorrect response derived from this prompt become child nodes. The root node stores not only the previous round’s responses but also self-reflection and feedback on the incorrect answers. This allows agents to integrate insights from past correct responses while addressing mistakes through refinement. Responses that evolve into correct answers (chosen nodes) illustrate how effective reasoning develops over multiple debate rounds. Conversely, responses that remain incorrect (rejected nodes) highlight lessons that should not be reinforced in the distillation process. Through this hierarchical structure, preference trees systematically map out reasoning trajectories, helping models internalize both successful strategies and critical learning points from past errors.

3.3 Distillation

Our distillation pipeline consists of two key stages: Supervised Fine-Tuning (SFT) and Tree-structured DPO (T-DPO). First, we extract problems and their corresponding correct responses from MAGs to

construct the SFT dataset. Each data instance in this dataset includes an instruction I , a problem x_i , and its correct answer y_i , but excludes the structured information from the previous debate round. The student model is optimized with maximum likelihood estimation:

$$\mathcal{L}_{SFT} = -\mathbb{E}_{(I,x,y)\sim\mathcal{D}_{SFT}} \log \pi_{\theta}(y|I, x),$$

where \mathcal{D}_{SFT} is the SFT training dataset and the student model π_{θ} is initialized as the original small model before distillation.

In the second stage, we apply T-DPO training to refine the student model’s reasoning. The root node of each preference tree serves as the input prompt, incorporating the instruction I , the problem x , and the previous structured information SI . The correct and incorrect child nodes act as the chosen and rejected responses, y_w and y_l , respectively. The student model is optimized with the T-DPO objective:

$$\mathcal{L}_{T-DPO} = -\mathbb{E}_{(I,x,SI,y_w,y_l)\sim\mathcal{D}_{T-DPO}} \left[\log \sigma \left(\beta \log \frac{\pi_{\theta}(y_w|I, x, SI)}{\pi_{ref}(y_w|I, x, SI)} - \beta \log \frac{\pi_{\theta}(y_l|I, x, SI)}{\pi_{ref}(y_l|I, x, SI)} \right) \right],$$

where \mathcal{D}_{T-DPO} is all preference trees in the T-DPO training dataset, σ is the logistic function, π_{θ} is the student model policy, π_{ref} is the base reference policy, and β is the hyperparameter controlling the deviation from the base reference policy. Both π_{θ} and π_{ref} are initialized with the above SFT-trained student model.

After undergoing SFT and T-DPO, the student model acquires the ability to construct correct reasoning paths while effectively avoiding past errors. During inference, the model processes problems step by step, systematically generating thoughts and analyses, ultimately leading to final answers.

4 Experiments

4.1 Experimental Settings

This section outlines the major settings. More setting details in [Appendix A](#).

Datasets. The MMLU Pro benchmark ([Wang et al., 2024b](#)) is a multi-discipline language understanding and reasoning benchmark, extending the MMLU benchmark ([Hendrycks et al., 2021a](#)) with more challenging problems. The dataset is

partially derived from MMLU, with additional samples sourced from the Internet and other datasets. It not only evaluates the models’ knowledge but also emphasizes reasoning capabilities. Notably, it expands the number of answer choices to 10. For our experiments, we choose three categories of MMLU Pro: computer science, physics, and biology. These categories vary in data size and complexity, making them ideal for evaluating both knowledge retention and logical reasoning. Since MMLU Pro does not provide a dedicated training set, we split the original test set into separate training and test sets. Statistical information of datasets is shown in Table 5.

The MATH dataset (Hendrycks et al., 2021b) is a benchmark designed to assess step-by-step mathematical reasoning in LLMs. It requires LLMs to think about the problems step by step and generate a final answer, such as a mathematical expression. In our experiments, we adopt the training data in Chen et al. (2024) but reconstruct the debates and MAGs with our teacher and student models.

Baselines. To evaluate our method, we compare it against various distillation baselines. (1) *No Distillation*: the original Mistral-7B-Instruct model without any distillation. For MMLU Pro, it is tested in a zero-shot setting. For MATH, we provide only a few simple demonstrations from (Hendrycks et al., 2021b) to illustrate the expected output format. (2) *Single-Teacher Distillation*: The student model is distilled using SFT with knowledge from only one teacher model at a time. (3) *Multi-Teacher Distillation*: MAGDI (Chen et al., 2024) uses both correct and incorrect answers of multiple teacher models for distillation. But it does not involve the student interacting with teachers and lacks reflection or feedback mechanisms. Following the official code of MAGDI, we reimplement it with our debate data. For all baseline models and D&R, We select GPT-4o¹, Claude 3.5², and Gemini 1.5 Pro³ as our teacher models, while the student model is Mistral-7B-Instruct (Jiang et al., 2023).

Evaluation Metrics. For MMLU Pro, we use accuracy as the evaluation metric. For MATH, we follow Hendrycks et al. (2021b) to use exact match after normalizing answers to get accuracy.

¹<https://openai.com/index/hello-gpt-4o>

²<https://www.anthropic.com/claude/sonnet>

³<https://ai.google.dev/gemini-api/docs/models/gemini>

Implement Details. To ensure high-quality distillation while controlling computational costs, we impose the following settings.⁴ When constructing debates, we limit debates to a maximum of 4 rounds to balance cost-efficiency and data richness. When collecting root node data for preference trees, we control the length of structured information since it is impractical to train the student model with too long sequences. We aim to include as much structured information in root nodes as possible and place the remainder into a new preference tree. For the first round, the preference trees have no structured information from the previous round. During distillation, we implement our distillation with LoRA (Hu et al., 2022). The prompt templates used in our experiments are shown in Appendix F.

