

# Logical Phase Transitions: Understanding Collapse in LLM Logical Reasoning

Xinglang Zhang\*, Yunyao Zhang\*, Zeliang Chen, Junqing Yu, Wei Yang, Zikai Song<sup>†</sup>

Huazhong University of Science and Technology

{normanspark, ikostar, skyesong}@hust.edu.cn

## Abstract

Symbolic logical reasoning is a critical yet underexplored capability of large language models (LLMs), providing reliable and verifiable decision-making in high-stakes domains such as mathematical reasoning and legal judgment. In this study, we present a systematic analysis of logical reasoning under controlled increases in logical complexity, and reveal a previously unrecognized phenomenon, which we term **Logical Phase Transitions**: rather than degrading smoothly, logical reasoning performance remains stable within a regime but collapses abruptly beyond a critical logical depth, mirroring physical phase transitions such as water freezing beyond a critical temperature threshold. Building on this insight, we propose **Neuro-Symbolic Curriculum Tuning**, a principled framework that adaptively aligns natural language with logical symbols to establish a shared representation, and reshapes training dynamics around phase-transition boundaries to progressively strengthen reasoning at increasing logical depths. Experiments on five benchmarks show that our approach effectively mitigates logical reasoning collapse at high complexity, yielding average accuracy gains of +1.26 in naive prompting and +3.95 in CoT, while improving generalization to unseen logical compositions. Code and data are available at: <https://github.com/AI4SS/Logical-Phase-Transitions>.

## 1 Introduction

Symbolic logical reasoning refers to the ability to draw correct conclusions by applying explicit logical rules to structured premises (Smith, 2003). It is a foundational aspect of human cognition and a core capability expected of large language models (LLMs) (Cohen et al., 2020; Vaswani et al., 2017). It supports a wide spectrum of tasks, including commonsense inference (Wang et al., 2024; Liu et al.,

\*Equal contribution. Order determined by coin flip.

<sup>†</sup>Corresponding author

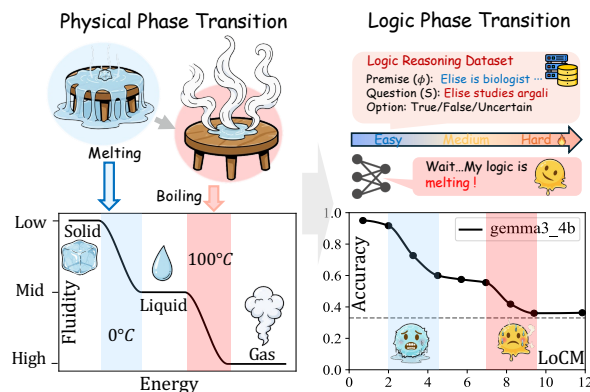


Figure 1: Overview of our work. **(Left)** Physical phase transitions: as temperature increases, matter undergoes abrupt state changes (solid → liquid at 0°C, liquid → gas at 100°C). **(Right)** Logical phase transitions: as the Logical Complexity Metric (LoCM) increases, LLM reasoning accuracy drops sharply, revealing a phase-transition-like behavior in logical reasoning.

2026), mathematical proof (Wang et al., 2023; Eisner et al., 2024; An et al., 2025), and philosophical analysis (Paul, 2013; Lin et al., 2026b).

Recent studies (Xu et al., 2024b,a; Pan et al., 2023) indicate that current LLMs perform well on logical reasoning tasks in simple scenarios, yet experience substantial degradation as reasoning becomes more challenging. While this gap is widely observed, how logical depth shapes reasoning ability in LLMs remains poorly understood, as there is still no clear characterization of reasoning behavior across increasing levels of logical complexity.

Motivated by this limitation, we analyze LLM logical reasoning from the perspective of logical complexity and uncover a previously unrecognized collapse in reasoning behavior. **(1) Measuring.** We begin by proposing the **Logical Complexity Metric (LoCM)**, a principled measure that quantifies logical difficulty in terms of symbolic structure and compositional depth. To support this complexity-aware analysis, we further construct

a logic-enhanced dataset that provides explicit first-order-logic (FOL) representations of propositions, premise sets, and reasoning chains, enabling fine-grained characterization of logical dependencies and structures. (2) *Discovery*. Based on this measuring, we reveal a previously unrecognized phenomenon, termed **Logical Phase Transitions (LPTs)**: rather than degrading smoothly, reasoning performance remains stable within specific LoCM regimes and then collapses abruptly at critical thresholds, forming multiple LPTs across the whole LoCM range. This behavior mirrors successive phase transitions in physical systems, such as the transitions from ice to water and from water to vapor (solid–liquid–gas) as temperature increases, as illustrated in Figure 1.

Building on this collapse phenomenon LPTs, we propose **Neuro-Symbolic Curriculum Tuning**, a principled framework that explicitly aligns neural representations with symbolic structures and systematically reshapes the training dynamics around these phase-transition boundaries, to mitigate reasoning collapse under increasing logical complexity. Specifically, our framework consists of two core components: (1) **Adaptive Neuro-Symbolic Alignment** establishes a shared representational space between natural language and logical symbols, enabling consistent reasoning across symbolic structures. (2) **Complexity-Aware Curriculum Optimization** organizes training into successive stages of progressively increasing logical complexity, allowing models to gradually adapt to deeper reasoning structures, particularly near critical transition regions where collapse is likely to occur.

In summary, our contributions are:

- We introduce the **Logical Complexity Metric** together with a logic-enhanced dataset, enabling fine-grained analysis of logical reasoning, and uncover a previously unrecognized collapse phenomenon, **Logical Phase Transitions**, characterized by performance collapse beyond critical complexity thresholds.
- We propose **Neuro-Symbolic Curriculum Tuning**, a principled framework that aligns neural and symbolic representations and restructures training around phase-transition boundaries, thereby mitigating reasoning collapse as logical complexity increases.

Table 1: Comparison of logical reasoning datasets. Complete metric definitions and full dataset comparisons are provided in the Appendix B.

Dataset	Creation	Scal.	Natural Language	Symbolic Rep.	Faithful Chains	Full FOL
ProntoQA	Synthetic	✓	✗	✓	✓	✗
ProofWriter	Synthetic	✓	✗	✓	✓	✗
FOLIO	Manual	✗	✓	✓	✗	✗
ProverQA	Synthetic	✓	✓	✓	✓	✗
NSA-LR(Ours)	Synthetic	✓	✓	✓	✓	✓

## 2 Related Work

A full discussion of related work appears in Appendix A; we highlight key directions here.

**LLM Logical Reasoning.** Prior work spans three main directions: (1) **Linear Reasoning (LR)** methods such as Naive Prompting and Chain-of-Thought (Wei et al., 2022); (2) **Aggregative Reasoning (AR)** approaches (Ryu et al., 2024; Zhang et al., 2023; Yao et al., 2023; Sun et al., 2024) that combine multiple reasoning trajectories; and (3) **Symbolic Reasoning (SR)** frameworks (Yang et al., 2023; Xu et al., 2024a,b; Zhang et al., 2025c) that integrate LLMs (Achiam et al., 2023; Guo et al., 2025; Zhao et al., 2023) with explicit logic modules (Pan et al., 2023). While these methods improve surface performance through guided or modular reasoning, they offer limited insight into how reasoning behavior changes with increasing logical complexity.

**Logical Reasoning Capacity Analysis.** Recent work analyzes LLM reasoning fidelity, compositionality, and failure modes (Huang and Chang, 2022; Zheng et al., 2024; Fang et al., 2025). Studies such as InfoQA (Wan et al., 2025), CoT-Valve (Ma et al., 2025), and symbolic Monte Carlo supervision (Tan et al., 2025) investigate capacity limits, chain-length controllability, and the role of symbolic trajectories, while *Apple* (Shojaee\* et al., 2025) shows that frontier LMMs suffer sharp reasoning collapses in tasks such as Tower of Hanoi as complexity increases. However, their setting focuses on procedural execution in structured puzzles, whereas we study symbolic logical reasoning under explicit propositional/FOL structure; the two works therefore differ in complexity definition, evaluation target, and intervention scope. More broadly, these analyses rely on task-specific or coarse difficulty proxies and do not provide a principled framework for quantifying logical complexity itself or characterizing collapse behaviors.

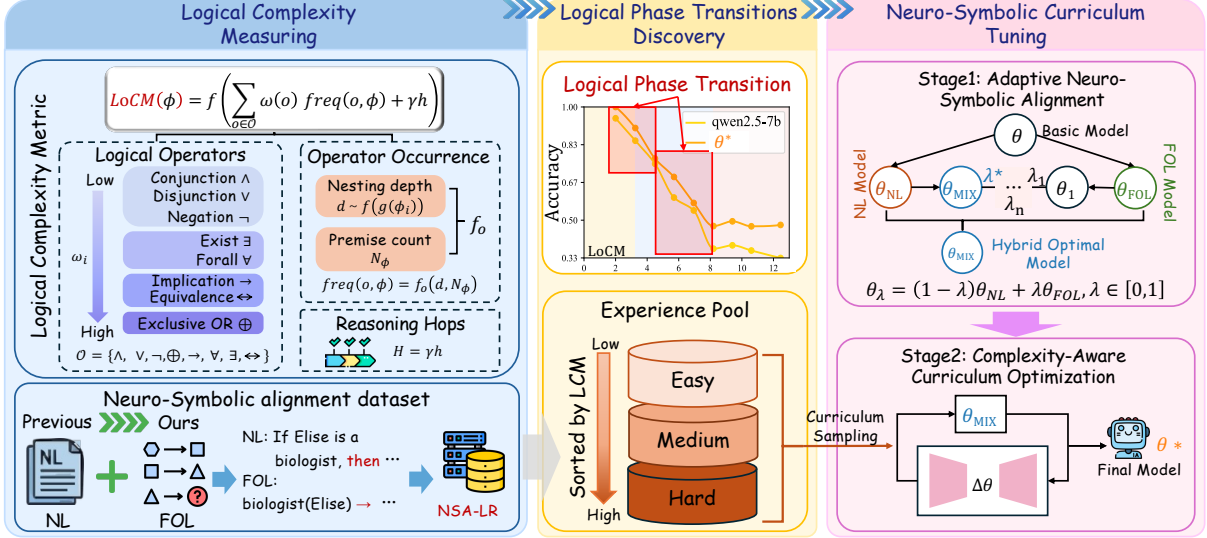


Figure 2: **Overview. Logical Complexity Measuring (left):** We construct a Neuro-Symbolic alignment dataset (NSA-LR) and define the Logical Complexity Metric (LoCM), which quantifies reasoning difficulty via logical operators, nesting depth  $d$ , premise count  $N_\phi$ , and reasoning hops  $H$ . **Logical Phase Transitions Discovery (center):** LoCM reveals a phase transition where accuracy collapses to random guessing beyond a critical complexity threshold and partitions samples into Easy, Medium, and Hard pools. **Neuro-Symbolic Curriculum Tuning (right):** Stage 1 trains a mixed-semantics model  $\theta_{\text{MIX}}$  and stage 2 applies curriculum optimization over the experience pool to obtain the final model  $\theta^*$ , mitigating reasoning collapse at high LoCM.

### 3 Methodology

We describe our method along three aspects: § 3.1 introduces logical complexity measuring, including logical complexity metric and the Neuro-Symbolic alignment dataset; § 3.2 presents the discovery of logical phase transitions; § 3.3 presents Neuro-Symbolic Curriculum Tuning. The overall framework is illustrated in Figure 2.

#### 3.1 Logical Complexity Measuring

To enable controlled analysis under increasing logical complexity, we introduce LoCM to quantify symbolic difficulty and construct a neuro-symbolic alignment dataset that provides the structured representations required to compute this metric.

##### Logical Complexity Metric

Given an input proposition  $Q$ , a model must integrate several symbolic factors such as the number of premises, the influence of logical operators, the depth of nested reasoning, and the overall length of the reasoning chain to determine its truth value. Existing complexity estimates are typically coarse and rely mainly on hop counts. To provide a more precise measure, we introduce the LoCM, which assigns each sample a scalar score that captures its logic difficulty.

**Definition 1 (LoCM).** For a reasoning instance  $\phi$  expressed in FOL, let  $\mathcal{P} = \{p_1, \dots, p_{N_\phi}\}$  denote the set of premises, where  $N_\phi = |\mathcal{P}|$ . Let  $\mathcal{O} = \{\wedge, \vee, \neg, \oplus, \rightarrow, \leftrightarrow, \forall, \exists\}$  denote the set of logical operators, including Boolean connectives and quantifiers. For each operator  $o \in \mathcal{O}$ , let  $\text{freq}(o, \phi) = f_o(d, N_\phi)$  denote its occurrence count in  $\phi$ , where  $d$  denotes the maximum syntactic nesting depth at which  $o$  appears. Let  $h$  denote the number of reasoning hops in the corresponding reasoning chain. The LoCM is defined as:

$$\text{LoCM}(\phi) = f\left(\sum_{o \in \mathcal{O}} \omega(o) \text{freq}(o, \phi) + \gamma h(\phi)\right) \quad (1)$$

where  $\omega(o)$  assigns a symbolic-complexity weight, and  $f(\cdot)$  is a monotonic transformation function used to stabilize scale and better fit empirical correlations. The metric  $\text{LoCM}(\phi)$  yields a single scalar score that quantifies the logical difficulty of a given reasoning instance.

