

Adaptive LLM-Symbolic Reasoning via Dynamic Logical Solver Composition

Lei Xu^{1,2*}, Pierre Beckmann^{1,2*}, Marco Valentino³, André Freitas^{1,4,5}

¹Idiap Research Institute, Switzerland

²École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

³School of Computer Science, University of Sheffield, United Kingdom

⁴Department of Computer Science, University of Manchester, United Kingdom

⁵Cancer Biomarker Centre, CRUK Manchester Institute, United Kingdom

{firstname.lastname}@idiap.ch

Abstract

Neuro-symbolic NLP methods aim to leverage the complementary strengths of large language models and formal logical solvers. However, current approaches are mostly static in nature, i.e., the integration of a target solver is pre-determined at design time, hindering the ability to employ diverse formal inference strategies. To address this, we introduce an adaptive, multi-paradigm, neuro-symbolic inference framework that: (1) automatically identifies formal reasoning strategies from problems expressed in natural language; and (2) dynamically selects and applies specialized formal logical solvers via autoformalization interfaces. Extensive experiments on individual and multi-paradigm reasoning tasks support the following conclusions: LLMs are effective at predicting the necessary formal reasoning strategies with an accuracy above 90%. This enables flexible integration with formal logical solvers, resulting in our framework outperforming competing baselines – by 17% and 6% compared to GPT-4o and DeepSeek-V3.1, respectively. Moreover, adaptive reasoning can even positively impact pure LLM methods, yielding gains of 10%, 5%, and 6% on zero-shot, CoT, and CoT_{sym} settings with GPT-4o. Finally, although smaller models struggle with adaptive neuro-symbolic reasoning, post-training offers a viable path to improvement. Overall, this work establishes the foundations for adaptive LLM-symbolic reasoning, offering a path forward for unifying material and formal inferences on heterogeneous reasoning challenges. The code and data are available online¹.

1 Introduction

Large language models (LLMs) have demonstrated sophisticated material inference properties but remain limited by hallucinations and logical inconsis-

tencies (Zhang et al., 2023; Guo et al., 2025; Madusanka et al., 2025). Formal reasoning systems, conversely, provide systematic guarantees with transparent and deterministic inference processes but require significant manual formalization (Nethercote et al., 2007; Frederiksen, 2008). Recent advances in LLM-driven auto-formalization (Wu et al., 2022; Zhang et al., 2025) have enabled the development of reasoning paradigms that bridge the material-formal gap (Pan et al., 2023; Quan et al., 2025a; Xu et al., 2024; Quan et al., 2024b,a; Górski et al., 2025; Quan et al., 2025b). However, existing approaches share a fundamental architectural constraint: they are designed for specific formal reasoning paradigms. For instance, LogicLLaMA (Yang et al., 2023a) specializes in first-order logic translation, Sat-LM (Ye et al., 2023) targets Boolean satisfiability problems, LINC (Olausson et al., 2023) handles first-order-logic problems, and Logic-LM (Pan et al., 2023), while supporting multiple logical solvers, requires specifying the reasoning strategy at design time. This static paradigm reflects an implicit assumption that the necessary strategies are known in advance, limiting the ability of neuro-symbolic/material-formal models to handle heterogeneous problems.

To address these challenges, we introduce a novel framework that enables a multi-paradigm adaptive reasoning. Our framework overcomes static solver integration through a context-aware reasoning strategy classification that automatically analyzes natural language inputs to identify required formal reasoning types in an unsupervised manner. Subsequently, the framework performs dynamic solver selection and composition, flexibly identifying appropriate formal logical solvers from an extensible library. This neuro-symbolic integration incorporates mechanisms to evaluate problem structure and adaptively select efficient formalization paradigms while maintaining expressiveness.

We conduct an extensive empirical evaluation

*Equal contribution.

