RETHINKING REASONING QUALITY IN LARGE LANGUAGE MODELS THROUGH ENHANCED CHAIN-OF-THOUGHT VIA RL

Anonymous authorsPaper under double-blind review

000

001

002

004

006

008 009 010

011 012

013

014

016

017

018

019

021

024

025

026

027

028

029

031

033

036

040

041

042

043

044

046

047

048

051

052

ABSTRACT

Reinforcement learning (RL) has recently become the dominant paradigm for strengthening the reasoning abilities of large language models (LLMs). Yet the rule-based reward functions commonly used on mathematical or programming benchmarks assess only answer format and correctness, providing no signal as to whether the induced Chain-of-Thought (CoT) actually improves the answer. Furthermore, such task-specific training offers limited control over logical depth and therefore may fail to reveal a model's genuine reasoning capacity. We propose Dynamic Reasoning Efficiency Reward (DRER) — a plug-and-play RL reward framework that reshapes both reward and advantage signals. (i) A Reasoning Quality Reward assigns fine-grained credit to those reasoning chains that demonstrably raise the likelihood of the correct answer, directly incentivising the trajectories with beneficial CoT tokens. (ii) A Dynamic Length Advantage decays the advantage of responses whose length deviates from a validation-derived threshold, stabilising training. To facilitate rigorous assessment, we also release *LogicTree*, a dynamically constructed deductive reasoning dataset that functions both as RL training data and as a comprehensive benchmark. Experiments show significant improvements in inference accuracy and logical consistency over the baseline methods at equal training steps, while the average confidence of CoT-augmented answers rises by 30%. The model further exhibits generalisation across diverse logical-reasoning datasets, and the mathematical benchmark AIME24. These results illuminate how RL shapes CoT behaviour and chart a practical path toward enhancing formal-reasoning skills in large language models. All code and data are available in our anonymous repository https://anonymous.4open.science/r/DRER-D34E.

1 Introduction

Recent reasoning models (DeepMind, 2024; Qwen, 2024; Team et al., 2025), including R1-like reproductions (Team et al., 2025; Mei et al., 2025; Yu et al., 2025; Shao et al., 2024; Hu, 2025; Kool et al., 2019; Ahmadian et al., 2024; Sutton et al., 1998), have adopted reinforcement learning (RL) to enhance chain-of-thought reasoning. By systematically exploring verifiable reasoning paths that lead to correct answers, these methods incrementally boost performance and deliver remarkable gains. Current RL-driven CoT approaches typically train on mathematics and programming benchmarks (OpenAI, 2024; Guo et al., 2025; Cobbe et al., 2021; Chen et al., 2021), whose inherently stepwise solution procedures serve as natural proxies for logical inference Wang et al. (2024a); Li et al. (2024), and they rely on rule-based reward OpenAI (2024); Guo et al. (2025) functions that assess only final answer correctness and formatting. This reliance stems from the straightforward evaluability of math and code tasks, where simple answer extraction or format checks suffice to assign reward signals and compute policy advantages.

However, this approach still faces two critical challenges. First, by relying solely on final-answer correctness as the reward signal, the model cannot distinguish which reasoning steps statistically boost the likelihood of the correct answers Paul et al. (2024), nor quantify each token's substantive contribution to the conclusion; instead, it may lean on "decorative" chains that diverge from genuine deductive paths Zhang et al. (2024), thereby undermining the accurate evaluation and effective training of its reasoning ability.

Second, the corpora used to reinforce "reasoning ability" are almost entirely drawn from execution-verifiable domains Sprague et al. (2024b)—such as mathematical problem sets and code synthesis tasks—while unified training data targeting pure formal logical inference remains severely lacking Morishita et al. (2024). Such constrained training regimens risk conceptual overextension Paul et al. (2024), whereby success on specific tasks is misconstrued as evidence of broadly applicable logical reasoning skills, potentially leading to an overestimation of the model's true inferential competence.



Figure 1: Overview of the Dynamic Reasoning Efficiency Reward (DRER) framework. Length₉₅ and Length₅ represent the 95th and 5th percentile lengths, respectively, computed from the validation set, and are used to normalize reasoning trajectory lengths according to task type or difficulty.

To address the limitations of outcome-only reward modeling in reasoning tasks, we propose *Dynamic Reasoning Efficiency Reward* (DRER), a plug-and-play reinforcement learning framework that reshapes both reward and advantage signals. DRER introduces two key mechanisms: (1) a *Reasoning-Quality Reward*, which assigns fine-grained credit to reasoning chains that statistically improve the likelihood of the correct answer, thereby reinforcing the utility of CoT tokens; and (2) a *Dynamic-Length Advantage*, which attenuates the policy advantage of responses whose lengths deviate from a validation-derived threshold, improving training stability. The overall framework is illustrated in Figure 1. In addition, we release *LogicTree*, a domain-agnostic deductive reasoning dataset carefully constructed to provide focused training supervision and to serve as a clean evaluation benchmark for identifying pathological reasoning behaviours.

Our experiments show that DRER delivers a marked leap in deductive skill: Qwen-2.5-7B-Instruct-1M model trained on LogicTree with 400 steps raises overall accuracy from 13% to 60%, still solves 31% of the hardest depth-8 items, and boosts answer-confidence by roughly 60% while reducing token consumption by 75% relative to DAPO/GRPO baselines.

The main contributions of this paper are summarized as follows:

- We propose DRER (Dynamic Reasoning Efficiency Reward), a novel reinforcement learning rewawrd framework that adaptively reshapes both reward and advantage signals to improve CoT reasoning.
- We release *LogicTree*, a domain-agnostic benchmark for formal deductive reasoning that serves dual purposes: functioning as both a focused training set and a clean evaluation benchmark, while providing highlight insights into LLMs reasoning behaviours.
- We systematically validate our approach through extensive experiments, confirming the effectiveness of our methodology in improving both reasoning quality and efficiency.

2 PRELIMINARY

Modeling Language Generation as a Token-Level MDP Reinforcement learning aims to learn a policy that maximizes cumulative reward through interaction with an environment. We model language generation as a sequential decision process within a Markov Decision Process (MDP) framework (Ouyang et al., 2022). Let $x=(x_0,\ldots,x_m)$ be the input prompt and $y=(y_0,\ldots,y_T)$ the generated response, with both drawn from a finite vocabulary \mathcal{A} . At step t, the state is $s_t=(x_0,\ldots,x_m,y_0,\ldots,y_t)$, and the action $a=y_{t+1}\in\mathcal{A}$ selects the next token. Transitions are deterministic: $\mathbb{P}(s_{t+1}\mid s_t,a)=1$, where $s_{t+1}=(x_0,\ldots,x_m,y_0,\ldots,y_{t+1})$. Generation ends upon producing a terminal token ω . The reward function R(s,a) provides scalar feedback on output quality. The initial state s_0 is the tokenized prompt, sampled from a distribution d_0 over inputs. This MDP formulation allows reinforcement learning—both value-based and value-free—to align language model generation with desired objectives and human preferences.

Group Relative Policy Optimization (GRPO) GRPO(Shao et al., 2024) removes the value function used in PPO(Schulman et al., 2017) and estimates the advantages within a group of G responses sampled by the behavior policy $\pi_{\theta_{\text{old}}}$ for each pair of questions-answers (q, a). GRPO maximizes a PPO-style clipped objective with an explicit KL penalty:

$$\begin{split} \mathcal{J}_{\text{GRPO}}(\theta) &= \mathbb{E}_{(q,a) \sim \mathcal{D}, \, \{o_i\} \sim \pi_{\theta_{\text{old}}}} \\ &\left[\frac{1}{G} \sum_{i=1}^{G} \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \left(\min (r_{i,t} \hat{A}_{i,t}, \, \text{clip}(r_{i,t}, 1-\epsilon, 1+\epsilon) \hat{A}_{i,t}) - \beta \, \mathcal{D}_{\text{KL}}(\pi_{\theta} \parallel \pi_{\text{ref}}) \right) \right], \end{split} \tag{1}$$

where

$$r_{i,t}(\theta) = \frac{\pi_{\theta}(o_{i,t} \mid q, o_{i, < t})}{\pi_{\theta_{\text{old}}}(o_{i,t} \mid q, o_{i, < t})}, \quad \hat{A}_{i,t} = \frac{R_i - \text{mean}(\{R_i\}_{i=1}^G)}{\text{std}(\{R_i\}_{i=1}^G)}.$$
(2)

GRPO first averages token-level losses within each response and then across the group, a sample-level aggregation that can implicitly favor longer responses and thus influence training dynamics (Liu et al., 2025).

Decouple Clip and Dynamic Sampling Policy Optimization (DAPO) DAPO(Yu et al., 2025) shares GRPO's group-based sampling and advantage normalization, but differs in two key aspects. First, it replaces GRPO's symmetric clipping with asymmetric clipping bounds, allowing for unbalanced exploration and conservative updates. Second, it introduces a dynamic sampling constraint that requires both correct and incorrect responses in the sampled group to ensure meaningful advantage shaping. The resulting objective is:

$$\mathcal{J}_{\text{DAPO}}(\theta) = \mathbb{E}_{(q,a) \sim \mathcal{D}, \{o_i\} \sim \pi_{\theta_{\text{old}}}} \left[\frac{1}{\sum_{i=1}^{G} |o_i|} \sum_{i=1}^{G} \sum_{t=1}^{|o_i|} \min \left(r_{i,t} \hat{A}_{i,t}, \operatorname{clip}(r_{i,t}, 1 - \varepsilon_{\text{low}}, 1 + \varepsilon_{\text{high}}) \hat{A}_{i,t} \right) \right],$$
(3)

where optimization is applied only if the sampled responses are not all equivalent to the reference answer. $r_{i,t}$ and $\hat{A}_{i,t}$ are defined as in Equation 2.