4.2 Main Results.

We show our experiment results in Table 1. The student model’s knowledge and reasoning capability across three disciplines, computer science, physics, and biology, are evaluated by MMLU Pro. The complex mathematic reasoning capability is evaluated by MATH.

D&R significantly enhances the student model through distillation. Before distillation, the original Mistral-7B-Instruct demonstrates moderate performance. On MMLU Pro, it surpasses the random guessing performance (10.00) by a meaningful margin, highlighting its basic knowledge and reasoning capability. However, on MATH, it achieves only 8.02, indicating that it can handle complex reasoning for only a small subset of MATH problems. All distillation methods lead to notable improvements over the original Mistral-7B-Instruct. D&R gets an average improvement of 14.18 (up to 18.41 in the biology category of MMLU Pro). This strong improvement underscores D&R’s ability to capture the valuable insights of debates and recognize the preference between the correct and incorrect answers.

D&R outperforms Single-Teacher Distillation. As baseline models, single-teacher distillation results vary across different teacher models. Among them, Single Teacher_{Claude 3.5} achieves the best average results of single-teacher distillation, while Single Teacher_{Gemini 1.5 Pro} lags behind, with an average performance gap of 3.38. Despite these variations, all single-teacher distillation methods

⁴We release our code in <https://github.com/zhouxiaofengshelf/D-R>.

Models	MMLU Pro			MATH	Average
	Computer Science (CS)	Physics	Biology		
<i>No Distillation</i>					
Mistral-7B-Instruct	20.49	18.46	48.95	8.02	23.98
<i>Single-Teacher Distillation</i>					
Single Teacher _{GPT-4o}	27.80	28.63	62.34	16.92	33.92
Single Teacher _{Claude 3.5}	30.73	30.77	62.76	16.56	35.21
Single Teacher _{Gemini 1.5 Pro}	25.85	24.92	60.67	15.86	31.83
<i>Multi-Teacher Distillation</i>					
MAGDi	22.93	20.00	54.39	11.26	27.15
D&R _{SFT}	32.68	29.23	64.44	17.64	36.00
D&R	33.17	34.77	67.36	17.32	38.16

Table 1: The main results of our experiments on three categories of MMLU Pro and MATH. We compare our D&R with three types of distillation baselines and SFT-only D&R_{SFT}. MAGDi are reimplemented with its official code and our debate data. **Bold** figures denote the best results.

lead to meaningful improvements, proving that knowledge and reasoning can be effectively transferred from teacher models.

Compared to the single-teacher distillation, D&R achieves better distillation results. On the average performance, D&R improves at least 2.95 (from 35.21 of Single Teacher_{Claude 3.5} to 38.16 of D&R). The key advantages of D&R include multiple teacher models for richer insights, debate mechanisms for iterative interactions, and T-DPO for structured reasoning preference learning. These factors collectively enable more effective knowledge and reasoning acquisition. The debate process enhances training data quality, focusing on the student model’s weaknesses, while T-DPO refines learning by emphasizing the contrast between correct and incorrect responses.

D&R achieves the best performance in Multi-Teacher Distillation. D&R_{SFT} already achieves strong performance, surpassing most baseline models. T-DPO further enhances it, leading to an additional 2.16 improvement in the average performance. On the MATH dataset, the result of D&R is slightly lower than the one of D&R_{SFT} but still higher than all other baseline models. We will conduct further exploration in Section 4.3.4. Compared to MAGDi, our D&R demonstrates a substantial average gain of 11.01 (up to 14.77 in the physics category of MMLU Pro). While both D&R and MAGDi utilize multiple teachers and debates, D&R generates richer debate content. Involving the student model in the debates can lead to higher-quality training data. The teacher models are aware of errors in the student model and can actively provide explicit feedback. The student model

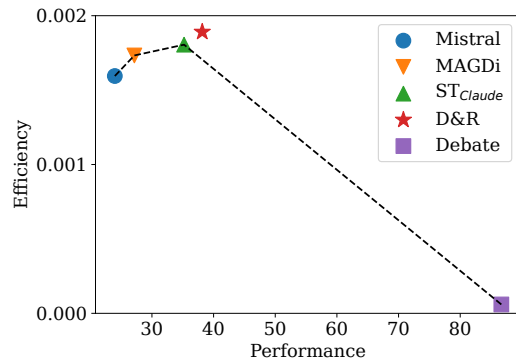


Figure 3: Trade-off between performance and efficiency. Inspired by Chen et al. (2024), we define efficiency as $1/\text{avg}(\text{tokens})$. D&R improves both performance and efficiency compared to previous distillation methods, surpassing their Pareto frontier. We choose Single Teacher_{Claude 3.5} (ST_{Claude}) to represent the Single-Teacher Distillation.

learns not only from teacher responses but also from the feedback and their own self-reflection to know where and why it makes errors. This process makes D&R’s distillation approach superior to passively learning from teacher responses without explicit error analysis, as done in MAGDi. Moreover, T-DPO provides a more structured and effective optimization framework than the margin-based objective (Cortes and Vapnik, 1995) and GCN (Kipf and Welling, 2017) of MAGDi. MAGDi underperforms compared to the single-teacher distillation, suggesting its joint learning of three training objectives may not model the debate data effectively.

4.3 Detailed Analysis

4.3.1 D&R Improves Inference Efficiency

D&R effectively balance efficiency and performance. Although LLMs and debate can provide

Models	Additional Correction	Accuracy
D&R	None	34.77
D&R		36.92
without SR data	Self-Reflect	33.85 (-3.07)
without TF data		32.31 (-4.61)
without SR & TF data		31.38 (-5.54)
D&R		37.54
without SR data	Teacher*	32.62 (-4.92)
without TF data	Feedback	32.00 (-5.54)
without SR & TF data		31.08 (-6.46)

* As during inference, we do not allow teacher models to involve, we let the student model role-play a teacher.