¹https://github.com/idiap/adaptive_symbolic_reasoning

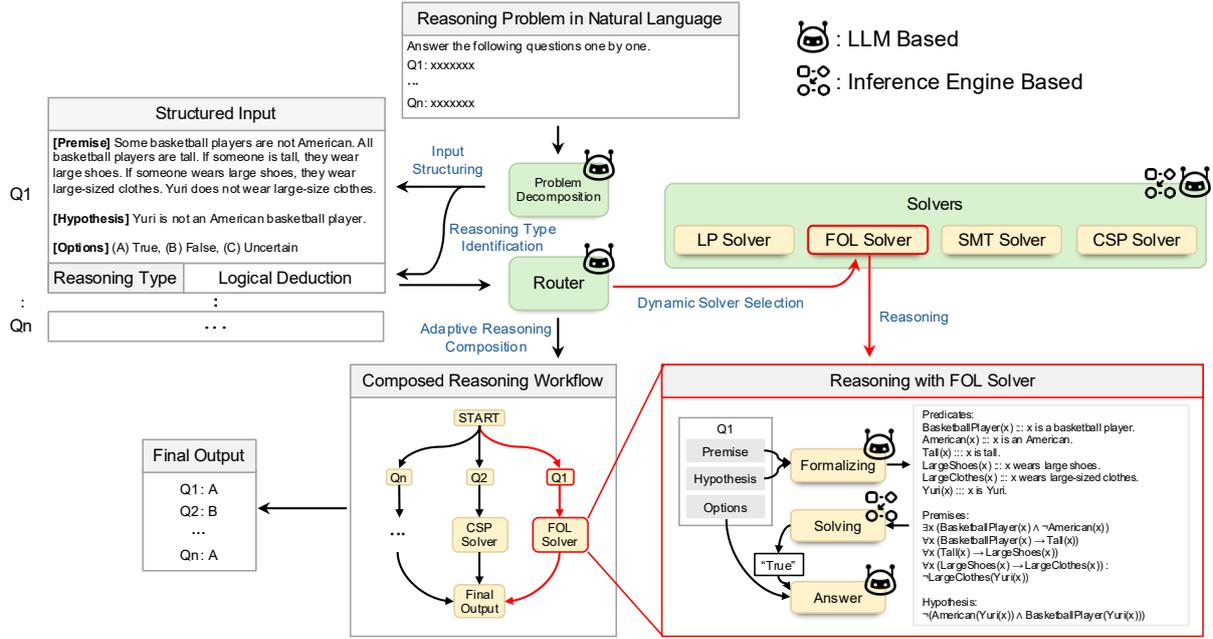


Figure 1: Overview of the proposed adaptive symbolic reasoning framework. Given a natural language reasoning problem, the system first performs problem decomposition to extract structured components and identify the corresponding reasoning type (LP, FOL, CSP, SMT). Based on this analysis, a Router dynamically selects the appropriate solver for each problem and orchestrates the reasoning process. Each solver performs autoformalization on the structured input and conducts formal reasoning to produce verified answers. Unlike static approaches that predetermine solver integration, our framework adaptively composes specialized solvers through problem-aware classification and dynamic orchestration, enabling robust handling of heterogeneous reasoning tasks.

on five diverse benchmarks spanning logical reasoning, constraint satisfaction, and clinical tria-patient matching (ProntoQA, ProofWriter, FOLIO, LogDed, $TREC_{trials}$), alongside newly designed multi-question stress tests to evaluate solver selection and composition. Our findings support the following conclusions:

1. Dynamic reasoning strategy selection using LLMs can effectively enable adaptive behavior, with large frontier models achieving over 98% accuracy.
2. Adaptive solver composition supports effective neuro-symbolic inferences. On a mixed dataset requiring diverse strategies, our framework achieves 92.1% accuracy, substantially outperforming zero-shot by 17.0%, Chain-of-Thought by 21.4%, and Symbolic CoT by 23.7%.
3. Dynamic solver composition is critical for robustness on sequential reasoning tasks. When solving tasks containing sequences of heterogeneous problems, pure LLM methods achieve only 27.3% overall accuracy while the

proposed multi-paradigm framework achieves up to 54.4%. This demonstrates that adaptive neuro-symbolic methods can offer a solution to substantially improve robustness on challenging reasoning problems which requires multi-paradigm formal inferences.