Reward Modeling Reward modeling in reinforcement learning (RL) for large language models (LLMs) is typically categorized into two approaches: rule-based rewards and learned reward models (RMs). Reward models, including outcome and process reward models (PRMs), learn a function through supervised learning, enabling finer-grained evaluation of intermediate reasoning steps. MATH-SHEPHERD Wang et al. (2024b) and OmegaPRM Luo et al. (2024) show that PRMs improve reasoning consistency and generalization, but they also raise annotation costs, introduce potential data bias (e.g., MCTS-generated traces), and reduce reliability in early-step evaluation, which can destabilize training.

Rule-based rewards are more widely adopted, where simple criteria such as answer correctness and syntactic validity are used to evaluate model outputs. Representative works Lyu et al. (2025); Xie et al. (2025); Li et al. (2025) like DeepSeek-R1 Guo et al. (2025) utilize correctness-based signals to construct efficient and interpretable training pipelines. The primary advantages of rule-based rewards are twofold: firstly, they exhibit low implementation cost and, secondly, they are characterised by high transparency. These properties render them well-suited for large-scale RL training. However, their limitations are also evident: these methods only evaluate final outcomes, ignoring the quality of intermediate reasoning steps. As a result, models may learn to "shortcut" reasoning, producing correct answers without coherent or logically valid chains of thought—leading to misalignment between reasoning processes and outputs Zhang et al. (2025).

3 Method

3.1 DRER

Rule-based rewards, such as answer correctness and format validity, minimal signals neglect to consider the reasoning trajectory that culminates in the ultimate response. Consequently, they

may permit verbose, irrelevant chains of thought, which compromise reasoning transparency and reliability.

In order to address this limitation, a novel reward framework, Dynamic Reasoning Efficiency Reward (DRER), is introduced. This plug-and-play system has been designed to shape not only the correctness of final outputs, but also the efficiency and utility of intermediate reasoning steps.

Given an input question x, the large-language model (LLM) π_{θ} produces an output sequence y autoregressively:

$$\pi_{\theta}(y \mid x) = \prod_{t=1}^{T} P_{\pi_{\theta}}(y_t \mid x, y_{< t}), \tag{4}$$

where the sequence y=[c,a] denotes the model's output sequence, where the first contiguous segment $c=(c_1,\ldots,c_{T_c})$ comprises the CoT tokens and the second segment $a=(a_1,\ldots,a_{T_a})$ contains the answer tokens. The overall sequence length satisfies $T=T_c+T_a$.

We believe that if the generated CoT tokens c are positive and coherent with the correct answer, it should *increase* the model's confidence in predicting ground-truth answer token:

$$\ell_{\text{CoT}} = \frac{1}{T_a} \sum_{t=1}^{T_a} \log \pi_{\theta} (a_t^{\star} \mid x_{CoT}, c, a_{< t}^{\star}), \quad \ell_{\text{NoCoT}} = \frac{1}{T_a} \sum_{t=1}^{T_a} \log \pi_{\theta} (a_t^{\star} \mid x_{NoCoT}, a_{< t}^{\star}), \quad (5)$$

CoT reasoning tokens that positively contribute to the model's ability to infer the correct answer should satisfy

$$\ell_{\text{CoT}} > \ell_{\text{NoCoT}}.$$
 (6)

where x_{CoT} and x_{NoCoT} denote the CoT and no CoT input question respectively; $c=(c_1,\ldots,c_{T_c})$ is the generated CoT of length T_c ; $a^\star=(a_1^\star,\ldots,a_{T_a}^\star)$ is the ground-truth answer consisting of T_a tokens, and $a_{< t}^\star$ stands for its prefix up to position t-1; Finally, π_θ is the autoregressive language model policy parameterised by θ .

To validate this hypothesis, we conduct experiments using Qwen2.5-7B-Instruct-1M on three benchmarks: GSM8K, Math500 (?) and LogicTree. For each dataset, we evaluate the quality of the generated Chain-of-Thought (CoT) traces using the same rubric by GPT-o4-mini. We further compute the log-probabilities assigned to the reference answers and group the samples based on final answer correctness. Results show that CoT traces rated as higher quality consistently yield significantly higher log-probabilities for the correct answers. Full experimental details and results are provided in section 4.4.

Reasoning Quality Reward To make the confidence-boosting property in equation 6 learnable, we define for each training instance \mathbf{x} the log-likelihood margin

$$\Delta(\mathbf{x}) = \ell_{\text{CoT}} - \ell_{\text{NoCoT}},\tag{7}$$

where ℓ_{CoT} and ℓ_{NoCoT} are given in equation 5. A positive $\Delta(\mathbf{x})$ indicates that the generated CoT reasoning tokens enhance the model's confidence in the correct answer, whereas a negative value reveals detrimental or spurious reasoning.

To obtain a numerically stable reward, we pass the margin through a smooth, bounded squashing function

$$R_q = \tanh(\Delta(\mathbf{x})), \tag{8}$$

yielding the *reasoning-quality reward*. The hyperbolic tangent preserves the sign of the margin, caps extreme values.

We incorporate R_q into the overall reinforcement-learning objective by maximising the expected composite return

$$R = R_{\text{task}} + \lambda_q R_q, \tag{9}$$

where $R_{\rm task}$ denotes the task-level reward (e.g., answer correctness) and $\lambda_q > 0$ is a weighting coefficient that balances task success and reasoning quality. This formulation directly rewards reasoning chains that demonstrably increase the likelihood of the correct answer while penalising uninformative or misleading chains, thereby systematically improving the model's logical reliability and interpretability.

Dynamic Length Advantage After every validation round we record the lengths $\{L_i\}$ of responses that are both correct and structurally valid within each difficulty bucket¹. The empirical 5% and 95% quantiles define a dynamic lower and upper length bound, $L_{\min}^{(d)}$ and $L_{\max}^{(d)}$, respectively, for bucket d. For a training sample i with effective response length ℓ_i , we introduce a multiplicative attenuation coefficient

$$g_i = \exp\left(-\frac{\max\{0, L_{\min}^{(d)} - \ell_i, \ell_i - L_{\max}^{(d)}\}}{\tau}\right), \qquad \tau > 0,$$
(10)

where $L_{\min}^{(d)}$ denotes the 5th-percentile response length observed in the previous validation step for bucket d, while $L_{\max}^{(d)}$ corresponds to the 95th percentile in the same distribution. The variable ℓ_i represents the effective response length of the current sample i, and $\tau \in [5, 10]$ is a temperature hyperparameter that controls the decay rate of the attenuation function.

The attenuation is then applied to the advantage computed by Group Computation, $\hat{A}_i = g_i \, A_i$, so that responses that are excessively short ($\ell_i < L_{\min}^{(d)}$) or verbose ($\ell_i > L_{\max}^{(d)}$) are exponentially down-weighted. This mechanism penalises pathological length behaviours while preserving the signal of well-sized, high-quality chains of thought. The complete algorithm procedure of DRER is detailed in Appendix 1.

3.2 Dataset

Most 'reasoning' benchmarks still fail to isolate formal deduction. Difficulty is inflated by injecting domain facts or arithmetic tricks, so logical skill is confounded with knowledge retrieval and calculation Lin et al. (2025); Sprague et al. (2024b). Logical depth and structure remain almost uncontrollable: items rarely reveal how accuracy decays as inference chains lengthen, and no systematic consistency checks can be run across paraphrased versions of the same proof pattern Saparov et al. (2023); Sprague et al. (2024a). Finally, intermediate steps are almost never evaluated; model capability is judged solely by the final answer Paul et al. (2024).




Figure 2: The framework of LOGICTREE automatic construction pipeline. We first sample atomic logic structures and sentences from seven deduction logic rules and four sentential logics, then fill it with natural statements in filtered AscentKB (Nguyen et al., 2021), and eventually construct the nested argument tree. Those intermediate will be hidden and transformed into questions.

LogicTree Therefore, we present the LogicTree dataset, based on nested deductive reasoning rules that poses significant challenges to state-of-the-art LRMs. Solving these problems requires models to not only recognize and correctly apply reasoning logic across diverse contexts but also to strategically plan hierarchical inference steps. Specifically, our dataset exhibits following key features: Programmatic Construction, The reasoning depth, breadth, and number of sub-questions are fully controllable. Beyond evaluating models' judgment on root conclusions, intermediate reasoning steps are extracted and expanded into sub-questions. Compared to prior deductive reasoning benchmarks, this enables granular assessment of models' hierarchical reasoning accuracy. Diverse Logical Forms, In contrast to grid puzzles or other logic games, LogicTree incorporates seven deductive reasoning rules and four sentential logic patterns, with each problem featuring distinct rule combinations. This significantly elevates the logical complexity. Probing LLMs' Foundational Reasoning, We

¹A bucket may correspond to a task type, question template, or any other granularity used in specific tasks.

 undertake multifaceted efforts to examine models' core logical capabilities. First, the dataset is decoupled from domain-specific knowledge to ensure models focus solely on pure logical reasoning. Second, we propose a logical consistency metric to evaluate models' ability to comprehend identical underlying logic across varying contextual representations.

LogicTree is constructed through three automated steps as shown in figure 2 and Appendix A.2

Table 1: Deductive reasoning rules statistics on LogicTree 9.6k problems spanning depth 8.

Deductive Rule	Logical Form	Amount
Modus Ponens	$(p \to q) \land p \implies q$	6760
Modus Tollens	$(p \to q) \land \neg q \implies \neg p$	6750
Hypothetical Syllogism	$(p \to q) \land (q \to r) \implies (p \to r)$	4 2 3 0
Disjunctive Syllogism	$(p \lor q) \land \neg p \implies q$	6 8 6 5
Reductio ad Absurdum	$(p \to q) \land (p \to \neg q) \implies \neg p$	6780
Conjunction Introduction	$(p \to q) \land (r \to s) \land (p \lor r) \implies (q \lor s)$	1 900
Conjunction Elimination	$(p \lor q) \land (p \to s) \land (q \to s) \implies s$	6 6 2 5

Evaluation The LogicTree dataset is programmatically generated with full control over logical depth, sub-problem quantity, and reasoning variations, which enables multifaceted analysis of models' logic mechanism from novel perspectives.