Table 2: Effect of self-reflection (SR) and teacher feedback (TF) in D&R. We also show D&R with additional self-reflection and teacher feedback during inference.

strong knowledge and reasoning, they come at a high inference cost. Distillation improves efficiency, making reasoning models more accessible and sustainable. Inspired by Chen et al. (2024), we measure inference efficiency by calculating token cost per problem across three categories in MMLU Pro (excluding MATH due to API cost). We report our detailed token cost in Appendix C. Figure 3 illustrates the trade-off between performance and efficiency. Debate incurs the highest token cost, making it the least efficient. After distillation, MAGDI and Single Teacher_{Claude 3.5} can generate better reasoning paths, improving their performance with high efficiency. D&R surpasses the boundaries of previous distillation methods, not only improving performance but also effectively enhancing efficiency. It becomes a new optimal balance of performance and inference efficiency.

4.3.2 D&R Can Correct Itself

D&R not only improves reasoning but also learns to identify and correct its own mistakes. In our debates, agents are explicitly required to generate self-reflection and feedback to identify and correct errors. During distillation, the student model learns this correction process. To evaluate its error correction ability, we conduct an additional correction step during inference. In correspondence with the self-reflection and feedback in the training debates, we design two additional correction methods. The first one is self-reflection. After generating an initial response, the student model self-reflects on its own response and updates its answer. The second one is teacher feedback simulation. The student model roleplays a teacher model, critiques its own response as if reviewing another agent’s answer, and then corrects errors (both the teacher and stu-

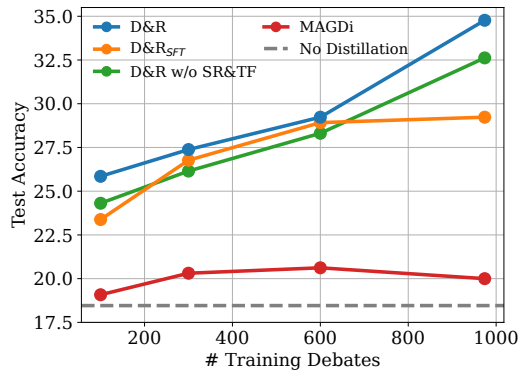


Figure 4: Performance with different training debates on MMLU Pro physics. The dashed line refers to the original Mistral-7B-Instruct without any distillation.

dent models are actually the same model). Table 2 shows how self-reflection and teacher feedback in training debates improve the student model’s performance. With these additional correction methods, D&R can leverage its learned error correction ability to refine its reasoning. However, without the self-reflection and feedback data from the training debates for distillation, the student model fails to effectively self-correct or simulate teacher feedback during inference. This highlights that beyond reasoning, self-reflection and teacher feedback in training debates are essential for transferring critical correction abilities to the student model, enhancing its performance through additional correction at inference.

4.3.3 D&R Works on Various Data Scales

To further investigate D&R’s adaptability to different training data size, we train D&R on the physics category of MMLU Pro using different debate scales: 100, 300, 600, and 974 debates (full set). The results, reported in Figure 4, reveal the following key findings.

At first, our D&R achieves effective distillation learning across all dataset sizes, while MAGDI struggles with smaller datasets and fails to distill knowledge and reasoning capabilities effectively. Then D&R outperforms MAGDI with all data scales. And T-DPO always enhances the SFT models, proving that self-reflection and teacher feedback in T-DPO are crucial. Without them, performance declines. The performance of D&R also improves as the data scale increases. Therefore, D&R has better utilization of debate data and still can capture preference with limited data. It can collect multiple SFT data instances and preference trees from a single debate MAG, improving data efficiency. The structured information in our pref-

Distillation Data	Student Models	MATH
The Same Training Debates for D&R	Mistral-7B-Instruct	8.02
	with D&R _{SFT}	17.64 (+9.62)
	with D&R	17.32 (+9.30)
	Llama-3.1-8B-Instruct	45.78
	with D&R _{SFT}	47.06 (+1.28)
	with D&R	48.02 (+2.24)

Table 3: Comparison between different student models. The performance of distillation can be affected by the capability of student models. Llama-3.1-8B-Instruct is able to learn more during D&R distillation.

erence trees provides rich knowledge, while self-reflection and feedback data can help the student model deeply understand complex reasoning preferences. Hence D&R can be applied in scenarios with limited data resources, achieving effective improvements with only a small number of debates. When given access to more data, it scales effectively, leading to further performance gains.

4.3.4 D&R with Different Student Models

D&R achieves improvements across student models of varying capabilities. The effectiveness of distillation depends not only on teacher models and algorithms but also on the capability of our student model. MATH is a challenging dataset that requires complex reasoning. And the basic performance of Mistral-7B-Instruct on MATH is a little weak. As shown in Table 1, while D&R_{SFT} achieves a strong distillation performance on the MATH dataset, D&R does not provide further improvements. To investigate whether this limitation stems from the student model’s capacity, we distill the same training debates into a more powerful student model, Llama-3.1-8B-Instruct. The results, presented in Table 3, confirm this hypothesis. On the MATH dataset, where Mistral may reaches a distillation bottleneck, Llama-3.1-8B-Instruct continues to improve with both SFT and T-DPO in D&R. The complex reasoning capacity of Mistral-7B-Instruct might not be sufficient for full distillation. It can absorb teacher models’ reasoning capabilities during SFT but reaches its capacity limit, preventing further gains from T-DPO. On the contrary, Llama-3.1-8B-Instruct is stronger and has a much better basic MATH performance than Mistral-7B-Instruct. It has much potential for distillation. Therefore, during the T-DPO, it can still learn more complex reasoning capabilities from our training debates, demonstrating that D&R can work well on the MATH dataset and fit different student

Variants	MMLU Pro			MATH	Avg.
	CS	Phys.	Biol.		
w/ <i>SFT</i>	32.68	29.23	64.44	17.64	36.00
w/ <i>RPO</i>	36.59	32.00	66.95	17.00	38.14
w/ <i>T-DPO</i>	33.17	34.77	67.36	17.32	38.16

Table 4: Comparison among D&R variants. Although w/ *RPO* is better than w/ *SFT*, it generally performs worse than w/ *T-DPO* version.

models. That proves D&R’s robust effectiveness across diverse student models.