4. Adaptive reasoning still poses challenges to smaller models, yet post-training enables significant improvements. In particular, small models show significant performance drops (8% on Llama-3.1-8B, 12.6% on Qwen-2.5-7B, 26.2% on Qwen-2.5-Coder-7B), with 60-80% of errors stemming from invalid formalization. However, fine-tuning using demonstrations from frontier models yields substantial improvements, reaching 41.0% on Llama-3.1-8B, 50.1% on Qwen-2.5-7B, and 59.6% on Qwen-2.5-Coder-7B.

2 Adaptive Multi-Paradigm Reasoning Framework

Given a natural language reasoning problem $x \in \mathcal{X}$ with ground truth answer $a \in \mathcal{A}$, our goal is to construct an end-to-end framework $\mathcal{F} : \mathcal{X} \rightarrow \mathcal{A}$ that

produces $\hat{\mathbf{a}} = \mathcal{F}(\mathbf{x})$ such that $\hat{\mathbf{a}} = \mathbf{a}$. Unlike traditional neuro-symbolic approaches that assume predetermined reasoning types, our framework handles raw natural language without prior knowledge of reasoning characteristics. As shown in Figure 1, the framework consists of three main stages: (i) problem decomposition, (ii) routing and (iii) reasoning, which formalizes the problem and calls the selected solver to produce predicted answers.

We thus formalize the end-to-end system $\mathcal{F} : \mathcal{X} \rightarrow \mathcal{A}$ as a compositional map:

$$\mathcal{F} = \text{Reason}(\text{Route}(\text{Decompose}(\mathbf{x}))).$$

For all LLM-based modules, we utilize frontier models (GPT-4o, DeepSeek-V3) and three small-scale open-source models (Llama-3.1-8B, Qwen-2.5-7B, Qwen-2.5-Coder-7B), to assess the scalability of our approach, see Section 3.1 for details.

Problem Decomposition. Existing neuro-symbolic solvers require rigid pre-specified problem types and structured inputs, which limits their practical applicability in scenarios where the reasoning strategy is not specified in advance. Moreover, real-world applications may involve multi-reasoning scenarios in which problems need to be addressed by solving multiple sequential questions. To address these challenges, we leverage a LLM-based parser component responsible for structured problem decomposition.

The framework first converts \mathbf{x} into a semi-structured representation capturing both decomposition and reasoning-type information. This step allows the system to handle inputs that may encompass multiple reasoning sub-tasks, each requiring a different formal paradigm:

$$\begin{aligned} (\mathcal{Q}, \mathcal{T}) &= \text{Decompose}(\mathbf{x}), \\ \mathcal{Q} &= \{Q_i\}_{i=1}^n, \quad \mathcal{T} = \{T_i\}_{i=1}^n, \\ T_i &\in \{\text{LP, FOL, CSP, SMT, } \dots\}. \end{aligned}$$

where each $Q_i = \{C_{i,j}\}_{j=1}^{m_i}$ represents the input \mathbf{x} structured into m_i key components $C_{i,j}$ required for reasoning under the corresponding paradigm $T_i \in \{\text{LP, FOL, CSP, SMT, } \dots\}$, where LP stands for Logic Programming, FOL stands for First-Order Logic, CSP stands for Constraint Satisfaction Problems, and SMT stands for Satisfiability Modulo Theories.

Inference Routing. Given the decomposed representation $\mathcal{Q}, \mathcal{T} = \text{Decompose}(\mathbf{x})$, the Router

component determines which solver should handle each sub-problem based on its reasoning type. Formally, let

$$\mathcal{S} = \{S_T \mid T \in \{\text{LP, FOL, CSP, SMT, } \dots\}\}$$

denote the portfolio of available solvers S_T . To leverage the reliability and interpretability of symbolic reasoning engines, we build each solver S_T upon its corresponding symbolic reasoning engine (e.g., Pyke for LP and Prover9 for FOL, see Section 3.1 for details).