We introduce three evaluation metrics: (1) Accuracy: Standard correctness rate, only credited when every sub-question is correctly answered; (2) Consistency Ratio: Reasoning stability across logically equivalent queries, measured as consistent correctness over several isomorphic questions; (3) $F\beta$ -Score: Balances Answer Rate (proportion of valid *TruelFalse* responses) and Precision (accuracy among valid responses) with parameter β .

Note that, unlike traditional NLI datasets with three-class classification (Cheng et al., 2025; Liu & Zhang, 2024) (*True*, *False*, or *Uncertain*), we restrict labels to *True*/*False* to mitigate semantic ambiguity that often artificially inflates accuracy by encouraging defaulting to *Uncertain*. LLMs may respond with *Unknown* during inference, reducing statistical noise from random guessing.

4 Experiment

4.1 Experimental Settings

In our experiments, we use the LogicTree dataset, which contains 9,600 questions spanning 8 levels of reasoning depth. The dataset is split into training, validation, and test sets with a ratio of 10:1:1. The most challenging problems involve a reasoning depth of 8 and a branching width of 6, comprising 7 sub-questions derived from intermediate conclusions. To evaluate logical consistency, five questions were constructed for each logical structure by varying the natural language phrasing, ensuring that the model cannot achieve high consistency scores through random guessing.

We post-train Qwen2.5-7B-Instruct-1M using the LogicTree dataset within the proposed DRER framework on DAPO for 400 steps. Two baseline algorithms, GRPO and DAPO, are adopted for comparison. The task reward (R_{task}) consists of two components: a format score and an answer score.

Format Score: The format score (S_{format}) evaluates whether the model's response adheres to the required output structure:

$$S_{\text{format}} = \begin{cases} 1, & \text{if format is correct} \\ -1, & \text{if format is incorrect} \end{cases}$$

Answer Score: The answer score (S_{answer}) evaluates the correctness of the response content against the ground truth:

$$S_{\rm answer} = \begin{cases} 2, & \text{if the answer fully matches the ground truth} \\ -1.5, & \text{if the answer partially mismatches the ground truth} \\ -2, & \text{if the answer cannot be parsed or is missing} \end{cases}$$

The total task reward is computed as:

$$R_{\text{task}} = S_{\text{format}} + S_{\text{answer}}$$

The complete training parameters are detailed in the Appendix A.4.

4.2 MAIN RESULTS

Training Throughout the 400 optimisation steps, we observe a monotonic rise in the model's accuracy on the LogicTree from 7% at the outset to nearly 60% in figure 5. Additionally, the reasoning steps are streamlined for greater conciseness and clarity. Detailed evaluation data are in Table 2. In both settings, DRER consistently improves final accuracy and accelerates convergence. Figure 3 and Figure 6indicates the step at which the baseline (DAPO or GRPO) reaches its final precision, showing that DRER achieves a significantly higher or comparable performance earlier, highlighting its efficiency in guiding learning through structured reasoning signals.




Figure 3: Accuracy on the LogicTree during post-training with DAPO (left) and GRPO (right), with and without DRER.

Evaluation As demonstrated Table 2, even advanced models such as GPT-o3-mini, deepseek-r1, and Claude3.7 achieve accuracy scores below 20% on problems with reasoning depths of 7-8 in the LogicTree test set. The best performing model, Qwen3-235B, maintains the highest accuracy of 25% on problems with reasoning depth of 7, with an average accuracy of 53%. This reveals significant deficiencies in the complex deductive reasoning capabilities of existing reasoning models. In contrast, our trained 7B model achieves state-of-the-art performance in terms of average accuracy, showing substantial improvement over the base model, and maintains a 31% accuracy rate even at maximum reasoning depth.

Additionally, our experiments reveal distinct Unknown response tendencies across models. While GPT-o4-mini exhibits stronger reasoning capability than GPT-40, their comparable accuracy stems from GPT-o4-mini's overcaution (excessive Unknown responses). However, GPT-o4-mini achieves significantly higher Precision and $F\beta$ -Score scores in valid responses (see Appendix for details).

4.3 Does model really learn the logical paradigm?

A key question remains whether models truly understand logic or merely memorize puzzles. While prior work (Cheng et al., 2025) reveals models' tendency for self-contradiction on logically equivalent propositions, LogicTree naturally evaluates this through problems sharing identical logical structures but varying linguistic instantiations. Our Consistency Ratio metric quantifies this capability.

As shown in Table 3, most models can understand simple deductive reasoning logic, but at reasoning depths of 7-8, even state-of-the-art models such as GPT-o3-mini, Qwen3-235B, deepseek-r1, and Claude3.7 demonstrate consistency rates approaching zero, revealing current models' insufficient capability for consistent extended thinking and complex combinatorial logic.

Table 2: Comparison of LRM's(above) and LLM's(below) accuracy on LogicTree across various logical depth.

Model	1	2	3	4	5	6	7	8	Avg.
Qwen3-235B-A22B	0.96	0.83	0.66	0.71	0.46	0.32	0.25	0.07	0.53
Deepseek-R1	0.85	0.76	0.61	0.47	0.36	0.18	0.19	0.07	0.44
Claude-3.7-Sonnet	0.76	0.67	0.21	0.10	0.07	0.02	0.02	0.00	0.23
Qwen3-8B	0.86	0.83	0.49	0.44	0.32	0.11	0.14	0.08	0.41
GPT-o4-mini	0.74	0.64	0.25	0.20	0.10	0.06	0.05	0.02	0.26
GPT-o3-mini	0.66	0.56	0.07	0.07	0.03	0.02	0.01	0.00	0.18
Qwen3-4B	0.74	0.74	0.39	0.29	0.29	0.06	0.09	0.04	0.33
Gemini-2.5-Flash-Preview	0.86	0.64	0.41	0.31	0.24	0.11	0.06	0.00	0.33
GPT-4o	0.63	0.60	0.28	0.13	0.13	0.00	0.00	0.00	0.22
Phi-4-14B	0.72	0.67	0.31	0.27	0.19	0.04	0.01	0.01	0.28
Gemma-3-27B	0.65	0.41	0.15	0.04	0.00	0.00	0.00	0.00	0.16
Deepseek-v3	0.39	0.24	0.05	0.06	0.00	0.00	0.00	0.00	0.09
GPT-4o-mini	0.44	0.24	0.27	0.11	0.12	0.02	0.02	0.01	0.15
Qwen2.5-7B-Instruct-1M	0.36	0.29	0.15	0.12	0.08	0.01	0.01	0.00	0.13
GRPO	0.81	0.71	0.58	0.42	0.45	0.20	0.20	0.11	0.45
DAPO	0.88	0.73	0.66	0.47	0.60	0.36	0.23	0.20	0.52
GRPO+DRER	0.87	0.75	0.69	0.54	0.61	0.35	0.27	0.22	0.54
DAPO+DRER (Ours)	0.90	0.83	0.76	0.59	0.67	0.45	0.31	0.31	0.60 $^{\uparrow 0.47}$

Table 3: Comparison of Consistency Ratio on LogicTree. For the complete results referring to Appendix.

Model	1	2	3	4	5	6	7	8	Avg.
Qwen3-235B-A22B	0.90	0.65	0.30	0.50	0.15	0.00	0.00	0.00	0.32
Deepseek-R1	0.70	0.55	0.20	0.15	0.10	0.00	0.00	0.00	0.22
Claude-3.7-Sonnet	0.65	0.35	0.00	0.00	0.00	0.00	0.00	0.00	0.12
GPT-o4-mini	0.50	0.35	0.00	0.05	0.00	0.00	0.00	0.00	0.11
Ours (DAPO+DRER)	0.70	0.70	0.60	0.35	0.50	0.35	0.05	0.10	0.41 $^{\uparrow 0.40}$

Additionally, we analyzed whether models explicitly utilized certain deductive reasoning rules in their responses. Results in the Appendix provide word-frequency statistics and examples for GPT-o4-mini, DeepSeek-R1, Qwen3-235B, and our model, indicating a drop in explicit paradigm mentions with growing logical complexity and uneven competence across paradigms. Moreover, there exhibits varying capabilities across different logical paradigms. For example, DeepSeek-R1 responses most frequently reference "Modus Tollens", while "Disjunction Elimination" appears substantially less often. This disparity may stem from either the inherent complexity of the latter rule or inadequate exposure during pre-training. Our framework shows improved rule identification capacity with increasing response length and logical complexity.

4.4 Does model's reasoning behaviour become more effective?

To isolate the effect of explanatory CoT on answer confidence, we test QWEN2.5-7B-INSTRUCT-1M on 500 randomly sampled GSM8K and 500 LogicTree problems, generating for each prompt (i) a direct answer (No-CoT) and (ii) a step-by-step CoT that ends with the answer in a delimited span.We mark a CoT as **effective** if the model is *incorrect* in the No-CoT setting but *correct* once the CoT is included.For each sample we compute $\ell_{\text{CoT}} - \ell_{\text{NoCoT}}$. The mean log-probability gain of the ground-truth answer tokens a_t^* when conditioning on the CoT. We further bucket the samples into four categories (WR) wrong No-CoT / right CoT, (RR) right No-CoT / right CoT, (WW) wrong in both, and (RW) right No-CoT / wrong CoT. Table 12 and Table 13 report the statistics.

We also grouped the experimental data obtained from the two test sets based on the sign of $\Delta \ell$ as shown in Tables 14 and 15. In the responses with positive $\Delta \ell$, the model generally shows a higher fix rate after generating the CoT trajectory (i.e., the proportion of WR is higher), and the probability of successfully transitioning from the wrong answer to the correct one significantly increases. In contrast, in the responses with negative $\Delta \ell$, the model exhibits a higher break rate and a lower fix rate, making it more likely to transition from the correct answer back to the wrong answer.

Across both benchmarks the **WR** group exhibits the largest positive $\Delta \log p$ on average +2.46 nats for GSM8K and +1.8 nats for LogicTree—confirming that *effective* CoTs substantially raise the model's confidence in the correct answer. By contrast, cases where the model is already correct (**RR**) yield only marginal gains, and **WW** often shows negative shifts, suggesting that spurious reasoning can even erode confidence. In (**RW**) cases ,it's clearly that CoT tokens are negative.