4.3.5 The Influence of Different Objectives

Our proposed T-DPO objective is well-suited for distilling reasoning abilities from debates. In reasoning tasks, various preference-learning objectives exist (Xiong et al., 2025; Zhang et al., 2024a). In this section, we investigate whether T-DPO is sufficient and if other training objectives could replace T-DPO in D&R. RPO (Pang et al., 2024) is a DPO-based training objective designed for reasoning tasks. Previous experiments (Pang et al., 2024) suggest that the sequence-level log probabilities for chosen completions with RPO are increasing, while the standard DPO suffers a decrease.

To compare these objectives, we apply RPO to our preference trees and report the results in Table 4. Following Pang et al. (2024), we set the hyperparameter α in the RPO loss as 1. We observe that D&R_{RPO} performs better than D&R_{SFT} and close to D&R_{T-DPO} in average performance. However, according to our experimental logs, the log probability does not show a consistent trend across different data categories. In some categories, such as physics, RPO increases log probabilities beyond T-DPO, yet performance does not improve. In others like computer science, log probabilities show little difference from T-DPO, yet RPO performs well. There is no clear correlation between log probability trends and model performance, suggesting that RPO’s effectiveness may be task-dependent. In conclusion, based on preference trees from debate graphs, T-DPO remains the most suitable objective for debate distillation.

5 Conclusion

In this paper, we proposed D&R, a novel distillation framework to enhance small models through debates. The debate process, involving teacher and student models, fully exposed the student model’s weaknesses, providing self-reflection and feedback. Then, we used T-DPO with preference trees from

debate MAGs, capturing structured information and preference relationships between different reasoning paths. D&R effectively distilled the teacher models' knowledge and reasoning capabilities into the student model through SFT and T-DPO. Our experiments demonstrate that D&R achieves superior distillation compared to baseline methods, bridging the gap between smaller student models and larger teacher models. Compared to online debate everytime, D&R significantly improved inference efficiency. Beyond reasoning, it also strengthened the student model's ability to self-correct. In future, we will explore long-term learning where the student model continues to improve over time through continuous debates and iterative feedback.

Acknowledgements

This research was supported by the National Research Foundation, Singapore under its National Large Language Models Funding Initiative (AISG Award No: AISG-NMLP-2024-002). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore. It was also supported by the National Natural Science Foundation of China (No. U21B2009) and the Lee Kong Chian Fellowship. This work was done during a visit to Singapore Management University funded by the China Scholarship Council.

Limitations

Our study has limitations in two aspects. First, while multi-agent debate and preference learning are explored across a wide range of tasks, such as evaluation and model alignment, our work focuses solely on the knowledge and reasoning tasks. Investigating D&R in other tasks would be valuable for assessing its generalizability. Second, our evaluation primarily measures the correctness of final answers rather than the validity of reasoning paths, which may lead to some false positive results. Process verification remains an open challenge, and we anticipate future studies that incorporate comprehensive evaluations of both reasoning paths and final outcomes.

References

Yu Bai, Xiyuan Zou, Heyan Huang, Sanxing Chen, Marc-Antoine Rondeau, Yang Gao, and Jackie CK

Cheung. 2024. [CITruS: Chunked Instruction-aware State Eviction for Long Sequence Modeling](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5908–5930, Miami, Florida, USA. Association for Computational Linguistics.

Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2024. [ChatEval: Towards Better LLM-based Evaluators through Multi-Agent Debate](#). In *The Twelfth International Conference on Learning Representations*.

Justin Chen, Swarnadeep Saha, Elias Stengel-Eskin, and Mohit Bansal. 2024. [MAGDi: Structured Distillation of Multi-Agent Interaction Graphs Improves Reasoning in Smaller Language Models](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 7220–7235. PMLR.

Yuqi Chu, Lizi Liao, Zhiyuan Zhou, Chong-Wah Ngo, and Richang Hong. 2024. [Towards Multimodal Emotional Support Conversation Systems](#). *arXiv preprint*.

Roi Cohen, May Hamri, Mor Geva, and Amir Globerson. 2023. [LM vs LM: Detecting Factual Errors via Cross Examination](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12621–12640, Singapore. Association for Computational Linguistics.

Corinna Cortes and Vladimir Vapnik. 1995. [Support-vector networks](#). *Machine Learning*, 20(3):273–297.

Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. 2024. [Improving Factuality and Reasoning in Language Models through Multiagent Debate](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 11733–11763. PMLR.

Yao Fu, Hao Peng, Tushar Khot, and Mirella Lapata. 2023. [Improving Language Model Negotiation with Self-Play and In-Context Learning from AI Feedback](#). *arXiv preprint*.

Yuxian Gu, Li Dong, Furu Wei, and Minlie Huang. 2024. [MiniLLM: Knowledge Distillation of Large Language Models](#). In *The Twelfth International Conference on Learning Representations*.

Tao He, Lizi Liao, Yixin Cao, Yuanxing Liu, Ming Liu, Zerui Chen, and Bing Qin. 2024. [Planning Like Human: A Dual-process Framework for Dialogue Planning](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4768–4791, Bangkok, Thailand. Association for Computational Linguistics.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021a. [Measuring Massive Multitask Language Understanding](#). In *International Conference on Learning Representations*.

- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021b. [Measuring Mathematical Problem Solving With the MATH Dataset](#). In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1.
- Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. [Distilling the Knowledge in a Neural Network](#). *CoRR*, abs/1503.02531.
- 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.
- 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](#). *arXiv preprint*.
- Thomas N. Kipf and Max Welling. 2017. [Semi-Supervised Classification with Graph Convolutional Networks](#). In *International Conference on Learning Representations*.
- Xin Lai, Zhuotao Tian, Yukang Chen, Senqiao Yang, Xiangru Peng, and Jiaya Jia. 2024. [Step-DPO: Step-wise Preference Optimization for Long-chain Reasoning of LLMs](#). *arXiv preprint*.
- Tian Lan, Wenwei Zhang, Chengqi Lyu, Shuaibin Li, Chen Xu, Heyan Huang, Dahua Lin, Xian-Ling Mao, and Kai Chen. 2024. [Training Language Models to Critique With Multi-agent Feedback](#). *arXiv preprint*.
- Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023. [CAMEL: Communicative Agents for "Mind" Exploration of Large Language Model Society](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 51991–52008. Curran Associates, Inc.
- Jiawei Li, Yizhe Yang, Yu Bai, Xiaofeng Zhou, Yinghao Li, Huashan Sun, Yuhang Liu, Xingpeng Si, Yuhao Ye, Yixiao Wu, Yiguan Lin, Bin Xu, Ren Bowen, Chong Feng, Yang Gao, and Heyan Huang. 2024. [Fundamental Capabilities of Large Language Models and their Applications in Domain Scenarios: A Survey](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11116–11141, Bangkok, Thailand. Association for Computational Linguistics.
- Jinggui Liang, Yuxia Wu, Yuan Fang, Hao Fei, and Lizi Liao. 2024a. [A Survey of Ontology Expansion for Conversational Understanding](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 18111–18127, Miami, Florida, USA. Association for Computational Linguistics.
- Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. 2024b. [Encouraging Divergent Thinking in Large Language Models through Multi-Agent Debate](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17889–17904, Miami, Florida, USA. Association for Computational Linguistics.
- Xinyue Liang, Jiawei Li, Yizhe Yang, and Yang Gao. 2024c. [Bit_numeval at SemEval-2024 Task 7: Enhance Numerical Sensitivity and Reasoning Completeness for Quantitative Understanding](#). In *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)*, pages 1830–1841, Mexico City, Mexico. Association for Computational Linguistics.
- Zimu Lu, Aojun Zhou, Ke Wang, Houxing Ren, Weikang Shi, Junting Pan, Mingjie Zhan, and Hongsheng Li. 2024. [Step-Controlled DPO: Leveraging Stepwise Error for Enhanced Mathematical Reasoning](#). *arXiv preprint*.
- Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng, Qingwei Lin, Shifeng Chen, Yansong Tang, and Dongmei Zhang. 2023. [WizardMath: Empowering Mathematical Reasoning for Large Language Models via Reinforced Evol-Instruct](#). *arXiv preprint*.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. 2023. [Self-Refine: Iterative Refinement with Self-Feedback](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 46534–46594. Curran Associates, Inc.
- Subhabrata Mukherjee, Arindam Mitra, Ganesh Jawahar, Sahaj Agarwal, Hamid Palangi, and Ahmed Awadallah. 2023. [Orca: Progressive Learning from Complex Explanation Traces of GPT-4](#). *arXiv preprint*.
- Richard Yuanzhe Pang, Weizhe Yuan, He He, Kyunghyun Cho, Sainbayar Sukhbaatar, and Jason E. Weston. 2024. [Iterative Reasoning Preference Optimization](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang

- Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2025. [Qwen2.5 Technical Report](#). *arXiv preprint*.
- 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*, volume 36, pages 53728–53741. Curran Associates, Inc.
- Mucheng Ren, Heyan Huang, Yuxiang Zhou, Qianwen Cao, Yuan Bu, and Yang Gao. 2022. TCM-SD: A Benchmark for Probing Syndrome Differentiation via Natural Language Processing. In *Chinese Computational Linguistics*, pages 247–263, Cham. Springer International Publishing.
- Kumar Shridhar, Alessandro Stolfo, and Mrinmaya Sachan. 2023. [Distilling Reasoning Capabilities into Smaller Language Models](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 7059–7073, Toronto, Canada. Association for Computational Linguistics.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. [Learning to summarize with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 3008–3021. Curran Associates, Inc.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti 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, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open Foundation and Fine-Tuned Chat Models](#). *arXiv preprint*.
- Rong-Cheng Tu, Wenhao Sun, Zhao Jin, Jingyi Liao, Jiaying Huang, and Dacheng Tao. 2024. [SPAgent: Adaptive Task Decomposition and Model Selection for General Video Generation and Editing](#). *arXiv preprint*.
- Bo Wang, Heyan Huang, Yixin Cao, Jiahao Ying, Wei Tang, and Chong Feng. 2024a. [QRMeM: Unleash the Length Limitation through Question then Reflection Memory Mechanism](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4837–4851, Miami, Florida, USA. Association for Computational Linguistics.
- Bo Wang, Heyan Huang, Xiaochi Wei, Ge Shi, Xiao Liu, Chong Feng, Tong Zhou, Shuaiqiang Wang, and Dawei Yin. 2023. [Boosting Event Extraction with Denoised Structure-to-Text Augmentation](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 11267–11281, Toronto, Canada. Association for Computational Linguistics.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhui Chen. 2024b. [MMLU-Pro: A More Robust and Challenging Multi-Task Language Understanding Benchmark](#). In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Kai Xiong, Xiao Ding, Yixin Cao, Ting Liu, and Bing Qin. 2023. [Examining Inter-Consistency of Large Language Models Collaboration: An In-depth Analysis via Debate](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 7572–7590, Singapore. Association for Computational Linguistics.
- Wei Xiong, Chengshuai Shi, Jiaming Shen, Aviv Rosenberg, Zhen Qin, Daniele Calandriello, Misha Khalman, Rishabh Joshi, Bilal Piot, Mohammad Saleh, Chi Jin, Tong Zhang, and Tianqi Liu. 2025. [Building Math Agents with Multi-Turn Iterative Preference Learning](#). In *The Thirteenth International Conference on Learning Representations*.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, Qingwei Lin, and Daxin Jiang. 2024. [WizardLM: Empowering Large Pre-Trained Language Models to Follow Complex Instructions](#). In *The Twelfth International Conference on Learning Representations*.
- Yizhe Yang, Heyan Huang, Yang Gao, and Jiawei Li. 2024. [Building knowledge-grounded dialogue systems with graph-based semantic modelling](#). *Knowledge-Based Systems*, 298:111943.
- Jiahao Ying, Mingbao Lin, Yixin Cao, Wei Tang, Bo Wang, Qianru Sun, Xuanjing Huang, and Shuicheng Yan. 2024. [LLMs-as-Instructors: Learning from Errors Toward Automating Model Improvement](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 11185–11208, Miami, Florida, USA. Association for Computational Linguistics.
- Lifan Yuan, Ganqu Cui, Hanbin Wang, Ning Ding, Xingyao Wang, Boji Shan, Zeyuan Liu, Jia Deng, Huimin Chen, Ruobing Xie, Yankai Lin, Zhenghao Liu, Bowen Zhou, Hao Peng, Zhiyuan Liu, and