##### Neuro-Symbolic Alignment Dataset

Building on the data-construction principles (Qi et al., 2025), we construct a **Neuro-Symbolic Alignment Dataset for Logical Reasoning (NSA-LR)** that provides paired natural language (NL) and FOL representations for every sample. All NL propositions, premises, and reasoning steps

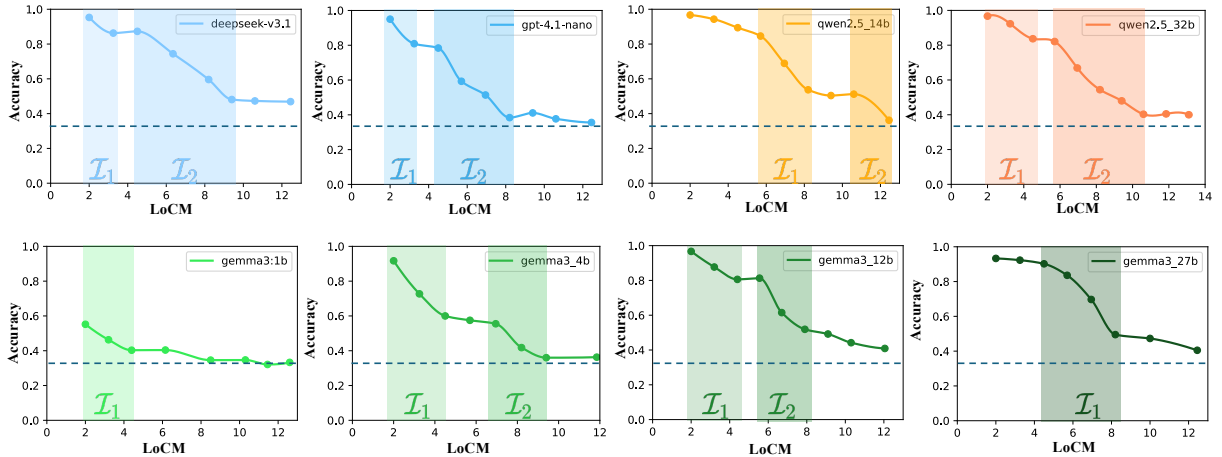


Figure 3: Logical phase-transition curves across different models. Shaded regions denote the identified transition intervals, where accuracy drops sharply as LoCM increases. The dashed line marks the 1/3 baseline corresponding to random guessing. Complete numerical results are provided in the Appendix C.2.

are translated into explicit predicates, quantifiers, connectives, and multi-step reasoning chains, following rules in Appendix Table 12. And each statement is independently translated by GPT-5 and Qwen3-Max (Yang et al., 2025a); matching outputs undergo CFG validation, while mismatches are manually adjudicated. A comparison with existing benchmarks, including ProofWriter (Tafjord et al., 2020), ProntoQA (Saparov and He, 2022), FOLIO (Han et al., 2022), and ProverQA (Qi et al., 2025), is provided in Table 1.

### 3.2 Discovery of LPT

Using LoCM, we evaluate the performance of LLMs, as shown in Figure 3, and observe a previously unrecognized phenomenon: rather than degrading smoothly, reasoning accuracy remains relatively stable over certain complexity ranges and then drops abruptly within specific regions as LoCM increases. This collapse behavior consistently appears across both open- and closed-source LLMs. More analysis is provided in Section 4.4.

Analogous to phase transitions in physics (Landau et al., 1937), where a system exhibits macroscopic changes once a control variable enters a critical region, we term this phenomenon a **Logical Phase Transition (LPT)** and formalize it as a collapse in model performance governed by logical complexity. Importantly, the transition is not characterized by a single threshold, but occurs over one or more *critical intervals*

$$\mathcal{I}_k = [\tau_{\min}^{(k)}, \tau_{\max}^{(k)}],$$

within which accuracy drops sharply as  $\text{LoCM}(\phi)$

enters the interval and stabilizes again once  $\text{LoCM}(\phi)$  exceeds its upper bound.

The discovery of LPTs indicates that direct exposure to high-complexity samples is ineffective, motivating curriculum learning (Wang et al., 2021) that organizes samples from easier to harder ones. By progressively increasing logical complexity, curriculum learning enables stable traversal of transition regions beyond the LPT regimes.

**Experience Pool.** Using the critical intervals from the phase-transition analysis, the dataset is stratified into three regimes: Easy samples satisfy  $\text{LoCM}(\phi) < \tau_{\min}^{(1)}$ ; Hard samples satisfy  $\text{LoCM}(\phi) > \tau_{\min}^{(K)}$ , where  $K$  is the number of transition intervals; all remaining samples fall into the Medium regime. This structured *Experience Pool* provides the basis for complexity-aware curriculum scheduling (Wang et al., 2021) in the Neuro-Symbolic curriculum tuning.

### 3.3 Neuro-Symbolic Curriculum Tuning

Motivated by the discovered LPTs, we design a Neuro-Symbolic Curriculum Tuning approach with two complementary components: (1) **Adaptive Neuro-Symbolic Alignment** for learning a hybrid-semantics model that aligns language and logic representations, and (2) **Complexity-Aware Curriculum**, which schedules training samples from low to high logical complexity.

#### Adaptive Neuro-Symbolic Alignment

Hybrid reasoning in LogicAgent (Zhang et al., 2025c) shows that NL and FOL provide complementary strengths: NL contributes semantic



Table 2: Performance comparison relative to Original. **Green** = improvement; **Red** = degradation; **Bold** = best.

Method	Dataset	Original	ProntoQA-tuned	ProofWriter-tuned	FOLIO-tuned	ProverQA-tuned	$\theta^*$ (Ours)
Naive	ProntoQA	55.20	55.40 (+0.20)	55.40 (+0.20)	56.20 (+1.00)	54.80 (-0.40)	<b>56.80 (+1.60)</b>
	ProofWriter	44.16	44.00 (-0.16)	44.33 (+0.17)	43.80 (-0.36)	44.30 (+0.14)	<b>44.66 (+0.50)</b>
	FOLIO	60.78	59.80 (-0.98)	59.80 (-0.98)	61.27 (+0.49)	59.80 (-0.98)	<b>62.25 (+1.47)</b>
	ProverQA	54.13	54.07 (-0.06)	53.20 (-0.93)	53.80 (-0.33)	54.73 (+0.60)	<b>55.47 (+1.34)</b>
	NSA-LR	49.55	49.09 (-0.46)	49.54 (-0.01)	49.09 (-0.46)	50.45 (+0.90)	<b>50.91 (+1.36)</b>
	<b>Naive Ave</b>	52.76	52.47 (-0.29)	52.45 (-0.31)	52.83 (+0.07)	52.82 (+0.06)	<b>54.02 (+1.26)</b>
CoT	ProntoQA	67.60	69.40 (+1.80)	70.60 (+3.00)	71.40 (+3.80)	71.20 (+3.60)	<b>72.00 (+4.40)</b>
	ProofWriter	55.16	54.33 (-0.83)	54.00 (-1.16)	56.33 (+1.17)	53.50 (-1.66)	<b>60.71 (+5.55)</b>
	FOLIO	66.17	65.68 (-0.49)	63.23 (-2.94)	63.72 (-2.45)	<b>66.67 (+0.50)</b>	65.20 (-0.97)
	ProverQA	60.70	55.70 (-5.00)	60.80 (+0.10)	60.70 (-0.00)	61.50 (+0.80)	<b>64.20 (+3.50)</b>
	NSA-LR	57.70	58.20 (+0.50)	54.70 (-3.00)	58.63 (+0.93)	59.62 (+1.92)	<b>65.00 (+7.30)</b>
	<b>CoT Ave</b>	61.47	60.66 (-0.81)	60.67 (-0.80)	62.16 (+0.69)	62.50 (+1.03)	<b>65.42 (+3.95)</b>

Table 3: Comparison with recent training-free reasoning methods on NSA-LR.

Method	Low	Medium	High	Overall
Naïve	60.6	49.8	38.6	49.6
+ NSCT	60.9 (+0.3)	50.2 (+0.4)	41.6 (+3.0)	50.9 (+1.3)
CoT	75.5	58.4	39.4	57.7
+ NSCT	84.0 (+8.5)	64.2 (+5.8)	46.8 (+7.4)	65.0 (+7.3)
ToT	82.0	64.8	42.4	63.2
+ NSCT	88.6 (+6.6)	68.3 (+3.5)	48.5 (+6.1)	68.4 (+5.2)
DetermLR	84.0	63.9	43.8	63.9
+ NSCT	87.5 (+3.5)	67.9 (+4.0)	49.2 (+5.4)	68.2 (+4.3)
SymbCoT	83.8	63.7	46.9	64.8
+ NSCT	89.2 (+5.4)	68.5 (+4.8)	50.5 (+3.6)	69.5 (+4.7)

and He, 2022) (5-hop subset), ProofWriter (Tafjord et al., 2020) (depth-5, OWA setting), FOLIO (Han et al., 2022) (full expert-curated split), and ProverQA (Qi et al., 2025) (1500 examples in total, 1–9 reasoning steps). Detailed benchmark descriptions are provided in Appendix B.1.

**Baselines.** We evaluate six model variants built upon Qwen2.5–7B: (1) the original model (*Orig.*), (2) a model fine-tuned on ProntoQA data, (3) a model fine-tuned on ProofWriter data, (4) a model fine-tuned on FOLIO data, (5) a model fine-tuned on ProverQA data, and (6) our multi-stage fine-tuned model ( $\theta^*$ ). For each model variant, we compare two prompting baselines: *Naive Prompting* and *Chain-of-Thought (CoT)* (Wei et al., 2022).

## 4.2 Evaluation Results

Tables 2 and 3 report the evaluation results, from which we draw three main observations.

**(1) Our method outperforms alternative fine-tuning strategies.** Under Naive prompting,  $\theta^*$  is

the only variant that improves over the Original model, achieving the best score on every evaluated dataset with an average gain of (+1.26). In contrast, single-dataset fine-tuning yields only marginal improvements or even degradation, indicating that LoCM-guided multi-stage training is more effective at mitigating reasoning collapse. Under CoT prompting, despite increased reasoning difficulty,  $\theta^*$  again achieves the highest average performance (+3.95) among all variants. While other methods exhibit mixed gains or notable drops on specific datasets,  $\theta^*$  consistently attains the best or near-best results.

**(2) Our method on NSA-LR generalizes to heterogeneous reasoning datasets.** Our model  $\theta^*$ , trained via Neuro-Symbolic Curriculum Tuning on the NSA-LR dataset, achieves strong performance on ProntoQA, ProofWriter, FOLIO, and ProverQA, consistently ranking among the top results. In contrast, alternative fine-tuning methods that are optimized on individual datasets exhibit limited gains on their training datasets and often suffer noticeable performance degradation when evaluated on other reasoning datasets, failing to generalize across heterogeneous tasks. These results indicate that curriculum tuning guided by neuro-symbolic alignment captures transferable reasoning structure, enabling performance improvements that extend beyond the training dataset to diverse logical forms and complexity regimes. One exception is FOLIO under CoT prompting, where the improvement is not maintained; we provide further analysis in Appendix C.4.

**(3) Our method consistently enhances diverse inference-time reasoning strategies.** As shown

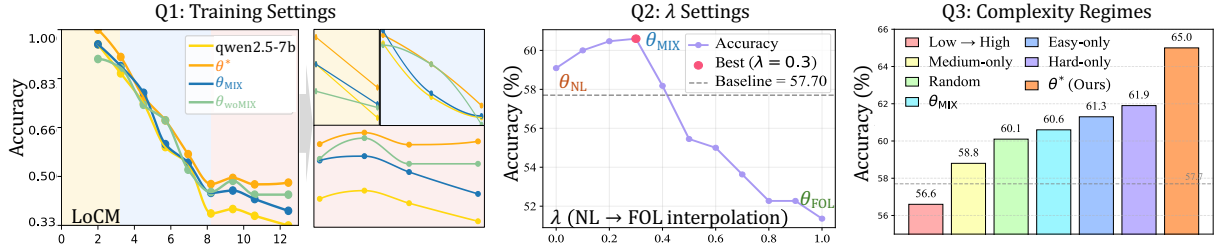


Figure 4: Ablation studies of the proposed multi-stage training framework. **(Left: Q1)** Comparison of curriculum learning and representation mixing. **(Middle: Q2)** Sensitivity analysis with respect to the mixing coefficient  $\lambda$ . **(Right: Q3)** Performance under different logical complexity regimes.

in Table 3, NSCT improves performance not only under Naive and CoT prompting, but also when combined with recent training-free reasoning methods such as ToT, DetermLR, and SymbCoT. These gains are consistent across all settings, indicating that the effect of NSCT does not depend on a specific prompting format or decoding heuristic. Notably, the improvements are most pronounced in the High-LoCM regime, where logical phase transitions are most severe: under CoT, ToT, and DetermLR, the gains in the High-LoCM band are +7.4, +6.1, and +5.4, respectively, all larger than those in the Low- or Medium-LoCM bands. Overall, these results support the view that NSCT is not merely a narrow fine-tuning heuristic tied to a specific prompting strategy. Instead, it behaves as a transferable training mechanism that consistently improves logical robustness across diverse decoding settings.