For each sub-question $Q_i \in \mathcal{Q}$ with associated reasoning type $T_i \in \mathcal{T}$, the router selects the corresponding solver S_{T_i} and invokes it to obtain a predicted answer:

$$\hat{a}_i = S_{T_i}(Q_i).$$

Reasoning: Autoformalization & Solving.

Each solver S_T in the portfolio operates on inputs that strictly conform to its formal syntax, which is not compatible with natural language components \mathcal{Q} . To bridge this gap, we equip each solver with an LLM-based autoformalization preprocessing step. Specifically, given $Q_i = \{C_{i,j}\}_{j=1}^{m_i}$ with reasoning type T_i , the solver S_{T_i} first employs an LLM to formalize each component $C_{i,j}$ into the syntax required by the reasoning engine for paradigm T_i . This representation is then fed into the symbolic reasoning engine to obtain inference results. Formally, the reasoning stage applies typed solvers via a formalize – solve pipeline:

$$\begin{aligned} \hat{a}_i &= S_{T_i}(Q_i) \\ &= \text{Solve}_{T_i}(\text{Formalize}_{T_i}(Q_i, T_i)) \end{aligned}$$

where $\text{Formalize}(Q_i, T_i)$ translates each component $C_{i,j} \in Q_i$ into the formal language of solver S_{T_i} . The formalized problem is then solved with the symbolic engine. Next, the output is transformed into a final answer via an LLM. Finally, the outputs of the different sub-tasks are aggregated $\hat{\mathbf{a}} = \{\hat{a}_i\}_{i=1}^n$. The prompts for autoformalization are presented in Appendix I. Algorithm 1 outlines the underlying end-to-end process, and Appendix B the implementation details for the Agentic workflow execution.

3 Empirical Analysis

3.1 Experimental Settings

Datasets. In order to evaluate the proposed dynamic reasoning framework, we select a heterogeneous set of datasets spanning logical reasoning, constraint satisfaction and clinical trial-patient

Methods	PrOntoQA		ProofWriter		FOLIO		LogDed7		TREC _{trials}		Mixed		Routing
	w/o routing	w/ routing	w/o routing	w/ routing	w/o routing	w/ routing							
GPT-4o													
Zero-shot	82.2	<u>91.6</u>	45.5	<u>60.3</u>	63.7	76.0	67.3	<u>74.4</u>	71.3	<u>78.0</u>	65.1	<u>75.1</u>	98.0
CoT	84.4	<u>90.2</u>	43.5	<u>52.5</u>	64.7	<u>71.1</u>	67.9	<u>71.0</u>	70.0	<u>73.7</u>	65.1	<u>70.7</u>	
CoT _{sym}	75.4	<u>97.8</u>	44.0	<u>53.7</u>	66.7	53.4	63.6	61.4	67.7	<u>75.7</u>	61.8	<u>68.4</u>	
Ours	-	99.4	-	94.7	-	69.1	-	95.6	-	82.3	-	92.1	
DeepSeek-V3.1													
Zero-shot	7.8	2.2	7.0	3.8	17.2	<u>18.1</u>	36.0	14.7	46.7	42.3	22.0	13.1	99.3
CoT	24.8	2.0	18.3	2.2	44.6	28.4	57.0	37.7	27.0	17.7	34.9	17.3	
CoT _{sym}	57.8	<u>91.2</u>	37.8	<u>57.0</u>	66.2	73.5	64.4	51.7	66.7	79.7	56.5	<u>67.2</u>	
Ours	-	97.6	-	81.3	-	52.0	-	57.7	-	68.0	-	73.4	

Table 1: Overall pass@1 performance on representative proprietary and open-source models. “w/o routing” and “w/ routing” denote the performance without and with the auxiliary reasoning selection strategy. Bold indicates the best setting for each dataset. Underline indicates improvement after adding routing information. The last column reports the Router’s prediction accuracy on the Mixed dataset. Since all settings use the same prompt and planning accuracy shows minimal variation across different settings, we report the mean accuracy for each backbone model.