- Maosong Sun. 2025. [Advancing LLM Reasoning Generalists with Preference Trees](#). In *The Thirteenth International Conference on Learning Representations*.
- Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Xian Li, Sainbayar Sukhbaatar, Jing Xu, and Jason E Weston. 2024. [Self-Rewarding Language Models](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 57905–57923. PMLR.
- Eric Zelikman, Georges Raif Harik, Yijia Shao, Varuna Jayasiri, Nick Haber, and Noah Goodman. 2024. [Quiet-STaR: Language Models Can Teach Themselves to Think Before Speaking](#). In *First Conference on Language Modeling*.
- Xuan Zhang, Chao Du, Tianyu Pang, Qian Liu, Wei Gao, and Min Lin. 2024a. [Chain of Preference Optimization: Improving Chain-of-Thought Reasoning in LLMs](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Zhihan Zhang, Yixin Cao, Chenchen Ye, Yunshan Ma, Lizi Liao, and Tat-Seng Chua. 2024b. [Analyzing Temporal Complex Events with Large Language Models? A Benchmark towards Temporal, Long Context Understanding](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1588–1606, Bangkok, Thailand. Association for Computational Linguistics.
- Xiaofeng Zhou, Heyan Huang, Zewen Chi, Mucheng Ren, and Yang Gao. 2024. [RS-BERT: Pre-training radical enhanced sense embedding for Chinese word sense disambiguation](#). *Information Processing & Management*, 61(4):103740.

A Experimental Setting Details

Datasets The data scale varies significantly across different categories in MMLU Pro. To ensure a sufficient amount of test data, we adopt different training-to-test split ratios: 1:1 for computer science, 3:1 for physics, and 2:1 for biology. When structuring information in the root nodes of preference trees, we filter out failed responses where the content is missing.

Table 5 presents the detailed statistics of our dataset. Most of our debates can proceed to the second round, resulting in meaningful interactions between the agents. In contrast, prior debate datasets contain a large proportion of debates that conclude after the initial round without further interaction. For the statistics of the previous debate data, please refer to Chen et al. (2024). With self-reflection and teacher feedback, agents in our framework can promptly correct errors after the second round, ensuring more consensual correct answers and reducing instances where debates terminate without reaching a valid solution. Furthermore, the number of SFT and T-DPO training instances generated from our debates is 2.5–7.7 times the number of debates, highlighting D&R’s efficient utilization of debate data for distillation.

Models To ensure the reproducibility of our work, we use the following versions of the teacher and student models: gpt-4o-2024-08-06 for GPT-4o, claude-3-5-sonnet-20241022 for Claude 3.5, gemini-1.5-pro-002 for Gemini 1.5 Pro, and Mistral-7B-Instruct-v0.2 for Mistral-7B-Instruct.

Implement Details For debate generation, we use the following generation configurations: temperature is 0.3, the hard limit max tokens is 700 (1000 for MATH and we also specify a soft limit in the prompt templates). For SFT, we train the student model on all datasets for 2 epochs with learning rate as $2e-4$ and batch size as 16.

For T-DPO, we train the student model on MMLU Pro for 3 epochs with learning rate as $5e-6$ and batch size as 16. On MATH, we train the student model for 1 epoch with learning rate as $1e-6$ and batch size as 16. We limit the total sequence length to 2000 and the prompt sequence length to 1400. For LoRA, we set the LoRA adapter rank as 16 and the alpha as 32.

For evaluation, our generation configurations are temperature as 0.3, top_p as 0.9, and the hard

limit max tokens as 700 (1000 for MATH). For the reimplementation of MAGDI, we follow its original training hyperparameters and generation configurations.

B Quality Analysis of Debate Data

Our debate data exhibits higher quality than traditional debate data, providing clearer reasoning structures and more targeted feedback. In D&R, we introduce an improved debate data generation method that actively involves the student model and explicitly incorporates self-reflection and feedback from teacher models. As a result, our training debate data is enriched with detailed reasoning and targeted guidance. To illustrate the advantages of our approach, Figure 10 presents a comparison between our debate data and traditional debate data from MATH.