### 4.3 Ablation Study

To evaluate our method, we conduct ablation experiments to answer the following questions:

1. **Is combining neuro-symbolic alignment with curriculum optimization effective?**
2. **How sensitive is the method to the interpolation coefficient  $\lambda$ ?**
3. **How does training on different complexity regimes affect performance?**

**Q1: Impact of combining neuro-symbolic alignment and curriculum optimization.** To address Q1, we compare four variants (Figure 4-Q1): (1) the base model, (2)  $\theta_{woMIX}$ , which applies curriculum optimization without neuro-symbolic alignment, (3)  $\theta_{MIX}$ , which uses neuro-symbolic alignment without curriculum optimization, and (4) our full model  $\theta^*$ , which integrates both components.

**(1) All variants exhibit consistent LPT trends.** Across all variants, reasoning accuracy follows a

similar trend as LoCM increases, with a sharp collapse beyond critical thresholds and subsequent convergence to random guessing. This indicates that LPTs are intrinsic to the underlying model and are not eliminated by neuro-symbolic alignment or curriculum optimization.

**(2) Fine-tuning improves accuracy at matched complexity levels, with the combined model performing best.** Although the overall LPT pattern remains unchanged, all variants outperform the base model at comparable complexity levels. The full model  $\theta^*$  consistently exceeds  $\theta_{MIX}$  and  $\theta_{woMIX}$ , indicating that neuro-symbolic alignment and curriculum optimization provide complementary benefits within each complexity regime.

**Q2: Impact of balancing FOL and NL supervision ( $\lambda$ ).** We sweep  $\lambda$  between pure NL supervision ( $\lambda = 0$ ) and pure FOL supervision ( $\lambda = 1$ ) to examine how the two modalities contribute to reasoning behavior (Figure 4-Q2).

**(1) Pure NL improves over baseline, while pure FOL severely underperforms.** At  $\lambda = 0$ , training with natural language alone already surpasses the untuned baseline, indicating that NL supervision enhances reasoning by preserving contextual and discourse-level cues. In contrast, at  $\lambda = 1$ , pure FOL training causes a marked accuracy drop: while enforcing strict symbolic structure, it removes lexical semantics, pragmatic signals, and implicit assumptions required to interpret natural-language queries. As a result, the model aligns tightly with formal expressions but fails to robustly map them back to their natural-language realizations, leading to degraded semantic reasoning under FOL-only supervision.

**(2) Moderate NL-FOL mixing yields the best trade-off.** Accuracy peaks around  $\lambda = 0.3$ , where natural language continues to provide semantic grounding and flexible understanding of varied

Table 4: Accuracy–complexity correlation under different monotone transforms.

Transform	Definition	Corr.
Linear	$f(C) = C$	-0.391
Log	$f(C) = \log(C + 1)$	-0.384
Square	$f(C) = C^2$	-0.343
Inverse	$f(C) = 1/(C + 1)$	+0.248
<b>Sqrt</b>	$f(C) = \sqrt{C}$	<b>-0.399</b>

phrasings, while FOL supplies explicit operator-level structure and removes many logically inconsistent patterns. In this regime, symbolic precision from FOL and semantic richness from NL are complementary rather than competing, leading to the strongest overall performance and more robust generalization across logically complex queries.

**Q3: Effect of training under different complexity regimes.** We examine how different complexity regimes and training schedules influence model performance (Figure 4-Q3).

**(1) Single-regime fine-tuning yields limited gains.** Training on a single complexity regime yields only marginal improvements over the mixed baseline (60.6), with accuracies of 61.3 (easy-only), 58.8 (medium-only), and 61.9 (hard-only), indicating that no individual regime alone supports robust generalization. Notably, medium-only training performs the worst (58.8), reflecting less stable learning signals.

**(2) Naive low-to-high training causes forgetting, while complexity-aware curriculum is effective.** A low-to-high training strategy leads to a substantial performance drop (56.6) due to catastrophic forgetting, as earlier learned patterns are overwritten by harder instances. In contrast, our complexity-aware curriculum achieves the best performance (65.0) by retaining earlier regimes while progressively introducing higher complexity, thereby enabling consistent gains across complexity levels. More details in Appendix C.3.

#### 4.4 Analysis and Discussion

We discuss the following three aspects:

- 1. Decomposition of LoCM.** We analyze the operator weights and nonlinearity  $f(\cdot)$  in LoCM.
- 2. LPTs in LLMs.** We examine the emergence of LPTs in LLMs across LoCM ranges.
- 3. Structured prompting on LPTs.** We discuss how CoT prompting influences LPTs.

Table 5: Accuracy–complexity correlation under LoCM operator ablations.

Setting	Operator(s)	Type	Corr.
Remove	$\neg$	Negation	-0.377
	$\wedge, \vee$	Basic connectives	-0.366
	$\forall, \exists$	Quantifiers	-0.374
	$\rightarrow, \leftrightarrow$	Conditional	-0.370
	$\oplus$	XOR	-0.368
	$h$	Hops	-0.385
Only	$\neg$	Negation	-0.263
	$\wedge, \vee$	Basic connectives	-0.358
	$\forall, \exists$	Quantifiers	-0.289
	$\rightarrow, \leftrightarrow$	Conditional	-0.352
	$\oplus$	XOR	-0.325
	$h$	Hops	-0.370
Full	All operators	Complete LoCM	<b>-0.391</b>

**D1: Decomposition of LoCM.** Based on the empirical results, we make the following observations. Additional analyses are provided in Appendix C.6.

**(1) Operator-weight justification.** As shown in Table 5, assigning weights to logical operators is necessary to reflect their symbolic roles. Removing or isolating any operator type consistently weakens the correlation between complexity and accuracy, indicating that no single connective dominates reasoning difficulty. The weighted formulation is therefore required to capture interaction structures among logical operators.

**(2) Choice of nonlinearity  $f(\cdot)$ .** As shown in Table 4, we compare monotone transformations applied to aggregated operator contributions. The square-root mapping achieves the strongest alignment with accuracy, while logarithmic and square transformations introduce distortions in low- or high-complexity regions. These trends motivate selecting the square-root function as the final instantiation of  $f(\cdot)$ .

**D2: LPTs in LLMs.** Figure 3 shows several consistent phenomena across open- and closed-source LLMs; see Appendix C.2 for full results.

**(1) Logical phase transitions universally occur across models.** As LoCM increases, all evaluated models exhibit sharp accuracy drops, typically forming one or two distinct transition intervals. This demonstrates that phase-transition behavior is a fundamental property of current LLMs rather than an artifact of model family or training paradigm.

**(2) Accuracy eventually converges to the random-guessing baseline.** Once the instance LoCM surpasses the upper bound of the final transition interval, model accuracy approaches the 1/3

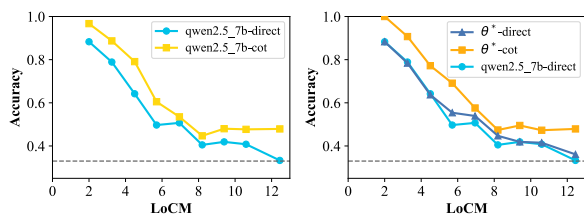


Figure 5: Effect of prompting strategies on logical phase transition (LPT) curves. **Left:** Base model (Qwen2.5-7B, without fine-tuning) under Direct and CoT prompting. **Right:** Comparison between the base model and the finetuned model under Direct and CoT prompting.

level indicated by the dashed line, implying that final predictions are essentially guesses rather than results of successful intermediate reasoning.

**(3) Larger models exhibit delayed degradation and stronger robustness.** In the Gemma 3 family (green curves), larger models maintain higher accuracy across the full complexity range, reflecting uniformly improved robustness rather than a mere horizontal shift of the transition region. In contrast, the 1B model collapses once LoCM exceeds 8 and quickly approaches the random baseline, showing its inability to handle high-complexity reasoning.

**D3: Structured prompting on LPTs.** As shown in Figure 5, both fine-tuning and CoT prompting consistently raise accuracy across logical-complexity bins. In particular, CoT yields uniform gains over Direct prompting at nearly all complexity levels, indicating improved inference efficiency and capacity utilization. However, despite these gains, the onset of the logical phase transition remains unchanged, and performance still degrades sharply once complexity exceeds the model’s intrinsic limit. Therefore, **neither fine-tuning nor structured prompting extends the complexity horizon; instead, they improve performance strictly within the same phase regime.**

## 5 Conclusion

We introduce the Logical Complexity Metric (LoCM) as a principled, model-agnostic measure for characterizing the limits of logical reasoning in large language models, and show that LLMs exhibit universal logical phase transitions marked by abrupt accuracy collapse once complexity exceeds a critical threshold. Guided by this diagnostic insight, we develop a complexity-graded reasoning corpus and a neuro-symbolic curriculum optimization that aligns training with complexity progression. Experiments demonstrate that this

approach consistently improves robustness across diverse benchmarks and prompting settings, mitigating post-transition failure and strengthening high-complexity logical reasoning.

## Limitations

**(1) Logical phase transitions persist under fine-tuning and structured prompting.** Neither fine-tuning nor CoT prompting delays, shifts, or eliminates LPTs. Although both improve accuracy within fixed complexity regimes, the critical thresholds remain unchanged, and performance still collapses sharply beyond them. This indicates that LPTs arise from intrinsic properties of current model architectures and inference mechanisms rather than from deficiencies addressable by tuning or prompting. **(2) Dependence on existing datasets and a focus on symbolic-dominant reasoning.** Our framework relies on re-annotating existing logical reasoning datasets with complete FOL representations, rather than introducing a new data generation paradigm, which constrains the diversity of logical forms to the underlying corpora. Moreover, LoCM targets reasoning expressible in FOL and does not directly cover reasoning driven by commonsense knowledge, world modeling, or probabilistic inference. **(3) Phase-transition boundaries require model-specific empirical calibration.** Although LPT phenomena are consistently observed across both open- and closed-source LLMs, the exact locations and shapes of critical intervals vary across model families and must be empirically calibrated. This suggests that LoCM captures relative rather than absolute difficulty, and that transferring phase boundaries across architectures should be done with caution.