We train the same base model with two popular base algorithms, GRPO and DAPO on DRER framework. Picture 4 shows that reasoning quality reward increases sharply ,which proves that the CoT tokens are becoming positive for model reasoning.

Figure 9 and figure 10 plot the prediction distribution for a difficulty-3 problem from 100 samples. Compared with the DAPO 400 step baseline, the DRER-trained 400 step policy produces a markedly sharper peak around the ground-truth answer, indicating that the learnt reasoning tokens help concentrate probability mass on the correct solution. Finally, Figure 7 and Figure 8 show that DRER keeps the average response length stable at fewer tokens, saving tokens per problem relative to the baseline while still achieving higher accuracy. This validates DRER ability to simultaneously improve reasoning quality and reduce inference cost.

4.5 Does model acquire the generalization ability?

Table 4: Performance of Qwen2.5 baseline and our model on various benchmarks.

Model	AIME24 (avg@32)	MMLU-redux (Wang et al., 2024c)	ZebraLogic (Lin et al., 2025) (Cell Acc)	ProntoQA (Saparov & He, 2022)
Qwen2.5-Instruct-1M	12.8	71.7	30.9	41.0
DAPO	13.9	72.4	32.3	42.0
Ours (DAPO+DRER)	16.5	73.3	33.4	45.0

Experiments in Table 4 test some benchmarks that may require similar logical capabilities. On logical benchmarks ZebraLogic (Lin et al., 2025) (grid-based SAT problems (Sempolinski, 2009) and ProntoQA (Saparov & He, 2022) (first-order logic with hierarchical ontology), our model (DAPO+DRER) shows consistent improvements.Notably, these gains extend to OOD datasets (AIME24 and MMLUredux (Wang et al., 2024c)), with AIME24 pass@32 scores matching QwQ32B (Yu et al., 2025). The model maintains 20% accuracy even on our LogicTree problems with unseen depth 9-10, indicating robust generalization to complex reasoning scenarios.

5 CONCLUSION AND FUTURE WORK

In this work, We present DRER, a token-level reward that ties each reasoning step to the model's confidence, and introduce LogicTree, a benchmark for formal deduction. On LogicTree and other tasks, DRER yields more consistent and accurate chain-of-thoughts than existing rewards. The post-trained model also generalizes to unseen datasets, indicating that broader logical training further strengthens its reasoning ability.

DRER builds reasoning quality directly into the RL objective. With bounded rewards, each policy update provably raises the expected return, while a length-aware advantage keeps responses from growing erratic. This pairing steers the search toward chains that truly boost answer confidence without inflating cost. Because DRER needs no value model, it drops into existing pipelines and scales to larger or multimodal models. As future work, our aim is to extend the framework to further refining reasoning-aligned training.

We release all code and the complete LogicTree corpus to ensure transparency and reproducibility. Together, DRER and LogicTree provide a lightweight, theoretically grounded basis for reasoning-aligned RL, enabling safer and more interpretable LLMs in logic-critical domains. Future work should extend this framework to richer logics and multimodal data.

REFERENCES

- Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin, Ahmet Üstün, and Sara Hooker. Back to basics: Revisiting reinforce style optimization for learning from human feedback in llms, 2024. URL https://arxiv.org/abs/2402.14740.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- Michael K Chen, Xikun Zhang, and Dacheng Tao. Justlogic: A comprehensive benchmark for evaluating deductive reasoning in large language models. *arXiv preprint arXiv:2501.14851*, 2025.
- Fengxiang Cheng, Haoxuan Li, Fenrong Liu, Robert van Rooij, Kun Zhang, and Zhouchen Lin. Empowering llms with logical reasoning: A comprehensive survey. *arXiv preprint arXiv:2502.15652*, 2025.
- Pengyu Cheng, Yong Dai, Tianhao Hu, Han Xu, Zhisong Zhang, Lei Han, Nan Du, and Xiaolong Li. Self-playing adversarial language game enhances llm reasoning. *Advances in Neural Information Processing Systems*, 37:126515–126543, 2024.
- Peter Clark, Oyvind Tafjord, and Kyle Richardson. Transformers as soft reasoners over language. *arXiv preprint arXiv:2002.05867*, 2020.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Google DeepMind. Gemini 2.0 flash thinking, 2024. URL https://deepmind.google/technologies/gemini/flash-thinking/.
- Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies. *Transactions of the Association for Computational Linguistics*, 9:346–361, 2021.
- Jiayi Gui, Yiming Liu, Jiale Cheng, Xiaotao Gu, Xiao Liu, Hongning Wang, Yuxiao Dong, Jie Tang, and Minlie Huang. Logicgame: Benchmarking rule-based reasoning abilities of large language models. *arXiv preprint arXiv:2408.15778*, 2024.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- Simeng Han, Hailey Schoelkopf, Yilun Zhao, Zhenting Qi, Martin Riddell, Wenfei Zhou, James Coady, David Peng, Yujie Qiao, Luke Benson, et al. Folio: Natural language reasoning with first-order logic. *arXiv preprint arXiv:2209.00840*, 2022.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1, 2021.
- Jian Hu. Reinforce++: A simple and efficient approach for aligning large language models. *arXiv* preprint arXiv:2501.03262, 2025.
- Mehran Kazemi, Quan Yuan, Deepti Bhatia, Najoung Kim, Xin Xu, Vaiva Imbrasaite, and Deepak Ramachandran. Boardgameqa: A dataset for natural language reasoning with contradictory information. *Advances in Neural Information Processing Systems*, 36:39052–39074, 2023.
- Wouter Kool, Herke van Hoof, and Max Welling. Buy 4 REINFORCE samples, get a baseline for free! In *Deep Reinforcement Learning Meets Structured Prediction, ICLR 2019 Workshop, New Orleans, Louisiana, United States, May 6, 2019.* OpenReview.net, 2019. URL https://openreview.net/forum?id=r11gTGL5DE.

- Chengshu Li, Jacky Liang, Andy Zeng, Xinyun Chen, Karol Hausman, Dorsa Sadigh, Sergey Levine,
 Li Fei-Fei, Fei Xia, and Brian Ichter. Chain of code: Reasoning with a language model-augmented
 code emulator. In *International Conference on Machine Learning*, pp. 28259–28277. PMLR, 2024.
 - Xuefeng Li, Haoyang Zou, and Pengfei Liu. Limr: Less is more for rl scaling, 2025. URL https://arxiv.org/abs/2502.11886.
 - Bill Yuchen Lin, Ronan Le Bras, Kyle Richardson, Ashish Sabharwal, Radha Poovendran, Peter Clark, and Yejin Choi. Zebralogic: On the scaling limits of llms for logical reasoning. *arXiv* preprint arXiv:2502.01100, 2025.
 - Hanmeng Liu and Yue Zhang. 大模型逻辑推理研究综述(survey on logical reasoning of large pre-trained language models). In Zhao Xin (ed.), *Proceedings of the 23rd Chinese National Conference on Computational Linguistics (Volume 2: Frontier Forum)*, pp. 48–62, Taiyuan, China, July 2024. Chinese Information Processing Society of China. URL https://aclanthology.org/2024.ccl-2.3/.
 - Hanmeng Liu, Leyang Cui, Jian Liu, and Yue Zhang. Natural language inference in context-investigating contextual reasoning over long texts. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pp. 13388–13396, 2021.
 - Hanmeng Liu, Jian Liu, Leyang Cui, Zhiyang Teng, Nan Duan, Ming Zhou, and Yue Zhang. Logiqa 2.0—an improved dataset for logical reasoning in natural language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:2947–2962, 2023.
 - Jian Liu, Leyang Cui, Hanmeng Liu, Dandan Huang, Yile Wang, and Yue Zhang. Logiqa: A challenge dataset for machine reading comprehension with logical reasoning. *arXiv* preprint *arXiv*:2007.08124, 2020.
 - Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. Understanding r1-zero-like training: A critical perspective, 2025. URL https://arxiv.org/abs/2503.20783.
 - Liangchen Luo, Yinxiao Liu, Rosanne Liu, Samrat Phatale, Meiqi Guo, Harsh Lara, Yunxuan Li, Lei Shu, Yun Zhu, Lei Meng, Jiao Sun, and Abhinav Rastogi. Improve mathematical reasoning in language models by automated process supervision, 2024. URL https://arxiv.org/abs/2406.06592.
 - Chengqi Lyu, Songyang Gao, Yuzhe Gu, Wenwei Zhang, Jianfei Gao, Kuikun Liu, Ziyi Wang, Shuaibin Li, Qian Zhao, Haian Huang, Weihan Cao, Jiangning Liu, Hongwei Liu, Junnan Liu, Songyang Zhang, Dahua Lin, and Kai Chen. Exploring the limit of outcome reward for learning mathematical reasoning, 2025. URL https://arxiv.org/abs/2502.06781.
 - Zhiyu Mei, Wei Fu, Kaiwei Li, Guangju Wang, Huanchen Zhang, and Yi Wu. Real: Efficient rlhf training of large language models with parameter reallocation. In *Proceedings of the Eighth Conference on Machine Learning and Systems, MLSys* 2025, Santa Clara, CA, USA, May 12-15, 2025. mlsys.org, 2025.
 - Terufumi Morishita, Gaku Morio, Atsuki Yamaguchi, and Yasuhiro Sogawa. Enhancing reasoning capabilities of Ilms via principled synthetic logic corpus. *Advances in Neural Information Processing Systems*, 37:73572–73604, 2024.
 - Tuan-Phong Nguyen, Simon Razniewski, and Gerhard Weikum. Advanced semantics for commonsense knowledge extraction. In *Proceedings of the Web Conference 2021*, pp. 2636–2647, 2021.
 - OpenAI. Learning to reason with llms, 2024. URL https://openai.com/index/learning-to-reason-with-llms/.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow
 instructions with human feedback. Advances in neural information processing systems, 35:27730–27744, 2022.