Algorithm 1 Overall Workflow

Input: Natural Language Input x

Output: Final Answer \hat{a}

```

/Stage 1|Problem Decomposition/
 $Q, T = \text{Parser}(x)$ 
/Stage 2|Solver Registration/
 $S = \text{Register}()$ 
/Stage 3|Dynamic Reasoning Composition/
 $\hat{a} = [ ]$ 
for  $Q \in \mathcal{Q}, T \in \mathcal{T}$  do
   $S_T = \text{Router}(T)$  // Solver Selection
   $l = \text{Formalize}_T(Q)$  // Formalization
   $\hat{a} = \text{Solve}_T(l)$  // Calling Solver
   $\hat{a}.\text{append}(\hat{a})$ 
end for

```

matching including: PrOntoQA (Saparov and He, 2023), ProofWriter (Tafjord et al., 2021), FOLIO (Han et al., 2024), LogDed7 (bench authors, 2023), TREC_{trials} (Soboroff, 2021), where PrOntoQA and ProofWriter represent LP problems, FOLIO represents FOL problems, LogDed7 represents CSP problems and TREC_{trials} represents SMT problems. Additional details on the datasets and their formal selection criteria can be found in Appendix C. We convert all instances from the five datasets into a unified natural language format using task-specific templates to simulate the pure natural language input scenarios encountered in real-world applications (see Appendix H), then shuffle and merge them into a single **Mixed** dataset. Under this setup, the framework must first identify the reasoning strategy, then compose the relevant solvers and perform reasoning accordingly.

Baselines. We use LLMs under different prompting settings as baselines. Specifically, we evaluate three LLM settings, including: Zero-shot, Chain-of-Thought (CoT) (Kojima et al., 2022), Symbolic CoT (CoT_{sym}) (Xu et al., 2024). Additionally, to validate the effectiveness of the routing mechanism within our framework, we integrate the Router’s outputs as auxiliary prompts into the aforementioned Zero-shot, CoT, and CoT_{sym} settings. We also analyse the benefits of adaptive reasoning over static neuro-symbolic approaches such as LogicLM (Pan et al., 2023) through dedicated ablation studies in Section 3.4.

Foundation Models. We use four reference foundation models across different families to evaluate the proposed approach: one large proprietary model (GPT-4o (OpenAI, 2024)), one large open-source model (DeepSeek-V3.1 (DeepSeek-AI, 2024)) and three smaller models (Llama-3.1-8B-Instruct (Llama team, 2024), Qwen-2.5-7B-Instruct (Qwen Team, 2024), and Qwen-2.5-Coder-7B-Instruct (Qwen Team, 2024)), under Zero-shot, CoT, CoT_{sym} settings. For smaller models, we use the pretrained versions for problem decomposition and routing in our framework as they demonstrate stronger performance. For the Solver component, we employ pretrained versions for autoformalization. Additionally, we experiment with fine-tuning the small models on formalizations generated by GPT-4o to investigate the potential for performance improvement via distillation methods.

To manage inference costs, we evaluate GPT-4o and DeepSeek-V3.1 using pass@1 metric, measuring accuracy based on a single inference run. For small models, we conduct experiments three times with different seeds and report the average