## Acknowledgments

This work is supported by the National Natural Science Foundation of China (Numbers 62272184 and 62402189), the China Postdoctoral Science Foundation (Numbers 2024M751012, 2025T180429, and GZC20230894), the Postdoctor Project of Hubei Province (Number 2024HBB-HCXB014), the Natural Science Foundation of Hubei Province No.JCZRMS202600758, and Sponsored by CIPS-SMP-Zhipu Large Model Fund (CIPS-SMP20250306). The computation is completed in the HPC Platform of Huazhong University of Science and Technology.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Shengnan An, Xunliang Cai, Xuezhi Cao, Xiaoyu Li, Yehao Lin, Junlin Liu, Xinxuan Lv, Dan Ma, Xuanlin Wang, Ziwen Wang, and 1 others. 2025. Amo-bench: Large language models still struggle in high school math competitions. *arXiv preprint arXiv:2510.26768*.
- Steven Bird. 2006. Nltk: the natural language toolkit. In *Proceedings of the COLING/ACL 2006 interactive presentation sessions*, pages 69–72.
- Xi Chen and Min Zeng. 2025. Prototype conditioned generative replay for continual learning in NLP. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, Albuquerque, New Mexico. Association for Computational Linguistics.
- Yao Chen, Yilong Chen, Yinqi Yang, Junyuan Shang, Zhenyu Zhang, Zefeng Zhang, Shuaiyi Nie, Shuo-huan Wang, Yu Sun, Hua Wu, HaiFeng Wang, and Tingwen Liu. 2026. **Sparse growing transformer: Training-time sparse depth allocation via progressive attention looping**. *Preprint*, arXiv:2603.23998.
- Yao Chen, Jiawei Sheng, Wenyuan Zhang, and Tingwen Liu. 2025. **Improving reasoning capabilities in small models through mixture-of-layers distillation with stepwise attention on key information**. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 4952–4971, Suzhou, China. Association for Computational Linguistics.
- Michael Cohen, Badri Vellambi, and Marcus Hutter. 2020. Asymptotically unambitious artificial general intelligence. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 2467–2476.
- Ben Eisner, Yi Yang, Todor Davchev, Mel Vecerik, Jonathan Scholz, and David Held. 2024. Deep se (3)-equivariant geometric reasoning for precise placement tasks. *arXiv preprint arXiv:2404.13478*.
- Zhen Fang, Zhuoyang Liu, Jiaming Liu, Hao Chen, Yu Zeng, Shiting Huang, Zehui Chen, Lin Chen, Shanghang Zhang, and Feng Zhao. 2025. Dualvla: Building a generalizable embodied agent via partial decoupling of reasoning and action. *arXiv preprint arXiv:2511.22134*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shitong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Ruiyan Han, Zhen Fang, XinYu Sun, Yuchen Ma, Ziheng Wang, Yu Zeng, Zehui Chen, Lin Chen, Wenxuan Huang, Wei-Jie Xu, and 1 others. 2026. Unicorn: Towards self-improving unified multimodal models through self-generated supervision. *arXiv preprint arXiv:2601.03193*.
- Simeng Han, Hailey Schoelkopf, Yilun Zhao, Zhenqing Qi, Martin Riddell, Wenfei Zhou, James Coady, David Peng, Yujie Qiao, Luke Benson, and 1 others. 2022. Folio: Natural language reasoning with first-order logic. *arXiv preprint arXiv:2209.00840*.
- Zhezhen Hao, Hong Wang, Haoyang Liu, Jian Luo, Jiarui Yu, Hande Dong, Qiang Lin, Can Wang, and Jiawei Chen. 2025. Rethinking entropy interventions in rlvr: An entropy change perspective. *arXiv preprint arXiv:2510.10150*.
- Zhezhen Hao, Hong Wang, Jian Luo, Jianqing Zhang, Yuyan Zhou, Qiang Lin, Can Wang, Hande Dong, and Jiawei Chen. 2026. Recreate: Reasoning and creating domain agents driven by experience. *arXiv preprint arXiv:2601.11100*.
- Yangliu Hu, Zikai Song, Na Feng, Yawei Luo, Junqing Yu, Yi-Ping Phoebe Chen, and Wei Yang. 2025. Sf2t: Self-supervised fragment finetuning of video-llms for fine-grained understanding. *arXiv preprint arXiv:2504.07745*.
- Jie Huang and Kevin Chen-Chuan Chang. 2022. Towards reasoning in large language models: A survey. *arXiv preprint arXiv:2212.10403*.
- Lev Davidovich Landau and 1 others. 1937. On the theory of phase transitions. *Zh. eksp. teor. Fiz*, 7(19-32):926.
- Wenbing Li, Zikai Song, Hang Zhou, Yunyao Zhang, Junqing Yu, and Wei Yang. 2025a. Lora-mixer: Coordinate modular lora experts through serial attention routing. *arXiv preprint arXiv:2507.00029*.
- Wenbing Li, Hang Zhou, Junqing Yu, Zikai Song, and Wei Yang. 2024. Coupled mamba: Enhanced multimodal fusion with coupled state space model. *arXiv preprint arXiv:2405.18014*.
- Yu Li, Tian Lan, and Zhengling Qi. 2025b. Inspo: Unlocking intrinsic self-reflection for llm preference optimization. *arXiv preprint arXiv:2512.23126*.
- Yu Li, Tian Lan, and Zhengling Qi. 2026a. When right meets wrong: Bilateral context conditioning with reward-confidence correction for grpo. *arXiv preprint arXiv:2603.13134*.
- Yu Li, Rui Miao, Zhengling Qi, and Tian Lan. 2026b. Arise: Agent reasoning with intrinsic skill evolution in hierarchical reinforcement learning. *arXiv preprint arXiv:2603.16060*.
- Honglin Lin, Zheng Liu, Yun Zhu, Chonghan Qin, Juekai Lin, Xiaoran Shang, Conghui He, Wentao Zhang, and Lijun Wu. 2026a. Mmfinereason: Closing the multimodal reasoning gap via open data-centric methods. *arXiv preprint arXiv:2601.21821*.

- Honglin Lin, Chonghan Qin, Zheng Liu, Qizhi Pei, Yu Li, Zhanping Zhong, Xin Gao, Yanfeng Wang, Conghui He, and Lijun Wu. 2026b. Scientific image synthesis: Benchmarking, methodologies, and downstream utility. *arXiv preprint arXiv:2601.17027*.
- Guoming Ling, Zhongzhan Huang, Yupei Lin, Junxin Li, Shanshan Zhong, Hefeng Wu, and Liang Lin. 2026. **Neural chain-of-thought search: Searching the optimal reasoning path to enhance large language models**. *Preprint*, arXiv:2601.11340.
- Zheng Liu, Honglin Lin, Chonghan Qin, Xiaoyang Wang, Xin Gao, Yu Li, Mengzhang Cai, Yun Zhu, Zhanping Zhong, Qizhi Pei, and 1 others. 2026. Chartverse: Scaling chart reasoning via reliable programmatic synthesis from scratch. *arXiv preprint arXiv:2601.13606*.
- Zheng Liu, Mengjie Liu, Jingzhou Chen, Jingwei Xu, Bin Cui, Conghui He, and Wentao Zhang. 2025. Fusion: Fully integration of vision-language representations for deep cross-modal understanding. *arXiv preprint arXiv:2504.09925*.
- Xinyin Ma, Guangnian Wan, Runpeng Yu, Gongfan Fang, and Xinchao Wang. 2025. **CoT-valve: Length-compressible chain-of-thought tuning**. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6025–6035, Vienna, Austria. Association for Computational Linguistics.
- Liangming Pan, Alon Albalak, Xinyi Wang, and William Yang Wang. 2023. Logic-lm: Empowering large language models with symbolic solvers for faithful logical reasoning. *arXiv preprint arXiv:2305.12295*.
- Richard W Paul. 2013. Critical and reflective thinking: A philosophical perspective. In *Dimensions of thinking and cognitive instruction*, pages 445–494. Routledge.
- Chengwen Qi, Ren Ma, Bowen Li, He Du, Binyuan Hui, Jinwang Wu, Yuanjun Laili, and Conghui He. 2025. Large language models meet symbolic provers for logical reasoning evaluation. *arXiv preprint arXiv:2502.06563*.
- Bertrand Russell. 2020. *Principles of mathematics*. Routledge.
- Hyun Ryu, Gyeongman Kim, Hyemin S Lee, and Eunho Yang. 2024. Divide and translate: Compositional first-order logic translation and verification for complex logical reasoning. *arXiv preprint arXiv:2410.08047*.
- Abulhair Saparov and He He. 2022. Language models are greedy reasoners: A systematic formal analysis of chain-of-thought. *arXiv preprint arXiv:2210.01240*.
- Parshin Shojaee\*, Iman Mirzadeh\*, Keivan Alizadeh, Maxwell Horton, Samy Bengio, and Mehrdad Farajtabar. 2025. **The illusion of thinking: Understanding the strengths and limitations of reasoning models via the lens of problem complexity**. In *NeurIPS*.
- Peter Smith. 2003. *An introduction to formal logic*. Cambridge University Press.
- Zikai Song, Ying Tang, Run Luo, Lintao Ma, Junqing Yu, Yi-Ping Phoebe Chen, and Wei Yang. 2024a. Autogenic language embedding for coherent point tracking. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 2021–2030.
- Zikai Song, Ying Tang, Run Luo, Lintao Ma, Junqing Yu, Yi-Ping Phoebe Chen, and Wei Yang. 2024b. Autogenic language embedding for coherent point tracking. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 2021–2030.
- Zikai Song, Junqing Yu, Yi-Ping Phoebe Chen, Wei Yang, and Xinchao Wang. 2026. **Hypergraph-state collaborative reasoning for multi-object tracking**. *Preprint*, arXiv:2604.12665.
- Hongda Sun, Weikai Xu, Wei Liu, Jian Luan, Bin Wang, Shuo Shang, Ji-Rong Wen, and Rui Yan. 2024. Determlr: Augmenting llm-based logical reasoning from indeterminacy to determinacy. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9828–9862.
- Oyvind Taffjord, Bhavana Dalvi Mishra, and Peter Clark. 2020. Proofwriter: Generating implications, proofs, and abductive statements over natural language. *arXiv preprint arXiv:2012.13048*.
- Xingwei Tan, Marco Valentino, Mahmud Elahi Akhter, Maria Liakata, and Nikolaos Aletras. 2025. **Enhancing logical reasoning in language models via symbolically-guided Monte Carlo process supervision**. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 31874–31888, Suzhou, China. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Kaiyang Wan, Lang Gao, Honglin Mu, Preslav Nakov, Yuxia Wang, and Xiuying Chen. 2025. A fano-style accuracy upper bound for llm single-pass reasoning in multi-hop qa. *arXiv preprint arXiv:2509.21199*.
- Chengbing Wang, Yang Zhang, Wenjie Wang, Xiaoyan Zhao, Fuli Feng, Xiangnan He, and Tat-Seng Chua. 2026a. **Think-while-generating: On-the-fly reasoning for personalized long-form generation**. *Preprint*, arXiv:2512.06690.
- Chengbing Wang, Wuqiang Zheng, Yang Zhang, Fengbin Zhu, Junyi Cheng, Yi Xie, Wenjie Wang, and Fuli Feng. 2026b. **Perm: Psychology-grounded empathetic reward modeling for large language models**. *Preprint*, arXiv:2601.10532.

- Ke Wang, Houxing Ren, Aojun Zhou, Zimu Lu, Sichun Luo, Weikang Shi, Renrui Zhang, Linqi Song, Mingjie Zhan, and Hongsheng Li. 2023. Math-coder: Seamless code integration in llms for enhanced mathematical reasoning. *arXiv preprint arXiv:2310.03731*.
- Weiqi Wang, Tianqing Fang, Chunyang Li, Haochen Shi, Wenxuan Ding, Baixuan Xu, Zhaowei Wang, Jiaxin Bai, Xin Liu, Cheng Jiayang, Chunkit Chan, and Yangqiu Song. 2024. **CANDLE: Iterative conceptualization and instantiation distillation from large language models for commonsense reasoning**. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2351–2374, Bangkok, Thailand. Association for Computational Linguistics.
- Xin Wang, Yudong Chen, and Wenwu Zhu. 2021. A survey on curriculum learning. *IEEE transactions on pattern analysis and machine intelligence*, 44(9):4555–4576.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Ludwig Wittgenstein. 2023. *Tractatus logico-philosophicus*.
- Jundong Xu, Hao Fei, Meng Luo, Qian Liu, Liangming Pan, William Yang Wang, Preslav Nakov, Mong-Li Lee, and Wynne Hsu. 2024a. Aristotle: Mastering logical reasoning with a logic-complete decompose-search-resolve framework. *arXiv preprint arXiv:2412.16953*.
- Jundong Xu, Hao Fei, Liangming Pan, Qian Liu, Mong-Li Lee, and Wynne Hsu. 2024b. Faithful logical reasoning via symbolic chain-of-thought. *arXiv preprint arXiv:2405.18357*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025a. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- An Yang, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoyan Huang, Jiandong Jiang, Jianhong Tu, Jianwei Zhang, Jingren Zhou, and 1 others. 2025b. Qwen2. 5-1m technical report. *arXiv preprint arXiv:2501.15383*.
- Yuan Yang, Siheng Xiong, Ali Payani, Ehsan Shareghi, and Faramarz Fekri. 2023. Harnessing the power of large language models for natural language to first-order logic translation. *arXiv preprint arXiv:2305.15541*.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems*, 36:11809–11822.
- Liliang Ye, Yunyao Zhang, Yafeng Wu, Yi-Ping Phoebe Chen, Junqing Yu, Wei Yang, and Zikai Song. 2025. Mvp: Winning solution to smp challenge 2025 video track. *arXiv preprint arXiv:2507.00950*.
- Yu Zeng, Wenxuan Huang, Zhen Fang, Shuang Chen, Yufan Shen, Yishuo Cai, Xiaoman Wang, Zhenfei Yin, Lin Chen, Zehui Chen, and 1 others. 2026. Vision-deepresearch benchmark: Rethinking visual and textual search for multimodal large language models. *arXiv preprint arXiv:2602.02185*.
- Peixuan Zhang, Zijian Jia, Kaiqi Liu, Shuchen Weng, Si Li, and Boxin Shi. 2025a. Stage: Storyboard-anchored generation for cinematic multi-shot narrative. *arXiv preprint arXiv:2512.12372*.
- Peixuan Zhang, Chang Zhou, Ziyuan Zhang, Hualuo Liu, Chunjie Zhang, Jingqi Liu, Xiaohui Zhou, Xi Chen, Shuchen Weng, Si Li, and Boxin Shi. 2026. A benchmark and multi-agent system for instruction-driven cinematic video compilation. *arXiv preprint arXiv:2604.10456*.
- Yifan Zhang, Jingqin Yang, Yang Yuan, and Andrew Chi-Chih Yao. 2023. Cumulative reasoning with large language models. *arXiv preprint arXiv:2308.04371*.
- Yunyao Zhang, Zikai Song, Hang Zhou, Wenfeng Ren, Yi-Ping Phoebe Chen, Junqing Yu, and Wei Yang. 2025b. **ga – s<sup>3</sup>: Comprehensive social network simulation with group agents**. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 8950–8970, Vienna, Austria. Association for Computational Linguistics.
- Yunyao Zhang, Xinglang Zhang, Junxi Sheng, Wenbing Li, Junqing Yu, Wei Yang, and Zikai Song. 2025c. **From ambiguity to verdict: A semiotic-grounded multi-perspective agent for llm logical reasoning**. *Preprint*, arXiv:2509.24765.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, and 1 others. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 1(2).
- Zi'ou Zheng, Christopher Malon, Martin Renqiang Min, and Xiaodan Zhu. 2024. Exploring the role of reasoning structures for constructing proofs in multi-step natural language reasoning with large language models. *arXiv preprint arXiv:2410.08436*.
- Yixiao Zhou, Yang Li, Dongzhou Cheng, Hehe Fan, and Yu Cheng. 2026. Look inward to explore outward: Learning temperature policy from llm internal states via hierarchical rl. *arXiv preprint arXiv:2602.13035*.