- Debjit Paul, Robert West, Antoine Bosselut, and Boi Faltings. Making reasoning matter: Measuring and improving faithfulness of chain-of-thought reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 15012–15032, 2024.
 - Qwen. Qwq-32b: Embracing the power of reinforcement learning, 2024. URL https://qwenlm.github.io/blog/qwq-32b/.
 - Abulhair Saparov and He He. Language models are greedy reasoners: A systematic formal analysis of chain-of-thought. *arXiv preprint arXiv:2210.01240*, 2022.
 - Abulhair Saparov, Richard Yuanzhe Pang, Vishakh Padmakumar, Nitish Joshi, Mehran Kazemi, Najoung Kim, and He He. Testing the general deductive reasoning capacity of large language models using ood examples. *Advances in Neural Information Processing Systems*, 36:3083–3105, 2023.
 - John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017. URL https://arxiv.org/abs/1707.06347.
 - Peter Sempolinski. Automatic solutions of logic puzzles. 2009.
 - Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, YK Li, Yu Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
 - Zayne Sprague, Xi Ye, Kaj Bostrom, Swarat Chaudhuri, and Greg Durrett. Musr: Testing the limits of chain-of-thought with multistep soft reasoning. *ICLR*, 2024a.
 - Zayne Sprague, Fangcong Yin, Juan Diego Rodriguez, Dongwei Jiang, Manya Wadhwa, Prasann Singhal, Xinyu Zhao, Xi Ye, Kyle Mahowald, and Greg Durrett. To cot or not to cot? chain-of-thought helps mainly on math and symbolic reasoning. *arXiv preprint arXiv:2409.12183*, 2024b.
 - Richard S Sutton, Andrew G Barto, et al. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
 - Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, et al. Kimi k1. 5: Scaling reinforcement learning with llms. *arXiv preprint arXiv:2501.12599*, 2025.
 - Nemika Tyagi, Mihir Parmar, Mohith Kulkarni, Aswin Rrv, Nisarg Patel, Mutsumi Nakamura, Arindam Mitra, and Chitta Baral. Step-by-step reasoning to solve grid puzzles: Where do llms falter? *arXiv preprint arXiv:2407.14790*, 2024a.
 - Nemika Tyagi, Mihir Parmar, Mohith Kulkarni, Aswin Rrv, Nisarg Patel, Mutsumi Nakamura, Arindam Mitra, and Chitta Baral. Step-by-step reasoning to solve grid puzzles: Where do llms falter? In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 19898–19915, 2024b.
 - Ke Wang, Junting Pan, Weikang Shi, Zimu Lu, Houxing Ren, Aojun Zhou, Mingjie Zhan, and Hongsheng Li. Measuring multimodal mathematical reasoning with math-vision dataset. Advances in Neural Information Processing Systems, 37:95095–95169, 2024a.
 - Peiyi Wang, Lei Li, Zhihong Shao, R. X. Xu, Damai Dai, Yifei Li, Deli Chen, Y. Wu, and Zhifang Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations, 2024b. URL https://arxiv.org/abs/2312.08935.
 - Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyan Jiang, et al. Mmlu-pro: A more robust and challenging multitask language understanding benchmark. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024c.
 - Chulin Xie, Yangsibo Huang, Chiyuan Zhang, Da Yu, Xinyun Chen, Bill Yuchen Lin, Bo Li, Badih Ghazi, and Ravi Kumar. On memorization of large language models in logical reasoning. In *The* 4th Workshop on Mathematical Reasoning and AI at NeurIPS'24.

- Tian Xie, Zitian Gao, Qingnan Ren, Haoming Luo, Yuqian Hong, Bryan Dai, Joey Zhou, Kai Qiu, Zhirong Wu, and Chong Luo. Logic-rl: Unleashing llm reasoning with rule-based reinforcement learning, 2025. URL https://arxiv.org/abs/2502.14768.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov, and Christopher D Manning. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. *arXiv preprint arXiv:1809.09600*, 2018.
- Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiaze Chen, Jiangjie Chen, Chengyi Wang, Hongli Yu, Weinan Dai, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu, Wei-Ying Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingxuan Wang. Dapo: An open-source llm reinforcement learning system at scale, 2025. URL https://arxiv.org/abs/2503.14476.
- Dan Zhang, Sining Zhoubian, Yisong Yue, Yuxiao Dong, and Jie Tang. Rest-mcts*: Llm self-training via process reward guided tree search. *arXiv preprint arXiv:2406.03816*, 2024.
- Yi-Fan Zhang, Xingyu Lu, Xiao Hu, Chaoyou Fu, Bin Wen, Tianke Zhang, Changyi Liu, Kaiyu Jiang, Kaibing Chen, Kaiyu Tang, Haojie Ding, Jiankang Chen, Fan Yang, Zhang Zhang, Tingting Gao, and Liang Wang. R1-reward: Training multimodal reward model through stable reinforcement learning, 2025. URL https://arxiv.org/abs/2505.02835.
- Qin Zhu, Fei Huang, Runyu Peng, Keming Lu, Bowen Yu, Qinyuan Cheng, Xipeng Qiu, Xuanjing Huang, and Junyang Lin. Autologi: Automated generation of logic puzzles for evaluating reasoning abilities of large language models. *arXiv preprint arXiv:2502.16906*, 2025.

A TECHNICAL APPENDICES AND SUPPLEMENTARY MATERIAL

Table 5: An example of a logictree puzzle.

An example of a logictree puzzle

Paragraph:

On the condition that coral reefs need light to grow so only occur in shallow waters, it is definitely the case that in addition to this, olive oil is also ideal for frying and is the most stable fat when heated. If in addition to this, olive oil is also ideal for frying and is the most stable fat when heated, then if ribbons relate post: How to sew trims, then titanium dioxide and zinc oxide also functioned in this fashion. It is a fact that either the anus of this invertebrate is located on top of its body or coral reefs need light to grow so only occur in shallow waters. The statement that 'the anus of this invertebrate is located on top of its body' is incorrect.

Ouestion:

It is a common misconception that if ribbons relate post: How to sew trims, then titanium dioxide and zinc oxide also functioned in this fashion.

Solution:

False

A.1 SEVEN DEDUCTIVE PARADIGMS IN LOGICTREE

LogicTree centres on seven classic deductive paradigms that constitute the atomic reasoning units of every sample. Each paradigm is implemented as a dedicated Python class (see logic.py) whose constructor generates the required premises and the logically entailed conclusion. The table below summarises their formal schemata together with bilingual surface examples.

705706

708

710

711

712

713

714

715

716

717

718

719

720

721

722

723

724

725

726 727

728 729

730

731

732

733

734

735

737

738

739

740

741

742

743

Table 6: Model Response of logictree.

Model Response of logictree

Paragraph:

When the notion that 'if the statement that states the worms also eat the food scraps and worm bin bedding is false, then the statement 'emergent wetland vegetation is rooted in soil that is under the water for most of the time' can be considered false' is untrue is true, it follows that hydrangeas need minimal care in well-drained, fertile soil, and are shade lovers. One may reasonably assume that if the notion that 'if the statement that states the worms also eat the food scraps and worm bin bedding is false, then the statement 'emergent wetland vegetation is rooted in soil that is under the water for most of the time' can be considered false' is untrue, then those that suggest hydrangeas need minimal care in well-drained, fertile soil, and are shade lovers are wrong. Assuming sometimes, cancer cells left behind after treatment eventually form new tumors, we know that if either eggs from common ducks require about 28 days to hatch or the island was originally inhabited by Ligures Ilvates, who gave it the ancient name Ilva, then if the statement that states the worms also eat the food scraps and worm bin bedding is false, then emergent wetland vegetation is rooted in soil that is under the water for most of the time. The statement that states sometimes, cancer cells left behind after treatment eventually form new tumors is right. Whenever the claim that either eggs from common ducks require about 28 days to hatch or the island was originally inhabited by Ligures Ilvates, who gave it the ancient name Ilva is false is true, the following statement is also true: the bowl acts as a carb hole. Whenever the claim that either eggs from common ducks require about 28 days to hatch or the island was originally inhabited by Ligures Ilvates, who gave it the ancient name Ilva is false is true, the following statement is also true: the statement 'the bowl acts as a carb hole' can be considered false.

Question:

The statement 'the worms also eat the food scraps and worm bin bedding' can be considered true.

It is a common misconception that if the statement that states the worms also eat the food scraps and worm bin bedding is false, then emergent wetland vegetation is rooted in soil that is under the water for most of the time.

Model Response:

- 1. Analyze the logical structure of the given statements.
- 2. Identify the conditions and implications in each statement.
- 3. Determine the truth value of the given statements and how they relate to each other.
- 4. Check the consistency of the statements with the provided premises.
- 5. Deduce the implications of the given statements on the conclusion.
- 6. Evaluate the final conclusion based on the logical consistency and the given premises.

</think>

<answer>[True, False]</answer>

Solution:

True

False

744745746747

748

749 750

751

752

753

754

755

A.2 LOGICTREE: TEMPLATE AND CONSTRUCTION

We construct LogicTree through three automated steps:

- 1. **Logical Node Sampling.** Atomic premises and target conclusions are sampled from seven classical deductive rules (e.g., Modus Ponens, Modus Tollens) and four sentential logics, generating symbolic propositions.
- 2. **Natural-Language Instantiation.** Each symbolic proposition is mapped to natural declarative statements retrieved from the filtered AscentKB corpus Nguyen et al. (2021), excluding ambiguous expressions or compound sentences to enhance lexical diversity while maintaining clarity.

Table 7: Full Chain-of-Thought (CoT) Prompt Template Used for DRER Training and Evaluation.

COT prompt

System Input:

system

You are a helpful assistant. The assistant first thinks step by step about the reasoning process in the mind and then provides the user with the answer.

The reasoning process and answer are enclosed within <think> ... </think> and <answer> ... </answer> tags, respectively, i.e.

<think> Write the reasoning process for the given paragraph here

<answer> Fill in the final answer list for {num_q} question(s) here: True, False or Unknown.