For the problem of finding the greatest common divisor, the teacher model in our debate data clearly outlines the use of the Euclidean algorithm, systematically performs the calculations, and verifies the result. In contrast, the debate data from Chen et al. (2024), constructed using a traditional approach, lacks an explicit mention of the prime factorization method and presents a less structured response.

In our debate, the student model initially produces an incorrect response. While it introduces the Euclidean algorithm, it actually applies prime factorization, making multiple calculation errors in the process. However, the inclusion of self-reflection and teacher feedback highlights these mistakes and provides explicit recommendations for correctly applying the Euclidean algorithm. This guidance plays a crucial role in shaping subsequent debate rounds. Traditional multi-agent debate frameworks provide agents only with prior responses, requiring them to independently compare answers and identify correct and incorrect elements. This process adds cognitive complexity, making it harder for the agents to reason effectively. In contrast, our method streamlines error correction by integrating self-reflection and feedback, offering clear suggestions for improvement. Consequently, in the following debate rounds, agents can more efficiently rectify previous mistakes and generate more accurate responses, leading to a more effective learning process.

C Token Cost

We report averaged per-problem token cost on the test sets of three categories of MMLU Pro in Table

Datasets	# Training Debates	Round (1/2/3/4)	# R	# SR	# TF	# SFT	# T-DPO	# Test Problems
MATH	1000	85/606/149/160	9524	2579	4111	6957	7744	5000
Physics	974	136/696/104/38	7968	1722	3016	6246	4238	325
Biology	478	151/295/24/8	3380	553	1061	2827	1214	239
Computer Science	205	32/151/17/5	1620	340	554	1280	670	205

Table 5: Detailed statistics of our datasets. Round (1,2,3,4) shows the number of debates with 1/2/3/4 rounds. R, SR, and TF represent Responses, Self-Reflection, and Teacher Feedback, respectively. # SFT is the number of SFT data constructed from debates. # T-DPO is the number of preference trees constructed from debates for T-DPO.

Models	MMLU Pro			Avg.
	CS	Physics	Biology	
Mistral	601.87	692.51	586.92	627.10
MAGDI	578.05	619.64	533.54	577.08
ST _{Claude}	555.38	550.27	556.16	553.94
D&R	520.24	584.17	482.07	528.83
Debate	17993.04	21789.26	10526.94	16769.75

Table 6: Token cost on the three categories of MMLU Pro. For each category, we report the averaged per-problem token cost on test sets.

Models	MMLU Pro		
	CS	Physics	Biology
GPT-4o	75.12	74.77	88.70
Claude 3.5	77.07	77.54	87.87
Gemini 1.5 Pro	78.05	76.31	84.94
Debate	84.29	85.23	90.38
Llama-3.1-8B-Instruct	46.83	39.38	65.27
with D&R _{SFT}	50.24	44.92	71.97
with D&R	52.68	46.46	74.06

Table 7: Results of teacher models, debates, and a more powerful student model, Llama-3.1-8B-Instruct, on the three categories of MMLU Pro.

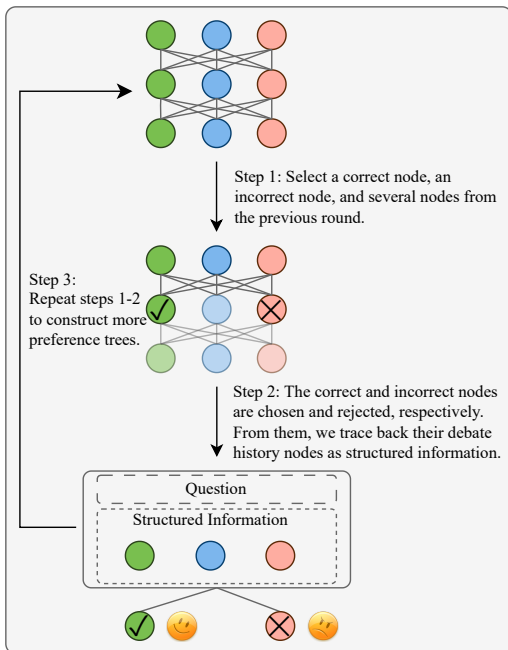


Figure 5: The schematic diagram showing how to construct preference trees from debates.

6, which is used to calculate efficiency in Section 4.3.1. D&R learns to reach answers with better reasoning paths and reduce failures during generation, hence has a better efficiency than Mistral, MAGDI, and ST_{Claude}. Since knowledge and reasoning capabilities have been distilled into the student model, there is no need to use multiple large models for debate. The token cost of D&R is much lower than that of debate.

D More Results on MMLU Pro

In Table 7, we report more results on the three categories of MMLU Pro, including three teacher models, debates, and Llama-3.1-8B-Instruct. The teacher models and debates achieve strong performance, ensuring the quality of distillation data. Besides MATH (as shown in Section 4.3.4), Llama-3.1-8B-Instruct also gain improvements on MMLU Pro. It demonstrates the broad applicability and robustness of D&R.

E The Detailed Schematic Diagram for Constructing Preference Trees

In Figure 5, we present a schematic diagram to complement Figure 2, illustrating the steps to construct preference trees:

1. Select a correct node, an incorrect node, and several nodes from the previous round.
2. The correct and incorrect nodes are chosen and rejected, respectively. From them, we trace back their debate history nodes as structured information.
3. Repeat steps 1-2 to construct more preference trees.