## Appendix

The appendix is organized as follows. § **A** reviews additional related work. § **B** describes datasets and baselines. § **C** provides experimental details. § **D** presents the formal setup. § **E** discusses future directions. § **F** includes full prompting templates.

### The Usage of LLM

In accordance with ACL guidelines, we used large language models solely for writing assistance and language refinement.

## A Related Work

### A.1 Benchmarks for Logical Reasoning

Existing logical reasoning benchmarks fall broadly into two categories. **Synthetic symbolic benchmarks** such as PrOntoQA (Saparov and He, 2022) and ProofWriter (Tafjord et al., 2020) evaluate deduction in controlled symbolic settings using rule-based or template-driven generation. **Natural-language logical benchmarks** such as FOLIO (Han et al., 2022) and ProverQA (Qi et al., 2025) provide English scenarios paired with verified reasoning chains produced via theorem-proving pipelines. Together, these datasets offer structured environments for assessing the logical validity of model outputs.

**Open Challenges.** Despite their utility, existing benchmarks often emphasize *logical form* in relatively controlled settings, typically using concrete and unambiguous predicates with fixed semantics. They may not provide complete end-to-end symbolic representations (e.g., full FOL contexts paired with verified reasoning chains) under a unified schema at scale. As a result, they offer limited support for examining how reasoning difficulty scales with increasing *logical complexity*, motivating our construction of a complexity-graded Neuro-Symbolic alignment dataset.

### A.2 LLM-Based Logical Reasoning

**Reasoning Methodologies.** Prior work on LLM (Liu et al., 2025; Hu et al., 2025; Ye et al., 2025; Zeng et al., 2026; Han et al., 2026; Li et al., 2025a) logical reasoning can be grouped into three main directions: (1) **Linear Reasoning (LR)** methods such as Naive Prompting and Chain-of-Thought (Wei et al., 2022; Chen et al., 2025; Lin et al., 2026a), which generate step-by-step natural-language explanations; (2) **Aggregative Reasoning (AR)** approaches (Ryu et al., 2024;

Zhang et al., 2023; Yao et al., 2023; Sun et al., 2024; Song et al., 2026, 2024a) that combine multiple reasoning trajectories through sampling, search, or majority voting; and (3) **Symbolic Reasoning (SR)** frameworks (Yang et al., 2023; Xu et al., 2024a,b; Pan et al., 2023) that integrate LLMs (Yang et al., 2025b; Wang et al., 2026b,a; Li et al., 2025b) with explicit logic modules or FOL-based reasoning. Together, these methodologies improve performance on a wide range of reasoning benchmarks through guided, multi-path, or logic-augmented reasoning.

**Open Challenges.** Despite these advances, many methodologies primarily improve prompting, reasoning-time scaffolds, or external reasoning modules (Ling et al., 2026; Chen and Zeng, 2025), rather than explicitly characterizing or strengthening robustness across varying levels of *logical complexity*. As a result, most approaches improve accuracy on fixed benchmarks but provide limited insight into how internal reasoning behavior changes along a fine-grained complexity continuum, which is the central focus of our work.

### A.3 Logical Reasoning Capacity Analysis

**Existing Capacity and Fidelity Analyses.** A substantial body of work investigates the reasoning (Li et al., 2026b,a; Hao et al., 2025, 2026; Zhang et al., 2025a) fidelity and compositionality of LLMs (Zhang et al., 2025b; Li et al., 2024; Song et al., 2024b; Chen et al., 2026; Zhang et al., 2026; Zhou et al., 2026). Prior studies show that models often hallucinate intermediate steps or deviate from valid derivations even when answers are correct (Huang and Chang, 2022; Zheng et al., 2024), revealing instability in multi-step reasoning and difficulty maintaining consistent logical chains. More recent capacity-oriented research examines how reasoning behavior changes as tasks grow more complex. InfoQA (Wan et al., 2025) formalizes a single-pass accuracy upper bound for multi-hop QA, demonstrating that performance can collapse once information load exceeds model capacity. CoT-Valve (Ma et al., 2025) shows that reasoning chains can be compressed or expanded through latent control, highlighting sensitivity to chain length and internal reasoning load. Symbolic Monte Carlo supervision (Tan et al., 2025) leverages symbolic trajectories to stabilize intermediate reasoning.

*Apple* (Shojaee\* et al., 2025) further demonstrates abrupt accuracy collapses in structured puzzles such as Tower of Hanoi as complexity in-

Table 6: Comparison of dataset statistics. NSA-LR scales up the ProverQA framework and adds verifiable symbolic reasoning chains. †*Statistics for ProofWriter refer to the depth-5 OWA subset used in our experiments.*

Dataset	Method	Scalability	Rich NL	Sym. Rep.	Verif. Chain (NL)	Verif. Chain (Sym.)	NL-Sym Mapping	Test Set Size	Train Set Size
ProverQA	Synthetic	✓	✓	✓	✓	✗	✓	1,500	5,000
FOLIO	Manual	✗	✓	✓	✗	✗	✗	204	1,000
ProofWriter†	Synthetic	✓	✗	✗	✓	✗	✗	600	5,000
ProntoQA	Synthetic	✓	✗	✓	✓	✗	✗	500	2,880
NSA-LR (Ours)	Synthetic	✓	✓	✓	✓	✓	✓	1,500	<b>15,800</b>

creases.

**Open Challenges.** Although informative, existing analyses often rely on coarse difficulty indicators such as hop count, chain length, or rule depth, offering only partial insight into how reasoning performance evolves under more nuanced forms of complexity. They typically lack a unified framework for quantifying *logical complexity* and for characterizing how LLM behavior changes across fine-grained complexity levels. Our work complements these analyses by introducing an LoCM and systematically evaluating model performance along a controlled complexity continuum.

## B Dataset Details

### B.1 Dataset Descriptions

We describe the logical reasoning datasets used in our evaluation, which together provide a systematic and fair testbed for assessing model reasoning across diverse logical structures and deductive complexity levels.

**FOLIO** FOLIO is an expert-curated benchmark for natural language reasoning with first-order logic (FOL) annotations. Compared with synthetic datasets, it offers greater linguistic diversity and naturalness while preserving strict logical validity. The task requires determining whether a conclusion follows from a set of premises, involving phenomena such as quantification and negation, and thus serves as a key testbed for evaluating the integration of semantic understanding and formal reasoning.

**ProofWriter** ProofWriter is a synthetic benchmark designed to evaluate systematic neural deduction under the Open World Assumption. It features rule-based natural language problems requiring multi-step reasoning over conjunctions and disjunctions. Its explicit proof supervision enables evaluation of both final-answer correctness and the validity of intermediate deductive steps.

**ProntoQA** ProntoQA is a synthetic benchmark for diagnosing deductive reasoning in a highly controlled symbolic setting. Each instance contains premises, a step-by-step chain of thought, and a conclusion, allowing precise analysis of multi-hop inference while minimizing linguistic ambiguity.

**ProverQA** ProverQA, built on ProverGen, combines LLM-generated linguistic diversity with formally verified theorem-proving chains. It is designed to assess deductive consistency and symbolic-linguistic alignment, with stratified reasoning depths that support scalable evaluation across increasing inference complexity.

**NSA-LR (Ours).** To enable fine-grained logical complexity analysis, we construct a Neuro-Symbolic alignment dataset extending the ProverGen framework used in ProverQA. Unlike the original framework, which focuses on verifying reasoning chains, our key modification exhaustively translates all natural language components, including propositions, premise sets, and intermediate reasoning steps, into explicit first-order logic. By strictly aligning each sentence with its logical form using unified translation rules, the dataset provides the symbolic grounding required to compute LoCM and systematically analyze how structural logical complexity affects model performance.

### B.2 Dataset Statistics and Comparison

Table 6 presents the statistics and key characteristics of the datasets. While FOLIO provides mappings between natural language and FOL, it lacks verifiable reasoning chains. In contrast, our NSA-LR dataset (an extension of ProverQA) uniquely combines scalability, linguistic richness, and fully verifiable symbolic reasoning chains, supporting rigorous logic-consistency analysis.

### B.3 Example of NSA-LR Data Instance

For a detailed walkthrough of a high-complexity reasoning instance ( $LoCM \approx 7.25$ ) that demon-

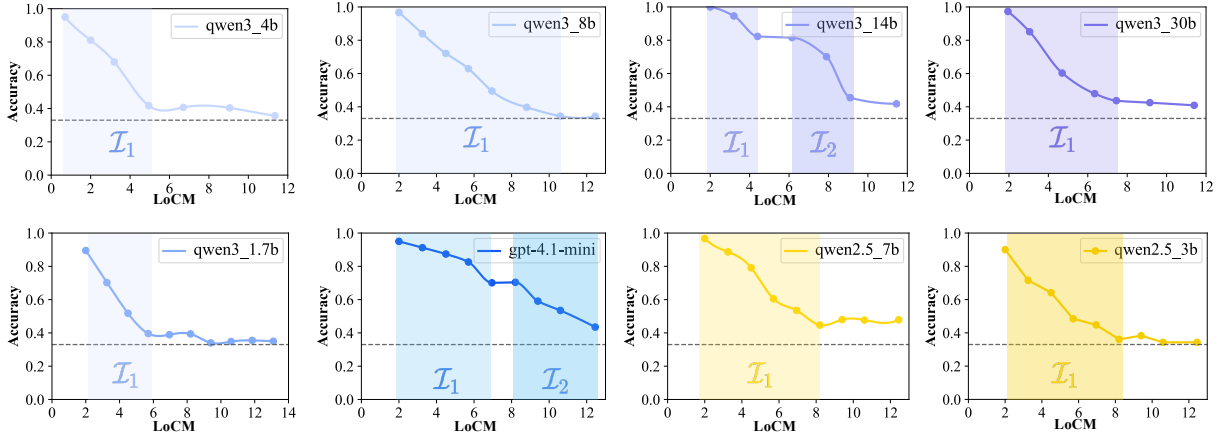


Figure 6: Supplementary logical phase-transition curves for additional models. Shaded bands indicate the detected transition intervals (e.g.,  $\mathcal{I}_1$ ,  $\mathcal{I}_2$ ), where accuracy exhibits an abrupt decline as LoCM increases. The dashed horizontal line denotes the  $1/3$  random-guess baseline.

strates the strict alignment between natural language and first-order logic across a multi-step deduction chain, please refer to the case study presented in Figure 12.

## C Experimental Details

**Overview.** This section provides detailed experimental settings and extended analyses to ensure reproducibility and support the main findings. We describe the evaluation protocol and implementation details (§C.1), report full results across *Model*  $\times$  *Complexity Bins* (§C.2), analyze the necessity of curriculum design (§C.3), include an analysis of the FOLIO CoT anomaly (§C.4), and present diagnostic analyses of the complexity metric and logical phase transitions (§C.5, §C.6). We further examine potential confounding factors, including complexity–premise count relations (§C.7) and reasoning effort (§C.8).

### C.1 Evaluation and Implementation Details

**Evaluation and answer extraction.** All models are prompted to output a JSON-formatted response containing the reasoning trace and final answer. When malformed outputs occur, we apply a robust post-processing strategy to recover predictions by prioritizing explicit final-answer fields when available, and otherwise extracting the answer from the trailing completion using heuristic patterns.

**Quantized evaluation for fair comparison.** Due to computational constraints, all open-source models used in our evaluation are tested under llama.cpp Q4 quantization. To ensure fair comparison, we quantize our fine-tuned checkpoint to

the same Q4 setting via llama.cpp before reporting results.

**Training setting.** We fine-tune Qwen2.5-7B-Instruct using parameter-efficient LoRA on a single Tesla V100S-PCIE-32GB GPU. Training is performed for one epoch with a learning rate of  $1 \times 10^{-4}$  and an effective batch size of 48, implemented with a per-device batch size of 1 and `gradient_accumulation_steps=48`. For LoRA, we set the rank to  $r = 8$  and the scaling parameter to  $\alpha = 16$ .

### C.2 Full Results

This section extends the analysis beyond the models reported in the main text. The supplementary phase-transition curves shown in Figure 6 include additional model families and sizes, providing complementary evidence for the observed logical phase-transition behavior across contemporary LLMs.

**Model-by-Bin Accuracy Table.** For detailed numerical comparison, Table 7 reports the accuracy of each tested model across the nine complexity bins. This tabular view enables precise inspection of performance at specific complexity levels and facilitates the identification of patterns that may not be immediately visible from the phase-transition curves alone.

**Transition patterns vary within a model family.** Even within the same model family, both the *number* of detected critical intervals and the *onset* of logical phase transitions can vary substantially across parameter scales. The Gemma series illustrates this variability: its 4B and 12B models exhibit two critical intervals, whereas the 1B and 27B models show only one. This observation indi-

Table 7: Full results (accuracy) across model families and complexity bins. Performance systematically degrades from Bin 1 to Bin 9 as task complexity increases.