Like this: [True, False...] </answer>

You must choose one of the following answers:

- TRUE: if the premises entail the statement
- FALSE: if the premises contradict the statement
- UNKNOWN: if you cannot determine the truth value of the statement from the premises You will be given a paragraph of logical premises and a statement. Perform logical reasoning **strictly based on the premises** using propositional logic.

User Input:

startl>user

Paragraph: {paragraph}
{current_question}

endl> <lim_startl>assistant <think>

Variable meanings:

{num_q}: Number of questions in the current prompt.

{paragraph}: The paragraph containing the logical premises.

{current_question}: The specific statement whose truth value is to be evaluated.

3. Nested-Tree Assembly. The instantiated nodes are recursively composed into reasoning trees with configurable depth and width. Intermediate conclusions are masked from given premises, then transformed into sub-questions to create multi-step problem instances. This design ensures the inference process depends solely on logical form rather than sentence semantics, effectively decoupling reasoning from world knowledge.

A.3 PRIMITIVE AND COMPOUND PROPOSITIONS

LogicTree expresses every deductive instance in terms of one *primitive statement* and four *compound connectives*. The primitive Statement captures an atomic fact—e.g. "Alice studies."— while the four connectives build larger formulas: *negation*, *conjunction*, *implication*, and *inclusive disjunction*. Each connective is implemented as a dedicated class whose method .nl() randomly selects a surface template from expressions.json. Table 10 summarises the five constructs, their formal notation, and representative English renderings.

Surface realisation. When generating a sample, the pipeline first creates Statement objects for the chosen entities, then composes them with the connectives above. For example, calling Negation(S).nl() yields a randomly chosen negated template such as "The claim that S is false."; calling Conditional(P,Q).nl() may return "Provided that P, we know that Q.". This template sampling, combined with optional adverb or negator insertion, gives LogicTree a high level of lexical diversity while preserving formal truth values.

Table 8: Full No-CoT Prompt Template used for DRER training and evaluation.

No-CoT Prompt

System Input:

<|im_start|>system

You are a helpful assistant. You answer questions by solely using logical reasoning.

You will be given a paragraph of logical premises and a statement. Perform logical reasoning **strictly based on the premises** using propositional logic.

Assume all premises are true. Do not rely on prior world knowledge.

<answer> Fill in the final answer list for {num_q} question(s) here:
True, False or Unknown. Like this: [True, False...] </answer>
You must choose one of the following answers:

- TRUE: if the premises entail the statement
- FALSE: if the premises contradict the statement
- UNKNOWN: if you cannot determine the truth value of the statement based on the premises

<|im_end|>

User Input:

<|im_start|>user
Paragraph: {paragraph}
{current_question}
<|im_end|>
<|im_start|>assistant
<answer>...</answer>

Variable meanings:

{num_q}: Number of questions in the current prompt.
{paragraph}: Paragraph containing the logical premises.
{current_question}: Statement whose truth value is to be evaluated.

A.4 TRAINING SETTING

Table 11 records important training parameters. Experiments are conducted on 4×H20 (80G) GPUs with CUDA 12.0, PyTorch 2.6.0, transformers 4.47.1. The Main Experiment phase (DAPO+DRER) trains for 400 training steps and takes approximately 50 hours. Training is carried out with a learning rate of 3×10^{-7} , a maximum response length of 4096 tokens, the batch size is 16 and 16 responses per prompt. For GRPO, the KL divergence coefficient is set to 0.001. In the DRER framework, we set $\lambda_q=1$ and $\tau=8$.

B RELATED DATASET

Logical reasoning datasets can broadly be categorized into three types. The first type focuses on deductive reasoning. The second type is based on grid-based logic puzzles. The third category comprises datasets based on multi-hop or strategic question answering. These datasets assess language models' logical capabilities from various perspectives, including formal logic, multi-step planning, structural induction, and strategy analysis. In addition, there are general-purpose reasoning datasets that are also frequently used to evaluate LLMs' logical reasoning abilities.

899 900 901

902 903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

```
Algorithm 1 DRER: Dynamic Reasoning Efficiency Reward.
              Require: Prompts P = \{q_b\}_{b=1}^B, ground-truth answers Y^\star = \{a_b^\star\}_{b=1}^B,
866
                1: policy \pi_{\theta}, rule reward R_{\text{rule}}(\cdot), reasoning weight \lambda_{q},
867
               2: bucket IDs \{d_b\}_{b=1}^B, bounds (L_{\min}^{(d)}, L_{\max}^{(d)}), temperature \tau
868
              Ensure: Advantages A \in \mathbb{R}^{B \times L}
                    (1) Build trajectories
870
               3: C \leftarrow \pi_{\theta}(P, \text{mode} = cot)
                                                                                                                                              871
               4: for b = 1 to B do
872
                          t_n[b] \leftarrow \text{NoCoTPROMPT}(q_b) \parallel \text{FORMATANSWER}(a_b^{\star})
873
                          Replace answer span in C[b] with a_b^{\star} \to t_c[b]; record span \mathcal{A}_b
874
               7: end for
875
                    (2) Reasoning-quality reward
876
                   for b = 1 to B do
                         \begin{aligned} &\ell_c = \frac{1}{|\mathcal{A}_b|} \sum_{t \in \mathcal{A}_b} \log p_{\theta}(a_{b,t}^{\star} \mid t_c[b]) \\ &\ell_n = \frac{1}{|\mathcal{A}_b|} \sum_{t \in \mathcal{A}_b} \log p_{\theta}(a_{b,t}^{\star} \mid t_n[b]) \\ &R_q[b] \leftarrow \tanh(\ell_c - \ell_n) \end{aligned}
877
878
              10:
879
              11:
880
                          R_{\text{seq}}^{\text{fill}}[b] \leftarrow R_{\text{rule}}(C[b]) + \lambda_q R_q[b]
              12:
                          Expand R_{\text{seq}}[b] to token reward r_b, on C[b]
              13:
              14: end for
883
                    (3) Group-wise normalisation
              15: for all prompt group g do
                          \mu_g \leftarrow \text{mean}(r_{m,.}), \ \sigma_g \leftarrow \text{std}(r_{m,.}) \quad (m \in g)
885
              16:
886
              17:
                                                                                                                                             \triangleright raw advantage A
                                \tilde{A}_{m,\cdot} \leftarrow \frac{r_{m,\cdot} - \mu_g}{\sigma_g + \varepsilon}
              18:
                          end for
              19:
889
              20: end for
890
                    (4) Dynamic-length attenuation
891
              21: for b = 1 to B do
892
893
                         g_b \leftarrow \exp\left(-\frac{\max\{0, L_{\min}^{(d)} - \ell_b, \ell_b - L_{\max}^{(d)}\}}{\tau}\right)
894
895
896
              25: end for
897
              26: return A
```

DEDUCTIVE REASONING

ConTRoL (Liu et al., 2021), consisting of 8,325 pairs of expert-designed datasets, is a challenging segment-level NLI dataset to evaluate model's contextual reasoning capacity from police recruitment tests. RuleTaker (Clark et al., 2020) is a benchmark dataset designed to test whether language models can logically reason about natural language rules and facts by determining whether the conclusions follow, do not follow, or are uncertain. LogiQA (Liu et al., 2020) is a benchmark of 8,678 civil service exam questions designed to evaluate models' reading comprehension and deductive reasoning across five logical types by requiring conclusion drawing from textual premises. LogiQA2.0 (Liu et al., 2023) is the enchanced version of LogiQA citeplogiqa, featuring improved translations, expert-verified annotations, and new NLI tasks, designed to evaluate logical reasoning and reading comprehension in MRC and NLI formats. FOLIO Han et al. (2022) is an maually annotated dataset containing 1,430 logically complex natural language reasoning examples with first-order logic (FOL) annotations, designed to evaluate and benchmark the deductive reasoning and NL-FOL translation capabilities of Large Language models. PrOntoQA (Saparov & He, 2022) is a benchmark proposed in 2022 to evaluate LLMs' reasoning by generating question-answer pairs from first-order logic, revealing their struggles with multi-step proof planning despite valid individual steps. Compared with PrOntoQA (Saparov & He, 2022), PrOntoQA-OOD (Saparov et al., 2023) is designed to evaluate the general deductive reasoning abilities of LLMs by testing their ability to

Table 9: Seven deductive paradigms that serve as the atomic reasoning units in LOGICTREE.

Paradigm	Formal Schema	Surface Realisation
Modus Ponens	$(p \! \to \! q) \land p \ \Rightarrow \ q$	If Alice studies, she will pass. Alice studies. Therefore, she will pass.
Modus Tollens	$\begin{array}{cc} (p \to q) \land \neg q & \Rightarrow \\ \neg p \end{array}$	If it rains, the road is wet. The road is not wet. Thus, it did not rain.
Hypothetical Syllogism	$\begin{array}{ccc} (p \to q) \land (q \to r) \\ r) \Rightarrow (p \to r) \end{array}$	If A wins, B celebrates. If B celebrates, C is happy. Hence, if A wins then C is happy.
Disjunctive Syllogism	$(p \lor q) \land \neg p \ \Rightarrow \ q$	Either today is Monday or Tuesday. Today is not Monday. Therefore, today is Tuesday.
Reductio ad Absurdum	$\begin{array}{c} (p \to q) \land (p \to \neg q) \Rightarrow \neg p \end{array}$	Assume the number is both even and odd. This leads to a contradiction. Thus, the number is not both even and odd.
Constructive Dilemma	$\begin{array}{c} (p \to q) \land (r \to s) \land \\ (p \lor r) \Rightarrow (q \lor s) \end{array}$	If it rains, we stay in; if it is sunny, we picnic. Either it rains or it is sunny. Hence, we either stay in or picnic.
Disjunction Elimination	$\begin{array}{c} (p \lor q) \land (p \to s) \land \\ (q \to s) \Rightarrow s \end{array}$	Either I study or I work. If I study, I will learn. If I work, I will learn. Thus, I will learn.

Table 10: Primitive and compound proposition types used in LOGICTREE.