F Designed Prompt Templates in D&R

In this section, we provide the prompt templates that are used in our D&R for generating debates and distillation in Figure 6, 7, 8, and 9.

```

### Instruction
You are a helpful assistant aiming to answer the following question with reasoning.
You are debating the input question with other agents. You are provided with all answers from the last round, along with
their self-reflections and the feedback received. Please carefully consider these answers and output a new answer.
Please output in the following format:
...

### Answer
<your step-by-step reasoning>
Final answer: <your final answer>
...

<choice prompt> Please keep your output concise and limit it to 500 words. Use the most short reasoning to get the
correct answer. Do not output unrelated content! Do not explicitly mention other agents; instead, internally consider their
answers.

### Question
<question>

### Last Round Answers
<structured information from last round>

```

Figure 6: The prompt template for generating debates.

```

### Instruction
You are a helpful assistant aiming to provide a feedback to yourself.
Unfortunately, your previous answer to the question was incorrect. Given the correct final answer, please reflect on it and
provide feedback on your incorrect response. You need to identify the potential incorrect knowledge, faulty reasoning or
format errors in your incorrect response. Please check if your incorrect final answer has any formatting errors compared
to the correct answer.
Please output in the following format:
...

### Feedback
<your feedback>
...

You don't need to apologize.
Only output your feedback. Please keep your output concise and limit it to 500 words. Provide your most concise
feedback. Do not output a new answer.

### Question
<question>

### Correct Final Answer
<gold answer>

### Your Incorrect Answer
<incorrect reasoning>
Final answer: <incorrect answer>

```

Figure 7: The prompt template for self-reflecting in debates.

```

### Instruction
You are a helpful assistant aiming to provide a feedback to other agents.
Congratulation! You answered the question correctly. Based on your correct answer, please help another agent who
answered incorrectly. You should provide feedback by pointing out the potential incorrect knowledge, faulty reasoning or
format errors in its response. Please check if its incorrect final answer has any formatting errors compared to the correct
answer.
Please output in the following format:
...

### Feedback
<your feedback>
...

Only output your feedback. Please keep your output concise and limit it to 500 words. Provide your most concise
feedback. Do not output a new answer.

### Question
<question>

### Your Correct Answer
<correct reasoning>
Final answer: <correct answer>

### Incorrect Answer from Another Agent
<incorrect reasoning>
Final Answer: <incorrect answer>

```

Figure 8: The prompt template for generating feedback in debates.

```

### Instruction
You are a helpful assistant aiming to answer the following question with step-by-step reasoning.
Please output in the following format:
...

### Answer
<your step-by-step reasoning>
Final answer: <your final answer>
...

<choice prompt>Do not output unrelated content!

### Question
<question>

```

Figure 9: The prompt template for D&R distillation.

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

Teacher response in traditional debates:
To find the greatest common divisor (GCD) of three numbers:
1. Find the GCD of the first two numbers
2. Then find the GCD of that number and the third number
GCD(128, 144) $128 = 2^7$ $144 = 2^4 \cdot 3^2$
The prime factors are 2 and 3. The greatest power of 2 that divides both numbers is 2^4 . GCD(128, 144) = 16
GCD(16, 480) $480 = 2^5 \cdot 3 \cdot 5$ 16 and 480 are both divisible by 16.
Therefore, GCD(128, 144, 480) = 16

Teacher response in D&R training debates:
Let's solve this step by step using the Euclidean algorithm:
1) First, let's find GCD(128, 144)
* $144 = 1 \times 128 + 16$
* $128 = 8 \times 16 + 0$
* So, GCD(128, 144) = 16
2) Now, let's find GCD(16, 480)
* $480 = 30 \times 16 + 0$
* So, GCD(16, 480) = 16
3) Therefore, GCD(128, 144, 480) = 16
* We can verify this:
* $128 = 16 \times 8$
* $144 = 16 \times 9$
* $480 = 16 \times 30$
Final answer: 16

Student response in D&R training debates:
1. Find the greatest common divisors of 128 and 144 using the Euclidean algorithm.
2. Set up the equation: $a = bq + r$, where a is the larger number, b is the smaller number, q is the quotient, and r is the remainder.
3. Repeat step 1 with the new pair of numbers (b, r) until r is zero.
4. The greatest common divisor (gcd) is the last non-zero value of b.
5. Apply the same process to find the gcd of the gcd of 128 and 144, and 480.
Final answer: $128 = 4 \cdot 32^2 \quad 144 = 4 \cdot 36^2 \quad 480 = 2^4 \cdot 5^2 \cdot 36^2$
Using the Euclidean algorithm: $128 = 4 \cdot 144 + 112 \quad 144 = 1 \cdot 112 + 32 \quad 112 = 3 \cdot 36 + 40 \quad 36 = 1 \cdot 40 + 0$
.....
Therefore, gcd(128, 144, 480) = 36.

Self-reflection in D&R training debates:
..... The prime factorization was then compared to find a common factor, but this approach is not necessary and can lead to incorrect results when dealing with larger numbers. The correct method is to use the Euclidean algorithm directly to find the greatest common divisor of 128, 144, and 480.

Feedback in D&R training debates:
The incorrect answer has several issues:
1. The agent listed steps of the Euclidean algorithm but didn't actually execute them. Instead, they attempted to use prime factorization, which was unnecessary and led to incorrect calculations.
2. The prime factorizations provided are incorrect:
* $128 = 2^7$ (not $4 \cdot 32^2$)
* $144 = 2^4 \cdot 3^2$ (not $4 \cdot 36^2$)
* $480 = 2^5 \cdot 3 \cdot 5$ (not $2^4 \cdot 5^2 \cdot 36^2$)
.....
The agent should have followed through with the Euclidean algorithm steps they outlined and provided a single number as the GCD.

Figure 10: An example from the MATH dataset. Given the question, we show a response from Chen et al. (2024), which is generated with a traditional debate method without any self-reflection, feedback, and student models, and then two responses, a self-reflection, and feedback from our D&R training debates. Our debate data has a better quality with detailed responses, self-reflection, and feedback.