Model	Complexity Bins								
	1	2	3	4	5	6	7	8	9
<i>Qwen 2.5 Series</i>									
Qwen2.5-3B	0.900	0.716	0.642	0.485	0.448	0.361	0.383	0.343	0.343
Qwen2.5-7B	0.950	0.851	0.749	0.600	0.543	0.374	0.389	0.366	0.333
Qwen2.5-14B	0.967	0.943	0.894	0.846	0.689	0.537	0.504	0.512	0.362
Qwen2.5-32B	0.967	0.923	0.837	0.821	0.669	0.544	0.480	0.403	0.404
<i>Qwen 3 Series</i>									
Qwen3-1.7B	0.895	0.702	0.518	0.396	0.389	0.394	0.340	0.348	0.353
Qwen3-4B	0.811	0.680	0.464	0.349	0.407	0.525	0.375	0.232	0.357
Qwen3-8B	0.967	0.839	0.721	0.629	0.495	0.397	0.344	0.370	0.333
Qwen3-14B	1.000	0.946	0.823	0.841	0.790	0.701	0.455	0.320	0.343
Qwen3-30B	0.973	0.851	0.634	0.565	0.479	0.437	0.373	0.425	0.409
<i>Gemma Series</i>									
Gemma3-1B	0.552	0.463	0.403	0.433	0.382	0.346	0.347	0.321	0.333
Gemma3-4B	0.917	0.727	0.600	0.575	0.555	0.418	0.360	0.314	0.500
Gemma3-12B	0.967	0.877	0.806	0.813	0.615	0.519	0.492	0.442	0.409
Gemma3-27B	0.933	0.923	0.902	0.836	0.697	0.495	0.508	0.419	0.405
<i>Other Models</i>									
GPT-4.1 Nano	0.950	0.809	0.784	0.593	0.514	0.384	0.411	0.377	0.356
GPT-4.1 Mini	0.950	0.912	0.874	0.826	0.701	0.704	0.591	0.535	0.435
DeepSeek V3.1	0.953	0.863	0.873	0.779	0.718	0.596	0.481	0.473	0.469

cates that transition dynamics are not governed by a simple monotonic function of model size.

**Scaling does not guarantee robustness in high-complexity regimes.** Increasing parameter count does not necessarily lead to higher overall accuracy or delayed phase transitions. For example, Qwen2.5-14B achieves higher overall accuracy (70.2%) than its 32B counterpart (68.3%) and consistently outperforms it in the mid-to-high complexity range (Bin 5–Bin 8). This suggests that scaling alone is insufficient for handling high-LoCM instances, and that phase-transition behavior may instead be influenced by factors such as pretraining data composition, instruction tuning, and optimization dynamics.

**Scaling consistently benefits low-complexity performance.** In contrast to the non-monotonic behavior observed at higher LoCMs, scaling yields stable improvements in the low-complexity regime. Moving from 1–4B to 12–14B models results in substantial and consistent gains in Bin 1–Bin 3. These results indicate that while scaling effectively improves performance on elementary instances, it

exhibits diminishing returns or instability when models approach the complexity threshold.

### C.3 Necessity of Curriculum Training

**Single-regime training is insufficient for robust reasoning across logical complexity levels.** To examine this, we conduct controlled fine-tuning experiments in which training is restricted to a single complexity interval. As shown in Figure 7, models trained exclusively on high-complexity instances (*hard-only*) achieve better accuracy in the high-complexity region but suffer substantial degradation on low- and medium-complexity samples, indicating catastrophic forgetting of reasoning patterns associated with simpler regimes. By contrast, training restricted to low- or medium-complexity instances does not improve performance on high-complexity samples and can even further degrade that regime. Taken together, these results show that single-regime training cannot sustain consistent reasoning performance across complexity levels, motivating the use of complexity-aware curriculum training.

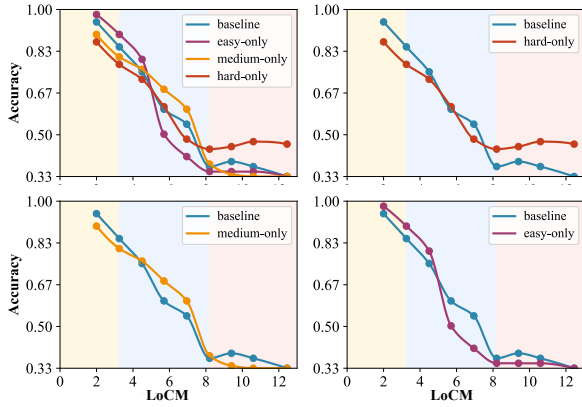


Figure 7: Effect of training on isolated complexity regimes.

#### C.4 Analysis of the FOLIO CoT Anomaly

As shown in Table 2, FOLIO is the main exception to the otherwise consistent gains of NSCT under CoT prompting. We do not view this as evidence of instability specific to NSCT, since similar drops on FOLIO also appear for several models fine-tuned on formal reasoning datasets, including the ProntoQA-tuned, ProofWriter-tuned, and FOLIO-tuned variants. This suggests that the anomaly more likely reflects a broader mismatch between formal-logic-oriented tuning and the reasoning demands of FOLIO.

Although FOLIO is annotated with first-order logic, it places heavier demands on natural-language interpretation and commonsense grounding than the more structurally controlled benchmarks. In contrast, NSCT is primarily designed to strengthen formal validity under increasing logical complexity. We therefore interpret the FOLIO result as a boundary case of the current framework rather than a contradiction of the main findings, and view it as a motivation for incorporating more commonsense-sensitive reasoning data in future work.

#### C.5 Logical Phase Transition Analysis

In this section, we further probe the logical phase transition phenomenon by examining whether commonly used interventions—Supervised Fine-Tuning (SFT), prompting strategies, and model scaling—can shift or delay the collapse region. Figure 8 summarizes the resulting comparisons.

**Observation 1: Supervised Fine-Tuning yields local accuracy gains without substantially delaying the phase transition.** As shown in the right panel of Figure 8, SFT improves performance

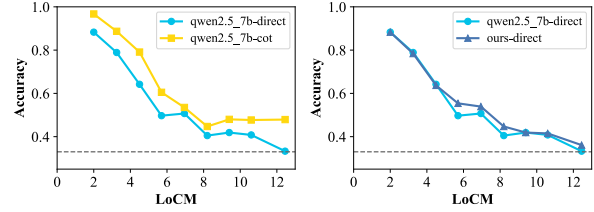


Figure 8: Logical phase-transition curves under different interventions. **Left:** Direct vs. CoT prompting on the base model. **Right:** base vs. fine-tuned model under Direct prompting.

across several complexity bins, while the sharp degradation still occurs within approximately the same LoCM range. This indicates that although SFT better aligns the model with the task format and reduces pre-threshold errors, it does not significantly shift the critical interval itself. In effect, fine-tuning primarily induces a *vertical lift* in accuracy rather than a *horizontal extension* of the complexity horizon.

**Observation 2: Advanced prompting strategies increase accuracy but do not extend the complexity horizon.** The left panel of Figure 8 shows that switching from Direct to CoT prompting consistently improves accuracy across bins, while the transition onset remains stable. This suggests that CoT acts as an inference-time facilitator, enabling the model to better utilize its existing capacity, but leaving the high-complexity failure mode intact once logical difficulty exceeds that capacity. Prompting therefore raises the overall accuracy profile without reliably shifting the phase-transition boundary.

**Observation 3: Scaling strengthens low-complexity robustness but does not eliminate high-complexity collapse.** Referring to the scaling analysis in Figure 6, increasing model size consistently stabilizes performance in low-LoCM regimes, forming a more robust pre-threshold plateau. However, the phase-transition behavior persists: even larger models eventually experience a rapid collapse once the complexity threshold is crossed. This pattern suggests that scaling improves performance on simple instances but does not prevent collapse under high logical complexity.

#### C.6 LoCM Diagnostics and Calibration

**Design Principle: Structural Orthogonality and Expressive Sufficiency.** We construct the operator basis  $\mathcal{O} = \{\neg, \wedge, \vee, \rightarrow, \leftrightarrow, \oplus, \forall, \exists\}$  not merely as a syntactic collection, but as a set of *orthogo-*

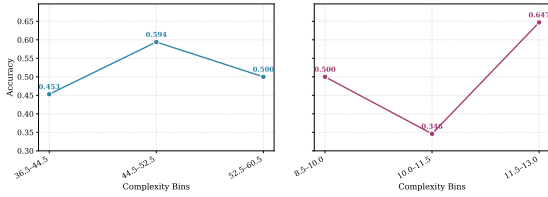


Figure 9: Accuracy vs. single-operator complexity proxies (Left:  $\wedge/\vee$ ; Right:  $\forall/\exists$ ).

*nal reasoning primitives* representing distinct cognitive operations in symbolic deduction. This design is inspired by Russell’s logical atomism, which views valid reasoning as discovering a finite set of irreducible logical forms whose truth-functional compositions constitute the "skeleton of the world." While certain connectives are mathematically redundant in propositional logic, they introduce unique inferential patterns and computational burdens for LLMs—such as branching case-splits for  $\oplus$  or variable binding for  $\forall$ —which are not captured by coarse complexity proxies. Consequently, our basis is *expressively sufficient* to cover the first-order semantics in our benchmarks while remaining *structurally minimal*. To account for global inference costs, we augment LoCM with a weighted hop term  $h(\phi)$  (weighted by  $\gamma = 2$ ). This formulation allows LoCM to distinguish between *deep-but-narrow* and *shallow-but-dense* reasoning chains, capturing diverse empirical failure modes that surface despite similar local structures.

**The final weights reflect both intuition and validation-set calibration.** We assign operator weights through a controlled grid search on a held-out validation set, using Pearson correlation between LoCM and model accuracy as a selection criterion. The resulting weights are summarized in Table 8. Importantly, the selected configuration aligns with intuitive differences in reasoning burden across operators. Exclusive OR ( $\oplus$ ) typically induces higher branching factors and explicit case splits, motivating its larger weight, while conditionals ( $\rightarrow, \leftrightarrow$ ) often require multi-branch implication handling. Quantifiers and negation introduce moderate structural bookkeeping, such as variable binding, scope management, or polarity switching, whereas basic connectives serve as the baseline. Beyond correlation-based selection, the necessity of the full weighted formulation is independently supported by the operator-weight ablations in Table 5, where removing or isolating any operator family consistently weakens the observed

correlation.

Table 8: Calibrated operator weights used in LoCM calculation.

Component	Operator(s)	Weight
Basic connectives	$\wedge, \vee$	1.0
Quantifiers	$\forall, \exists$	2.0
Negation	$\neg$	2.0
Hop term	$h(\phi)$	2.0
Conditionals	$\rightarrow, \leftrightarrow$	3.0
XOR	$\oplus$	3.5

**Multi-operator composition is necessary to avoid length-driven artifacts.** Single-operator heuristics tend to conflate *surface length* with *logical difficulty*. As shown in Figure 9, accuracy curves derived from individual operator counts (e.g., quantifier-only metrics) often exhibit **non-monotonic fluctuations** and spurious spikes in higher-complexity bins. Such artifacts arise because longer derivations mechanically accumulate specific operators without necessarily increasing the underlying reasoning challenge. By jointly modeling the full operator set in Eq. 1, the complete LoCM mitigates operator-specific proxy leakage and produces smoother, more monotonic degradation trends. This behavior more reliably exposes the underlying phase-transition structure, rather than artifacts induced by length or operator frequency alone.

**The square-root transform stabilizes binning by compressing heavy-tailed scores.** The raw score

$$S(\phi) = \sum_{o \in \mathcal{O}} \omega(o) \text{freq}(o, \phi) + \gamma h(\phi)$$

exhibits a heavy-tailed distribution, in which a small fraction of deeply nested instances attain disproportionately large values. Directly binning samples based on  $S(\phi)$  would therefore concentrate most instances into a narrow range, reducing resolution around the transition region. To address this issue, we apply a sub-linear monotone transformation, defining

$$\text{LoCM}(\phi) = \sqrt{S(\phi)}$$

, which preserves ordinal relationships while compressing extreme values. As shown in Table 3, this square-root transformation yields the strongest (most negative) correlation with accuracy among the Linear, Log, and Square variants considered, supporting its use as a stable operationalization of logical complexity.

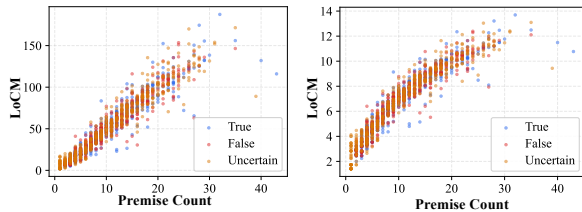


Figure 10: Premise count vs. complexity. **Left:** linear raw complexity score  $S(\phi)$ . **Right:** transformed metric  $\text{LoCM}(\phi) = \sqrt{S(\phi)}$ .