Construct	Logical Form	Example Surface Realisation (EN)
Statement (atomic)	p	Alice studies.
Negation	$\neg p$	It is not true that Alice studies.
Conjunction	$P \wedge q$	Alice studies and Bob plays chess.
Implication (Conditional)	$P \rightarrow q$	If it rains, then the road becomes wet.
Inclusive Disjunction	$P \lor q$	Either today is Monday or Tuesday.

generalize to more complex, compositional proofs, particularly those that are out-of-distribution (OOD). JustLogic (Chen et al., 2025) a generated deductive reasoning benchmark designed to evaluate LLMS, featuring high complexity, being independent of prior knowledge, and conducting in-depth error analysis in terms of reasoning depth and argumentative form.

However, the existing logical reasoning datasets still have some limitations. Most datasets have fixed or limited reasoning depth and breadth, which limits their ability to conduct a comprehensive evaluation of complex multi-step reasoning models. Many datasets entwine semantic information with logic, which may lead the model to rely on semantic cues rather than pure logical reasoning.

Furthermore, the majority focus only on final answer correctness, lacking assessment of the intermediate reasoning process and overall explanation quality.

In contrast, the LogicTree dataset we proposed has significant advantages: it is programmed and dynamically constructed, allowing for flexible control over the depth, breadth, and difficulty of inference; It separates semantics from logic to precisely evaluate pure deductive reasoning; It introduces a new logical consistency metric across multiple logical equivalence problems to measure the model's grasp of the underlying logical structure.

B.2 GRID-BASED LOGIC PUZZLES

BoardgameQA(Kazemi et al., 2023) is a dataset designed to evaluate the reasoning ability of language models when dealing with contradictory information.GridPuzzle(Tyagi et al., 2024a) is a dataset of grid-based logic puzzles designed to evaluate LLMs' structured, multi-step reasoning abilities through both final answers and detailed reasoning chains. The Knights and Knaves(Xie et al., 2025) dataset is an reasoning dataset designed to test logical deduction, where characters are either knights (truth-tellers) or knaves (liars), featuring controlled difficulty levels, procedural generation, and verifibility.

Table 11: Important Training Parameters.

Algorithm	Train Batch Size	Rollout N	KL Coef	Max Response Len
GRPO	16	16	0.001	4096
DAPO	16	16	-	4096

Existing datasets, such as GridPuzzle (Tyagi et al., 2024b), Knights and Knaves (KK) (Xie et al.) provide valuable reasoning benchmarks, but they all have limitations. For example, KK (Xie et al.) entangles logical reasoning with semantic cues, taking the risk of rapid learning through keyword associations. Some logic puzzle focuses on the final answer without verifying the intermediate steps, allowing the model to guess without sufficient reasoning.

On the contrary, LogicTree evaluates the final and intermediate steps and executes the complete reasoning chain. It also introduces a logical consistency rate among variants of the same logical form and uses semantic-logical unentanglement to ensure that the model relies on reasoning rather than superficial clues.

B.3 MULTI-HOP OR STRATEGIC QUESTION ANSWERING

HotpotQA (Yang et al., 2018) is a multi-hop question-answering dataset that requires reasoning across multiple documents and provides supporting facts to enhance the interpretability of the QA system.StrategyQA (Geva et al., 2021) is a benchmark dataset designed to evaluate implicit multi-step reasoning in LLMs across 15 domains and 13 strategies. SPAG (Cheng et al., 2024) is self-laying based adversarial language game dataset designed to enhance and evaluate the reasoning ability through a game involving indirect communication and strategic reasoning about hidden target words. LOGICGAME (Gui et al., 2024) is a benchmark designed to evaluate LLMs' ability to understand, execute, and plan based on predefined rules through diverse, verifiable game scenarios requiring multi-step logical reasoning.AutoLogi (Zhu et al., 2025) is benchmark test for open-ended logic puzzles with controllable difficulty and program-based verification, designed to evaluate the reasoning ability of LLM.

Compared with datasets such as HotpotQA (Yang et al., 2018), StrategyQA (Geva et al., 2021), they emphasize various forms of multi-step or strategic reasoning across natural language problems, but there are still obvious limitations: The reasoning strategies in existing datasets are often broad and empirical rather than based on formal logical deduction frameworks (for example, StrategyQA (Geva et al., 2021) relies on heuristic and empirical categories). Many datasets focus on language pattern matching or cross-document evidence aggregation rather than verifying the true formal reasoning process (for example, HotpotQA (Yang et al., 2018)).LogicTree, on the other hand, strictly adheres to classical mathematical logic, adopting clear and well-defined deduction rules, and does not rely on common sense knowledge, providing a purest logical reasoning environment.

B.4 GENERAL PURPOSE DATASET

MMLU-Pro(Wang et al., 2024c) is an advanced benchmark of 12,000 expert-reviewed, 10-option questions across 14 disciplines, designed to better evaluate LLM performance with greater difficulty and reduced noise than the original MMLU (Hendrycks et al., 2021). However, it primarily evaluates broad knowledge and reasoning abilities rather than focusing on strong formal logical reasoning. Thus, it is not specifically designed to test models' capabilities in complex multi-step logical dedu

C PROMPT TEMPLATES

Tables 7 and 8 list the exact prompts used in our experiments: a Chain-of-Thought (CoT) version that elicits step-by-step reasoning, and a No-CoT variant that asks for the final answer only. Curly-braced placeholders are replaced at runtime ({paragraph}, {current_question}, {num_q}). The two prompts share identical task instructions, so performance differences isolate the effect of showing or hiding the reasoning chain.

D TRAINING DETAILS

Table 12: Average $\ell_{\text{CoT}} - \ell_{\text{NoCoT}}$ by answer transition in GSM8K.

Original \downarrow / With CoT \rightarrow	Wrong (W)	Correct (R)
Wrong (W)	-4.32	2.46
Correct (R)	-5.00	-0.47

Table 13: Average $\ell_{\text{CoT}} - \ell_{\text{NoCoT}}$ by answer transition in LogicTree.

Original \downarrow / With CoT \rightarrow	Wrong (W)	Correct (R)
Wrong (W)	-1.13	1.81
Correct (R)	-3.79	-4.76




Figure 4: Reasoning quality reward on the LogicTree during post-training with DRER.

E SUPPLEMENTARY EVALUATION

Table 16 presents the entire evaluation data of Consistency Ratio. Figure 11 shows the distribution of those deduction logical key words in LLMs response. Figure 12 plot the complete evaluation data of $F\beta$ -Score, which provides a balanced metric to compare the comprehensive performance across those LLMs. Tables 4 records the evaluation results on other ood benchmarks, including AIME24, MMLU-redux (Wang et al., 2024c), ZebraLogic (Lin et al., 2025) and ProntoQA (Saparov & He, 2022).

F LIMITATIONS

Despite the empirical gains achieved by DRER and LogicTree, several limitations remain:

Table 14: Answer-transition proportions conditioned on the sign of $\Delta \ell = \ell_{CoT} - \ell_{NoCoT}$ on GSM8K (N=500). $p(W \rightarrow R)$ is the fix rate; $p(R \rightarrow W)$ is the break rate.

Group by $\Delta \ell$ sign	#Instances	Mean $\Delta \ell$	$p(W \rightarrow R)$	$p(R \rightarrow W)$
$\Delta \ell > 0$ (CoT favored)	140	+2.20	0.74	0.02
$\Delta \ell < 0$ (NoCoT favored)	360	-2.60	0.02	0.24

Table 15: Answer-transition proportions conditioned on the sign of $\Delta \ell = \ell_{\text{CoT}} - \ell_{\text{NoCoT}}$ on LogicTree (N=500). $p(W \rightarrow R)$ is the fix rate; $p(R \rightarrow W)$ is the break rate.

Group by $\Delta \ell$ sign	#Instances	Mean $\Delta \ell$	$p(W \rightarrow R)$	$p(R \rightarrow W)$
$\Delta \ell > 0$ (CoT favored)	120	+1.70	0.67	0.04
$\Delta \ell < 0$ (NoCoT favored)	380	-3.90	0.05	0.23

- Logic coverage. LogicTree is limited to the deductive reasoning paradigm, while more diverse
 forms such as analogical reasoning, inductive reasoning, or traceable reasoning have not yet been
 evaluated.
- Model scale and cost. All experiments use Qwen-2.5-7B-Instruct-1M as backbone. The memory and latency overhead of token-level rewards on 70 B-scale or MoE models is unknown and may be prohibitive.
- Evaluation bias. Training and evaluation rely on an automatic logic verifier and confidence scores; no human preference or chain-quality annotation is included, which may overlook subjective aspects of reasoning quality.
- Synthetic corpus and social bias. LogicTree sentences are synthetically generated; potential
 social biases or misuse risks in real-world deployments have not been systematically analysed.

In future work we plan to extend DRER to higher-order logic, explore low-cost reward approximations, and incorporate human evaluation and bias auditing to mitigate these limitations.

G Broader Impact

Our work aims to align large language models with formal logical principles, potentially improving the reliability and interpretability of machine reasoning. By releasing the LOGICTREE dataset and DRER code under a permissive licence, we enable researchers and practitioners to build verifiable agents for education, scientific discovery, and safety-critical auditing, where transparent deductive chains are preferable to opaque heuristics.

G.1 POSITIVE SOCIETAL OUTCOMES.

A reasoning-aligned model can serve as a didactic tutor in introductory logic courses, assist engineers in detecting faulty assumptions in software specifications, and support legal or medical professionals by highlighting which premises lead to a conclusion rather than merely producing an answer. The synthetic nature of LOGICTREE limits exposure to personal data and reduces the risk of privacy leaks.

G.2 POTENTIAL RISKS.

More persuasive and logically consistent outputs could be weaponised for misinformation or overly authoritative automation. Over-reliance on synthetic benchmarks might also hide biases that appear in real-world discourse. Furthermore, token-level reward signals expose fine-grained model behaviour, which could be exploited to reverse-engineer proprietary system prompts.

G.3 MITIGATIONS.