Methods	PrOntoQA		ProofWriter		FOLIO		LogDed7		TREC _{trials}		Mixed		Routing
	w/o routing	w/ routing	w/o routing	w/ routing	w/o routing	w/ routing							
Llama-3.1-8b													
Zero-shot	65.2	58.6	41.8	36.3	46.7	42.6	37.0	25.3	63.2	39.7	48.7	38.8	76.8
CoT	57.9	54.1	45.3	35.4	38.2	<u>42.3</u>	19.5	<u>24.0</u>	63.6	48.3	42.0	38.3	
CoT _{sym}	10.1	0.0	10.8	0.2	12.1	0.2	3.9	0.6	2.6	<u>3.1</u>	7.6	0.7	
Ours	-	3.9	-	5.3	-	0.8	-	7.0	-	27.6	-	8.0	
Ours _{soft}	-	75.7	-	51.5	-	31.5	-	10.1	-	40.9	-	41.0	
Qwen-2.5-7b													
Zero-shot	57.2	40.9	37.8	29.8	32.4	<u>40.0</u>	42.9	38.5	59.3	53.2	45.9	38.8	97.7
CoT	95.3	72.8	59.8	59.1	48.0	47.7	46.8	20.6	47.0	25.3	60.8	45.0	
CoT _{sym}	55.1	46.1	41.7	29.7	59.2	52.3	8.0	49.5	60.2	61.9	38.3	<u>45.5</u>	
Ours	-	0.0	-	0.0	-	0.0	-	31.5	-	23.1	-	12.6	
Ours _{soft}	-	71.8	-	51.9	-	6.4	-	45.0	-	51.9	-	50.1	
Qwen-2.5-Coder-7b													
Zero-shot	71.5	67.8	53.6	56.5	45.3	25.7	45.8	44.1	66.2	66.7	56.0	53.8	92.9
CoT	70.5	64.8	51.3	<u>53.3</u>	45.4	24.8	40.5	40.4	63.0	<u>65.2</u>	53.2	50.9	
CoT _{sym}	78.2	29.8	46.8	20.6	52.1	13.2	14.0	11.3	38.0	29.3	43.0	20.3	
Ours	-	6.1	-	12.8	-	2.8	-	60.2	-	23.2	-	26.2	
Ours _{soft}	-	89.4	-	44.9	-	3.1	-	74.2	-	43.3	-	59.6	

Table 2: Average performance of small open-source models over three random seeds. Since all settings use the same routing result from the pretrained model, we only report the accuracy once under the same backbone model.

performance across three runs. Additional implementation details can be found in Appendix D.

Logical Solvers. We instantiate our framework with a set of symbolic solvers to support logical reasoning tasks. Four back-end engines are integrated: Pyke (Frederiksen, 2008) for LP problems, Prover9 (McCune, 2005–2010) for FOL problems, Z3 (de Moura and Bjørner, 2008) for SMT problems, and MiniZinc (Nethercote et al., 2007) for CSP problems, respectively. Pyke performs rule-based inference over formalized premises and hypotheses, while Prover9 conducts first-order deduction. Z3 checks satisfiability of symbolic constraints within logical theories, and MiniZinc solves combinatorial search problems over variable domains to produce concrete assignments. See Appendix A for further details.

3.2 Results on Frontier Models

Table 1 compares the performance before and after incorporating the multi-reasoning planning, under different foundation models. In addition, the last column reports the routing accuracy, i.e., the proportion of test cases where the Router correctly selects the expected solver. Based on the results in Table 1, we draw the following conclusions:

The proposed adaptive neuro-symbolic framework achieves the best overall performance compared to all baselines. Overall, our framework combined with larger models achieves the best performance, reaching 92.1% accuracy on the Mixed dataset. Our framework also yields significant improvements across different reasoning types, which

demonstrate the superior performance of combining LLM with symbolic solvers.

Dynamic routing strategy is effective using large frontier models. As shown in the last column, both GPT-4o and DeepSeek-V3.1 achieve routing accuracy exceeding 98%, with DeepSeek-V3.1 even slightly outperforming GPT-4o. This demonstrates that current frontier models are capable of dynamically parsing and identifying reasoning types, which validates the feasibility of the proposed dynamic routing strategy.

Dynamic routing strategy is effective on improving reasoning performance on frontier models. Comparing performance with and without reasoning strategy selection under the same setting, we observe that the dynamic strategy significantly improves model performance in most cases, particularly for GPT-4o. This improvement is especially pronounced for the CoT-Sym method, which, leveraging formalization for reasoning. Using adaptive reasoning, this method shows substantial gains on both GPT-4o and DeepSeek-V3.1 (e.g., 33.4% improvement on ProntoQA for DeepSeek-V3.1 and 22.4% for GPT-4o). This demonstrates that the dynamic routing strategy not only plays a crucial role in enabling adaptive reasoning within a neuro-symbolic framework, but also serves as effective guidance that can enhance the reasoning capabilities of language models themselves.