**LoCM is invariant to semantics-preserving serialization changes.** As a sanity check, we verify that LoCM reflects underlying *logical structure* rather than surface-level formatting. Specifically, the metric remains unchanged under semantics-preserving perturbations such as variable renaming, whitespace variation, and equivalent parenthesizations. This confirms that LoCM captures operator topology and hop structure, rather than artifacts of serialization.

### C.7 LoCM vs. Premise Count

**LoCM correlates with premise count but is not reducible to length.** A natural concern is that LoCM may act primarily as a proxy for input length, measured here by premise count. Figure 10 plots premise count against (Left) the linear raw complexity score  $S(\phi)$  and (Right) the transformed metric  $\text{LoCM}(\phi) = \sqrt{S(\phi)}$ . As expected, premise count and complexity exhibit a positive correlation, since longer derivations mechanically accumulate more logical operators. However, the relationship is clearly non-deterministic: instances with similar premise counts can span a wide range of complexity values. This dispersion indicates that LoCM captures substantial structural variation beyond length alone, reflecting differences in operator composition, nesting, and reasoning depth.

**The large vertical spread at fixed premise count indicates an independent structural signal.** As shown in Figure 10, for nearly any fixed premise-count value, instances span a wide range of raw complexity scores and corresponding LoCM values. This vertical dispersion reflects differences between *long-but-simple* structures (e.g., shallow chains dominated by  $\wedge/\vee$ ) and *structurally dense* reasoning patterns, such as deeper nesting, broader quantifier scopes, or higher-hop derivations, which premise count alone cannot distinguish. The square-root transformation compresses extreme

Table 9: Accuracy vs. LoCM within fixed premise-count intervals

Premise-count interval	LoCM bin	Accuracy
[1.0, 5.2]	1.4–2.4	96.77%
	2.4–3.4	95.30%
	3.4–4.4	93.08%
	4.4–5.4	85.39%
[13.6, 17.8]	5.1–6.3	77.78%
	6.3–7.4	66.67%
	7.4–8.5	61.67%
	8.5–9.7	49.57%

raw scores while preserving relative ordering, stabilizing the scale (Right) without collapsing this structural variation.

**Within narrow premise-count intervals, accuracy still decreases monotonically with LoCM.** To assess whether LoCM captures performance variation beyond length, we perform a stratified control-variate analysis. Specifically, we restrict instances to narrow premise-count intervals and examine accuracy trends across increasing LoCM bins. As shown in Table 9, accuracy consistently degrades as LoCM increases even when premise count is held approximately constant. For instance, within the range [1.0, 5.2], accuracy declines from 96.77% to 85.39% as LoCM increases, while in [13.6, 17.8], it drops from 77.78% to 49.57%. These results indicate that LoCM captures structural reasoning difficulty beyond what can be explained by premise count alone.

### C.8 Reasoning Effort Analysis

**Observable decoding statistics provide a practical proxy for reasoning effort.** We analyze reasoning effort using inference-time observable signals, primarily completion length (measured in pseudo-tokens), and relate them to answer accuracy and logical complexity (LoCM). While response length is an imperfect proxy for internal computation, it provides a practical indicator of model behavior on challenging instances.

**Longer generations correlate with lower accuracy, signaling difficulty rather than improved reasoning.** As shown in Table 10, accuracy decreases monotonically with increasing completion length, dropping from 94.66% in the shortest bin to below 40% in the tail. Rather than reflecting better reasoning, extended generations more often indicate that the model has entered a difficult regime, expanding intermediate steps without reli-

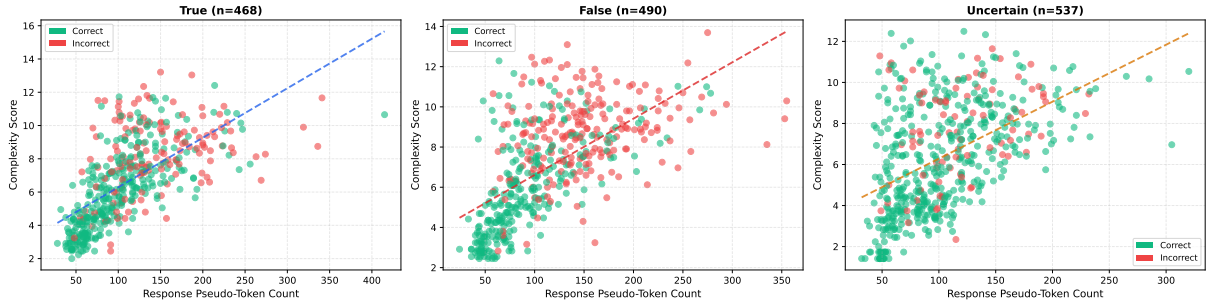


Figure 11: Completion length vs. LoCM, stratified by ground-truth label (Left: True; Middle: False; Right: Uncertain). Points are colored by correctness, and dashed lines denote linear fits.

Table 10: Accuracy as a function of completion length. Accuracy decreases sharply as generation length increases.

token range	Accuracy
24–63	94.66%
63–102	80.29%
102–141	63.51%
141–180	49.07%
180–220	39.77%
220–259	40.48%
259–415	42.86%

ably reaching a correct conclusion.

**Effort increases with LoCM for definite labels, indicating length tracks structured difficulty in those regimes.** We examine how completion length scales with LoCM, stratified by the ground-truth label (True / False / Uncertain). As shown in Figure 11, the True and False subsets exhibit clear positive correlations between completion length and LoCM (Pearson  $r = 0.652$  and  $r = 0.618$ , respectively), indicating that higher structural complexity tends to elicit longer generations when a definitive verdict exists.

**The Uncertain subset exhibits a high-variance, two-regime effort pattern.** In contrast, the Uncertain subset shows substantially higher dispersion and a weaker correlation between completion length and LoCM (Pearson  $r = 0.485$ ). The scatter reveals two regimes: (i) a short-response regime, where the model outputs Uncertain with minimal generation even at moderate-to-high LoCM, and (ii) a long-response regime, where extended derivations are produced with mixed correctness. This bimodal pattern is consistent with heterogeneous decision behaviors, which we report descriptively without attributing it to a specific internal mechanism.

## D Formal Setup and Notation

### D.1 Notation and Basic Syntax

**Core symbols.** We adopt standard first-order logic (FOL) notation. Variables ( $x, y, z$ ) range over a domain of discourse, constants ( $a, b, c$ ) denote specific objects, predicates ( $P(\cdot), R(\cdot, \cdot)$ ) denote properties or relations, and functions ( $f(\cdot), g(\cdot, \cdot)$ ) produce terms. Logical connectives are drawn from the operator set  $OP = \{\oplus, \vee, \wedge, \rightarrow, \leftrightarrow\}$ , and quantifiers are  $\forall$  and  $\exists$ . Table 11 summarizes the key syntax elements used throughout the paper.

**Terms, formulas, and well-formedness.** A *term* is either a variable, a constant, or a function applied to terms. An *atomic formula* is a predicate applied to a sequence of terms. A *formula* is built from atomic formulas using connectives and quantifiers. A formula is *well-formed* (WFF) if it is syntactically valid under our grammar and can be evaluated as true or false under an interpretation.

### D.2 Well-Formedness and CFG Validation

**CFG grammar.** To ensure the well-formedness of all FOL expressions, we implement a symbolic parser using `nltk` (Bird, 2006) and validate formulas with a context-free grammar (CFG):

S	$\rightarrow F \mid Q \mid F$
Q	$\rightarrow \text{QUANT VAR} \mid \text{QUANT VAR Q}$
F	$\rightarrow \text{'\(\rightarrow\'} \text{'( F '\')} \mid \text{'( F '\')} \mid \text{F OP F} \mid \text{L}$
OP	$\rightarrow \text{'\(\oplus\'} \mid \text{'\(\vee\'} \mid \text{'\(\wedge\'} \mid \text{'\(\rightarrow\'} \mid \text{'\(\leftrightarrow\}'$
L	$\rightarrow \text{'\(\rightarrow\'} \text{ PRED (' TERMS '\')} \mid \text{PRED (' TERMS '\')}$
TERMS	$\rightarrow \text{TERM} \mid \text{TERM '\('} \text{TERMS '\)'$
TERM	$\rightarrow \text{CONST} \mid \text{VAR}$
QUANT	$\rightarrow \text{'\(\forall\'} \mid \text{'\(\exists\}'$

**Dynamic instantiation.** Non-terminals PRED, CONST, and VAR are instantiated per example from the symbols appearing in the corresponding FOL context. This allows the same CFG to validate diverse instances while enforcing a consistent syntactic structure.

**Syntactic validation.** We apply CFG-based validation as a strict quality-control step before any

downstream symbolic reasoning. **Quantifier scope.** We require that each quantified variable is properly bound within its scope and appears as a valid VAR in the subsequent subformula. **Predicate–argument form.** We enforce that every predicate application matches the pattern  $\text{PRED}(\text{TERMS})$  with comma-separated terms, optionally preceded by  $\neg$ . **Operator placement.** We ensure that each binary connective in OP combines two valid subformulas and that parentheses yield an unambiguous parse under the CFG.

### D.3 NL-to-FOL Translation Rules

**Rule-based mapping.** We translate NL premises into FOL using deterministic rules that map entities to constants, properties to unary predicates, and relations to  $n$ -ary predicates. Logical operators and quantifiers are induced from explicit linguistic cues (e.g., “every”, “some”, “if–then”, “and/or”, “not”). Table 12 lists the translation rules used in our pipeline.

**Coreference and grounding.** When NL contains pronouns or referential expressions (e.g., “he”, “this person”), we first resolve them to an existing entity mention in the local context and then reuse the corresponding constant in the translated formula to maintain symbol consistency.

**Negation patterns.** We handle both local predicate negation (e.g., “not honest”  $\mapsto \neg \text{Honest}(x)$ ) and universal negation patterns (e.g., “no citizen is armed”  $\mapsto \forall x (\text{Citizen}(x) \rightarrow \neg \text{Armed}(x))$ ), ensuring that the resulting formulas remain WFF under the CFG in Appendix D.2.

**End-to-end check.** Each translated formula is passed through CFG validation. Only WFF expressions are retained for subsequent reasoning; malformed outputs are rejected and regenerated by the pipeline.

## E Future Directions

**Complexity-aware curriculum learning for reinforcement learning.** LoCM provides a structural foundation for curriculum design in logic-based reinforcement learning. By organizing instances according to LoCM, future training schedules can foster stable reasoning in low-complexity regimes before approaching the phase-transition boundary. Such a curriculum mitigates model collapse in difficult tasks by systematically bridging the gap between stable reasoning and the model’s intrinsic complexity threshold

**Process-level supervision via verifiable symbolic reasoning chains.** Standard RL for logic often struggles with sparse rewards and difficult credit assignment. Our Neuro-Symbolic alignment dataset addresses this by providing fully verifiable symbolic trajectories, enabling the development of Process Reward Models (PRMs). These models provide dense, step-wise feedback, replacing binary outcome-based rewards with structured verification signals that guide the model through complex reasoning paths.

**LoCM as a diagnostic tool for reasoning robustness.** Building on the view of Russell (Russell, 2020) and Wittgenstein (Wittgenstein, 2023) that logic underlies the structure of the world, we treat LoCM not merely as a metric for symbolic tasks but as a general framework for measuring reasoning complexity across domains. Any domain with an explicit logical structure can be formalized into symbolic reasoning chains. By encoding domain-specific premises and inference rules in an LoCM-compatible form, researchers can identify model-specific complexity thresholds, positioning LoCM as a diagnostic tool for reasoning robustness and for analyzing how LLMs engage with the logical structure of knowledge.

**Phase transitions as operational capacity boundaries.** Our results suggest that logical phase transitions reflect intrinsic reasoning limits rather than artifacts of specific optimization. Future research should treat phase-transition behavior as an operational notion of reasoning capacity. A key direction is to investigate how different paradigms—such as reinforcement learning or architectural scaling—alter the sharpness or onset of this boundary without assuming it can be arbitrarily shifted.

## F Full Prompts and Examples

Below are detailed prompts used in our evaluation. The prompts for constructing our datasets can be found in ProverGen (Qi et al., 2025).

Table 11: Key Syntax Elements in First-Order Logic

Name	FOL Notation	Explanation
<b>Variable</b>	$x, y, z$	Placeholder symbols representing arbitrary elements in the domain of discourse.
<b>Constant</b>	$a, b, c$	Refer to specific, fixed objects in the domain.
<b>Operators (OP)</b>	$\{\oplus, \vee, \wedge, \rightarrow, \leftrightarrow\}$	Defines the set of logical connectives used to combine or relate propositions, including exclusive or, or, and, implication, and biconditional. Used in building compound formulas.
<b>Function</b>	$f(x), g(x, y)$	Maps input objects to an output object; returns a term.
<b>Predicate</b>	$P(x), R(x, y)$	Express properties or relations; returns true or false.
<b>Negation</b>	$\neg P(x)$	Logical NOT: $P(x)$ is not true.
<b>Conjunction</b>	$P(x) \wedge Q(x)$	Logical AND: both $P(x)$ and $Q(x)$ must be true.
<b>Disjunction</b>	$P(x) \vee Q(x)$	Logical OR: at least one of $P(x)$ or $Q(x)$ must be true.
<b>Implication</b>	$P(x) \rightarrow Q(x)$	Logical implication: if $P(x)$ is true, then $Q(x)$ must be true.
<b>Biconditional</b>	$P(x) \leftrightarrow Q(x)$	Logical equivalence: $P(x)$ and $Q(x)$ are true or false together.
<b>Universal Quantifier</b>	$\forall x P(x)$	“For all $x$ , $P(x)$ is true” — generalization.
<b>Existential Quantifier</b>	$\exists x P(x)$	“There exists $x$ such that $P(x)$ is true” — existential claim.
<b>Term</b>	$x, a, f(a, x)$	The basic expressions referring to objects (variables, constants, or functions).
<b>Atomic Formula</b>	$P(a, x)$	A predicate applied to terms — indivisible logical unit.
<b>Complex Formula</b>	$\forall x(P(x) \rightarrow Q(f(x)))$	A formula built from atoms using connectives and quantifiers.
<b>WFF (Well-formed)</b>	—	A syntactically valid FOL formula interpretable as true or false.