We distribute our resources with an explicit no-malicious-use clause, encourage downstream users to apply bias and misinformation audits, and recommend human oversight for high-stakes deployment. Future work will extend DRER to real-world corpora and incorporate human preference feedback, allowing broader yet safer adoption of reasoning-aligned reinforcement learning.




Figure 5: Training dynamics of the DAPO baseline with the DRER framework over 400 steps.




Figure 6: Comparison of DeepSeek-R1's and Our Model's accuracy and average response token on LogicTree.

H CASE STUDY: CHAIN-OF-THOUGHT QUALITY

To complement the quantitative evaluation, we present a case study based on the example shown in Table 18, which compares reasoning traces produced by different models on the same LogicTree instance. This case study highlights how DRER improves both the clarity and efficiency of chain-of-thought (CoT) reasoning.

H.1 DAPO+DRER (OURS)

Driven by DRER's token-level reward and dynamic-length mask, the trace is both concise and transparent: six orderly steps map directly onto the formal pipeline identify rule \rightarrow resolve disjunction \rightarrow propagate truth. Expressions such as "apply transitivity" signal an explicit shift toward symbolic reasoning, the behaviour DRER is designed to promote. At roughly \sim 70 tokens—far shorter than the 100+ tokens typical of vanilla DAPO—the chain remains fully verifiable, demonstrating DRER's combined gains in effectiveness and efficiency.

H.2 O4MINI

Although it yields the correct answer, steps 2–4 compress several entailments into a single sentence and omit rule names, reducing transparency. With a length of about \sim 40 tokens, it illustrates the "shallow-but-correct" pattern observed in §4.2.




Figure 7: Comparison of response lengths over training steps between DAPO and DAPO+DRER. The integration of DRER leads to a reduction in response length, indicating enhanced efficiency with concise output.




Figure 8: Comparison of response lengths over training steps between GRPO and GRPO+DRER. The integration of DRER leads to a reduction in response length, indicating enhanced efficiency with concise output.

H.3 QWEN2.5-7B-INSTRUCT

This trace shows the greatest *length drift*: more than 110 tokens, many of them descriptive filler unrelated to logic, matching the "decorative-token inflation" failure mode in our diagnostics. Despite some correct premise restatement, the model ultimately outputs Unknown, confirming that verbosity does not equal confidence.

H.4 DEEPSEEK-R1

Provides only a meta-level remark ("break down each conditional") before jumping to the answer, leaving the derivation invisible; such hidden reasoning yields the lowest Consistency Ratio in our evaluation across paraphrastic variants.




Figure 9: Prediction distribution comparison between DAPO and DAPO+DRER under Chain-of-Thought (CoT) prompting. The DAPO+DRER model demonstrates significantly higher confidence in correct answers, as shown by a strong concentration of predictions on the fully correct label set ([true, true, true]). In contrast, the baseline DAPO model produces more scattered outputs, indicating lower certainty. This highlights the effectiveness of DRER in combination with CoT reasoning for improving answer consistency and correctness.

Table 16: Comparison of Consistency Ratio on LogicTree across various logical depth.

Model	1	2	3	4	5	6	7	8	Avg.
Qwen3-235B-A22B	0.90	0.65	0.30	0.50	0.15	0.00	0.05	0.00	0.32
Deepseek-R1	0.70	0.55	0.20	0.15	0.10	0.00	0.05	0.00	0.22
Claude-3.7-Sonnet	0.65	0.35	0.00	0.00	0.00	0.00	0.00	0.00	0.12
Qwen3-8B	0.65	0.70	0.05	0.05	0.05	0.00	0.00	0.00	0.19
GPT-o4-mini	0.50	0.35	0.00	0.05	0.00	0.00	0.00	0.00	0.11
GPT-o3-mini	0.45	0.30	0.00	0.00	0.00	0.00	0.00	0.00	0.09
Qwen3-4B	0.40	0.30	0.05	0.05	0.00	0.00	0.00	0.00	0.10
Gemini-2.5-Flash-Preview	0.75	0.50	0.15	0.05	0.00	0.00	0.00	0.00	0.18
GPT-4o	0.40	0.35	0.00	0.05	0.00	0.00	0.00	0.00	0.10
Phi-4-14	0.35	0.35	0.05	0.05	0.00	0.00	0.00	0.00	0.10
Gemma-3-27B	0.25	0.20	0.00	0.00	0.00	0.00	0.00	0.00	0.06
Deepseek-v3	0.15	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.02
GPT-4o-mini	0.10	0.10	0.00	0.00	0.00	0.00	0.00	0.00	0.03
Qwen2.5-7B-Instruct-1M	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01
GRPO	0.55	0.50	0.40	0.25	0.45	0.20	0.15	0.00	0.29
DAPO	0.65	0.45	0.45	0.20	0.50	0.10	0.05	0.10	0.31
GRPO+DRER	0.65	0.50	0.25	0.25	0.25	0.10	0.00	0.00	0.25
Ours (DAPO+DRER)	0.70	0.70	0.60	0.35	0.50	0.35	0.00	0.10	0.41 ^{†0.40}

H.5 SUMMARY.

dapo+drer offers the clearest, rule-grounded, and length-controlled chain of thought; **o4mini** is concise but omits warrants and lacks length regularisation; **Qwen2.5-7B-Instruct** is verbose yet uncertain; and **DeepSeek-R1** lacks an explicit chain. The contrast underscores DRER's targeted improvements in symbolic clarity, response economy, and process—outcome alignment.




Figure 10: Prediction distribution comparison between our model (DAPO+DRER) and Qwen2.5-7B-Instruct-1M under Chain-of-Thought (CoT) prompting. The DAPO+DRER model produces highly concentrated predictions on the fully correct label ([true, true, true]), indicating strong confidence and consistency. In contrast, Qwen2.5-7B-Instruct-1M predictions are widely dispersed across incorrect and partially correct categories, reflecting lower answer certainty. This highlights the effectiveness of DRER combined with CoT in guiding the model toward accurate and confident output.

Table 17: Details of the organization and model source (model version for proprietary models, and Huggingface model name for open-source models) for the LLMs evaluated in LogicTree.

Model	Organization	Size	Notes	Source
DeepSeek-R1	DeepSeek	671B	MoE	deepseek-ai/DeepSeek-R1
DeepSeek-V3	DeepSeek	671B	MoE	deepseek-ai/DeepSeek-V3
Claude 3.7 Sonnet	Anthropic	-		claude-3-7-sonnet-20250219
Gemini 2.0 Flash Thinking Preview	Google	-		gemini-2.5-flash-preview-04-17
Gemma-3-27B	Google	27B		google/gemma-3-27b-it
Qwen3-235B-A22B	Alibaba	235B	MoE	qwen3-235b-a22b
Qwen3-30B-A3B	Alibaba	30B	MoE	qwen3-30b-a3b
Qwe3-8B	Alibaba	-		qwen3-8b
Qwen3-4B	Alibaba	-		qwen3-4b
Qwen2.5-7B-Instruct-1M	Alibaba	-	MoE	qwen2.5-7b-instruct-1m
Phi-4-14B	Microsoft	14B		microsoft/phi-4
GPT-o4-mini	OpenAI	-		04-mini-2025-04-16
GPT-o3	OpenAI	-		o3-mini-2025-01-31
GPT-4o-mini	OpenAI	-		gpt-4o-mini-2024-07-18
GPT-4o	OpenAI	-		gpt-4o-2024-11-20




Figure 11: Word frequencies of seven deductive reasoning terms explicitly mentioned in LLMs response DRER.




Figure 12: $F\beta$ -Score, Answer Rate and Precision metrics Comparison across various models.

1458 1459 1460 1461 1462 1463 Table 18: Four-model CoT Comparison (Prompt-style) 1464 1465 Four-model CoT Comparison 1466 1467 Question 1468 It is a common misconception that if ribbons relate post: How to sew trims, then 1469 titanium dioxide and zinc oxide also functioned in this fashion. 1470 1471 dapo+drer 1472 Chain-of-Thought 1473 1. Analyze the logical structure of the premises. 1474 2. Identify conditional statements and their parts. 1475 3. Resolve the disjunction via the invertebrate–coral premise. 1476 1477 4. Propagate truth to the olive-oil statement. 1478 5. Apply transitivity to obtain *ribbons* \rightarrow TiO₂/ZnO. 1479 6. Compare with the question's claim. 1480 Answer: False 1481 1482 Qwen2.5-7B-Instruct-1M 1483 Chain-of-Thought 1484 Let's analyze step by step: 1) Coral reefs need light and thus occur only 1485 in shallow waters. 2) Olive oil is heat-stable. 3) If ribbons \rightarrow TiO₂/ZnO. 1486 4) "Anus on top" is false \Rightarrow coral-reef statement true. 5) No evidence of 1487 misconception. Hence information is insufficient. 1488 Answer: Unknown 1489 o4mini 1490 Chain-of-Thought 1491 1492 1. "Anus on top" is false. 1493 2. Therefore coral-reef-light is true. 1494 3. Thus olive-oil statement is true. 1495 4. Hence conditional *ribbons* \rightarrow TiO₂/ZnO is true. 1496 1497 5. The question claims this conditional is a misconception. 1498 6. Contradiction \Rightarrow question is false. 1499 **Answer: False** 1500 1501 deepseekr1 1502 Chain-of-Thought 1503 I need to analyze consistency; I will break down each conditional and test 1504 with the premises. 1505 Answer: False 1506

I LLM USAGE

Large Language Models (LLMs) were used to aid in the writing and polishing of the manuscript. Specifically, we used an LLM to assist in refining the language, improving readability, and ensuring clarity in various sections of the paper. The model helped with tasks such as sentence rephrasing, grammar checking, and enhancing the overall flow of the text.

It is important to note that the LLM was not involved in the ideation, research methodology, or experimental design. All research concepts, ideas, and analyses were developed and conducted by the authors. The contributions of the LLM were solely focused on improving the linguistic quality of the paper, with no involvement in the scientific content or data analysis.

The authors take full responsibility for the content of the manuscript, including any text generated or polished by the LLM. We have ensured that the LLM-generated text adheres to ethical guidelines and does not contribute to plagiarism or scientific misconduct.