Table 12: Natural Language to First-Order Logic (FOL) Translation Rules

Linguistic Element	NL Example	FOL Translation Rule and Output
<b>Entities / Objects</b>	“Moriarty is a cat.”	Nouns map to constants; adjectival properties map to unary predicates: $Cat(Moriarty)$
<b>Attributes / Properties</b>	“Moriarty is fluffy.”	Adjectives become unary predicates: $Fluffy(Moriarty)$
<b>Binary Relations</b>	“Moriarty comforts a customer.”	Verb relations become binary predicates: $Comforts(Moriarty, cust)$
<b><math>n</math>-ary Relations</b>	“Colt contributes to conservation.”	Multi-argument verbs become $n$ -ary predicates: $Contributes(Colt, cons)$
<b>Universal Quantification</b>	“All playful cats are fluffy.”	Quantified NP induces $\forall$ : $\forall x (Playful(x) \wedge Cat(x) \rightarrow Fluffy(x))$
<b>Existential Quantification</b>	“Some cats are playful.”	Existential statements induce $\exists$ : $\exists x (Cat(x) \wedge Playful(x))$
<b>Implication</b>	“If a cat is playful, then it loves attention.”	Conditionals map to implication: $Playful(x) \rightarrow LovesAttn(x)$
<b>Conjunction</b>	“Colt is dedicated and makes discoveries.”	Conjunction markers map to $\wedge$ : $Dedicated(Colt) \wedge Discovers(Colt)$
<b>Disjunction</b>	“A cat is either warm or calming.”	Inclusive disjunction: $Warm(x) \vee Calming(x)$
<b>Exclusive-Or</b>	“A cat is either warm or calming, but not both.”	Exclusive contrast: $Warm(x) \oplus Calming(x)$
<b>Negation</b>	“Moriarty is not aggressive.”	Negated predicates: $\neg Aggressive(Moriarty)$
<b>Universal Negation</b>	“No cat is aggressive.”	“No X is Y” $\rightarrow$ universal implication with negation: $\forall x (Cat(x) \rightarrow \neg Aggressive(x))$
<b>Atomic Reasoning Step</b>	“If a cat is fluffy and warm, then it comforts people.”	Multi-premise reasoning: $(Fluffy(x) \wedge Warm(x)) \rightarrow Comforts(x, ppl)$
<b>Multi-step Chain</b>	“If a cat is fluffy, then it is warm; if it is warm, then it comforts people.”	Sequential implications: $Fluffy(x) \rightarrow Warm(x), Warm(x) \rightarrow Comforts(x, ppl)$
<b>Coreference Resolution</b>	“He is fluffy.” (referring to Moriarty)	Resolve entity first: $Fluffy(Moriarty)$
<b>Role Assignment</b>	“Colt studies plants.”	Action roles become predicates: $Studies(Colt, plants)$

## Naive Prompting

System:

Given a problem statement as contexts, the task is to answer a logical reasoning question. Your answer should be in JSON format with key: answer.

Context:

has\_dense\_fur(Wynter)::Wynter has dense fur.

is\_playful(Wynter)::Wynter is playful.

has\_dense\_fur(Wynter)  $\rightarrow$  (is\_playful(Wynter)  $\oplus$  soft\_fur(Wynter))::If Wynter has dense fur, then she is either playful or has soft fur, but not both.

Question: Based on the above information, is the following statement true, false, or uncertain?

Wynter is not beloved by others:: $\neg$ beloved\_by\_others(Wynter)

Options:

A) True

B) False

C) Uncertain

The correct option is: {

"answer": "C"

}

Other Examples

Context:

[[CONTEXT]]

Question:

[[QUESTION]]

Options:

[[OPTIONS]]

The correct option is:

## CoT Prompting

System:

Given a problem statement as contexts, the task is to answer a logical reasoning question. Your answer should be in JSON format with keys: reasoning, answer.

Context:

chooses\_words\_carefully(Peyton)::Peyton chooses words carefully.

listens\_attentively(Peyton)::Peyton listens attentively.

$\forall x ((chooses\_words\_carefully(x) \vee listens\_attentively(x)) \rightarrow reserved\_speaker(x))$ :::If a person chooses words carefully or listens attentively, then they are a reserved speaker.

$\forall x ((reserved\_speaker(x) \vee thoughtful\_listener(x)) \rightarrow builds\_meaningful\_connections(x))$

:::For all humans, if they are reserved speakers or thoughtful listeners, then they build meaningful connections.

Question: Based on the above information, is the following statement true, false, or uncertain?

Peyton does not build meaningful connections:: $\neg$ builds\_meaningful\_connections(Peyton)

Options:

A) True

B) False

C) Uncertain

The correct option is:

```
{
"reasoning":
"fact1: chooses_words_carefully(Peyton):::Peyton chooses words carefully.
fact2: listens_attentively(Peyton):::Peyton listens attentively.
rule:
 $\forall x ((chooses\_words\_carefully(x) \vee listens\_attentively(x)) \rightarrow reserved\_speaker(x))$ 
:::If a person chooses words carefully or listens attentively, then they are a reserved speaker.
conclusion: reserved_speaker(Peyton):::Peyton is a reserved speaker.
fact1: reserved_speaker(Peyton):::Peyton is a reserved speaker.
rule:
 $\forall x ((reserved\_speaker(x) \vee thoughtful\_listener(x)) \rightarrow builds\_meaningful\_connections(x))$ 
:::For all humans, if they are reserved speakers or thoughtful listeners, then they build meaningful
connections.
conclusion: builds_meaningful_connections(Peyton):::Peyton builds meaningful connections.
Therefore, it is false that Peyton does not build meaningful connections. The correct option is: B.",
"answer": "B"
}
```

---

#### Other Examples

---

Context:

[[CONTEXT]]

Question:

[[QUESTION]]

Options:

[[OPTIONS]]

The correct option is:

### Context & Symbolic Mapping:

Let  $K := \text{Kaizen}$ .

NL: Kaizen is a philosophy.

FOL:  $\text{Philosophy}(K)$

NL: Kaizen does not inspire cooperation.

FOL:  $\neg \text{inspires\_cooperation}(K)$

NL: Kaizen does not encourage learning.

FOL:  $\neg \text{encourages\_learning}(K)$

NL: If a philosophy leads to progress, then it inspires cooperation and encourages learning.

FOL:  $\forall x ((\text{Philosophy}(x) \wedge \text{leads\_to\_progress}(x)) \rightarrow (\text{inspires\_cooperation}(x) \wedge \text{encourages\_learning}(x)))$

NL: Kaizen's philosophy is transformative.

FOL:  $\text{is\_transformative}(K)$

NL: Kaizen challenges the status quo.

FOL:  $\text{challenges\_status\_quo}(K)$

NL: If Kaizen's philosophy is transformative, then it either challenges the status quo or embraces small changes, but not both.

FOL:  $\text{is\_transformative}(K) \rightarrow (\text{challenges\_status\_quo}(K) \oplus \text{embraces\_small\_changes}(K))$

NL: Any philosophy that embraces small changes or inspires continuous change leads to progress.

FOL:  $\forall x ((\text{Philosophy}(x) \wedge (\text{embraces\_small\_changes}(x) \vee \text{inspires\_continuous\_change}(x))) \rightarrow \text{leads\_to\_progress}(x))$

NL: Kaizen's philosophy either inspires continuous change or supports growth, but not both.

FOL:  $\text{inspires\_continuous\_change}(K) \oplus \text{supports\_growth}(K)$

NL: Kaizen advocates equality.

FOL:  $\text{advocates\_equality}(K)$

NL: If Kaizen's philosophy supports growth, then it advocates equality and encourages fair sharing.

FOL:  $\text{supports\_growth}(K) \rightarrow (\text{advocates\_equality}(K) \wedge \text{encourages\_fair\_sharing}(K))$

NL: For any philosophy, it either encourages fair sharing or fosters social balance, but not both.

FOL:  $\forall x (\text{Philosophy}(x) \rightarrow (\text{encourages\_fair\_sharing}(x) \oplus \text{fosters\_social\_balance}(x)))$

NL: If Kaizen's philosophy promotes redistribution, then it fosters social balance.

FOL:  $\text{promotes\_redistribution}(K) \rightarrow \text{fosters\_social\_balance}(K)$

### Question:

NL: Based on the above information, is the following statement true, false, or uncertain?

Kaizen's philosophy promotes redistribution.

Target FOL:  $\text{promotes\_redistribution}(K)$

### Reasoning Trace (Dual-Track)

#### -Step 1-

NL: If leading to progress entails inspiring cooperation and encouraging learning, then the failure of either consequence rules out leading to progress. Since Kaizen does not inspire cooperation (and also does not encourage learning), Kaizen does not lead to progress.

FOL:  $\text{Philosophy}(K), \neg \text{inspires\_cooperation}(K), \forall x ((\text{Philosophy}(x) \wedge \text{leads\_to\_progress}(x)) \rightarrow (\text{inspires\_cooperation}(x) \wedge \text{encourages\_learning}(x))) \vdash \neg \text{leads\_to\_progress}(K)$

#### -Step 2-

NL: Being transformative enforces an exclusive choice: challenge the status quo or embrace small changes. Since Kaizen challenges the status quo, it cannot embrace small changes.

FOL:  $\text{is\_transformative}(K), \text{challenges\_status\_quo}(K), \text{is\_transformative}(K) \rightarrow (\text{challenges\_status\_quo}(K) \oplus \text{embraces\_small\_changes}(K)) \vdash \neg \text{embraces\_small\_changes}(K)$

#### -Step 3-

NL: Any philosophy that embraces small changes or inspires continuous change would lead to progress. Since Kaizen does not lead to progress, and it does not embrace small changes, it also cannot inspire continuous change.

FOL:  $\text{Philosophy}(K), \neg \text{embraces\_small\_changes}(K), \neg \text{leads\_to\_progress}(K), \forall x ((\text{Philosophy}(x) \wedge (\text{embraces\_small\_changes}(x) \vee \text{inspires\_continuous\_change}(x))) \rightarrow \text{leads\_to\_progress}(x)) \vdash \neg \text{inspires\_continuous\_change}(K)$

#### -Step 4-

NL: Kaizen must be exactly one of: inspires continuous change or supports growth. Since it does not inspire continuous change, it supports growth.

FOL:  $\neg \text{inspires\_continuous\_change}(K), (\text{inspires\_continuous\_change}(K) \oplus \text{supports\_growth}(K)) \vdash \text{supports\_growth}(K)$

#### -Step 5-

NL: Supporting growth implies advocating equality and encouraging fair sharing; therefore Kaizen encourages fair sharing. The standalone fact that Kaizen advocates equality is consistent but not required for this derivation.

FOL:  $\text{supports\_growth}(K), \text{supports\_growth}(K) \rightarrow (\text{advocates\_equality}(K) \wedge \text{encourages\_fair\_sharing}(K)) \vdash \text{encourages\_fair\_sharing}(K)$

#### -Step 6-

NL: For any philosophy, fair sharing and social balance are mutually exclusive. Since Kaizen encourages fair sharing, it does not foster social balance.

FOL:  $\text{Philosophy}(K), \text{encourages\_fair\_sharing}(K), \forall x (\text{Philosophy}(x) \rightarrow (\text{encourages\_fair\_sharing}(x) \oplus \text{fosters\_social\_balance}(x))) \vdash \neg \text{fosters\_social\_balance}(K)$

#### -Step 7-

NL: Promoting redistribution would imply fostering social balance. Since Kaizen does not foster social balance, it does not promote redistribution.

FOL:  $\neg \text{fosters\_social\_balance}(K), (\text{promotes\_redistribution}(K) \rightarrow \text{fosters\_social\_balance}(K)) \vdash \neg \text{promotes\_redistribution}(K)$

### Answer:

Therefore, the statement is **False**. The correct option is B.

Figure 12: Detailed case study of a high-complexity instance ( $LoCM \approx 7.25$ ) from the NSA-LR dataset. The example demonstrates strict NL-FOL alignment across a 7-step reasoning chain involving diverse logical operators ( $\neg, \rightarrow, \oplus, \forall$ ). The derivation is carried out under classical first-order logic, including contraposition-based inferences where needed.