Reasoning strategy and model size affect formalization and reasoning performance. Although our framework demonstrates strong performance with both GPT-4o and DeepSeek-V3.1, DeepSeek-

V3.1 lags noticeably behind GPT-4o despite its better routing accuracy. This performance gap stems primarily from differences in formalization quality (see Section 3.7 for detailed analysis), which is influenced by both reasoning type complexity and model capability. The impact of reasoning type is evident in the FOLIO results in Table 1, where both models perform significantly worse on FOLIO (FOL problems) compared to ProntoQA and ProofWriter (LP problems), despite all involving deductive reasoning. This can be explained by the difference in formalizing first-order vs propositional logic. The effect of model size on formalization is observable not only in the comparison between GPT-4o and DeepSeek-V3.1, but will be further demonstrated in the next section through experiments with smaller open-source models.

3.3 Results on Smaller Models

We experimented with small-scale open-source models to examine their performance and limitations in adaptive reasoning, with results shown in Table 2 (see Appendix E for complete results with standard deviations). We observe that the dynamic strategy selection maintains strong performance on small-scale models, achieving over 90% accuracy on Qwen-series models and 76.8% on Llama-3.1-8B. Overall, the neuro-symbolic framework achieves the best performance using Qwen-2.5-Coder-7B. However, a notable performance gap exists compared to closed-source models, supporting our earlier conclusion on model size’s effect on autoformalization. Notably, while Qwen2.5-7B’s Routing Accuracy (97.7%) nearly matches GPT-4o (98.0%), its reasoning performance lags significantly behind. Upon further analysis, we find that:

Pre-trained smaller models lack structural abstraction over formal languages necessary for external solver integration. Recall that the LP solver, FOL solver, CSP solver, and SMT solver are implemented in Prover9, Pyke, MiniZinc, and Z3, all ‘under-resourced’ formal languages. Without knowledge of the target syntax, small models often produce invalid formal representations.

However, performance of smaller models can be significantly improved via supervised fine-tuning. Despite fine-tuning with only 5K samples, we observe significant performance improvements across all models (33% for Llama-3.1-8B, 37.5% for Qwen-2.5-7B, and 33.4% for Qwen2.5-

Models	Mixed	
	random routing	w/ routing
GPT-4o	29.0	92.1
DeepSeek-V3.1	32.8	73.4
Llama-3.1-8b	26.0	41.0
Qwen-2.5-7b	29.6	50.1
Qwen-2.5-Coder-7b	20.5	59.6

Table 3: Ablation performance with random routing.

Coder-7B). This further demonstrates the critical impact of autoformalization quality on our framework’s capabilities and highlights persisting deficiencies of small-scale models for effective neuro-symbolic reasoning.

3.4 Ablation Study on Adaptive Reasoning

In this section, we analyze the effectiveness of two essential components of our framework: the adaptive reasoning (Routing) process and the symbolic reasoning process.

Effectiveness of adaptive reasoning. A core feature of the adaptive neuro-symbolic framework lies in the ability to dynamically select and integrate an appropriate solver for reasoning through the Router. To demonstrate the effectiveness of this component, we ablate it by randomly assigning solvers to problems. Note that in a random solver setting, the decomposed problem components may be incompatible with the assigned solver (e.g., a CSP solver cannot extract the required constraints from an LP problem). Therefore, we allow the solver to directly formalize the original input text. We evaluate GPT-4o, DeepSeek-V3.1, and three small models. For each model, we conduct experiments with 3 random seeds, with average results shown in Table 3. The ablations demonstrate that the dynamic routing mechanism is fundamental for effective neuro-symbolic integration, and ablating it leads to significant performance degradation across all models, which demonstrates the importance of incorporating adaptive reasoning strategies for addressing heterogeneous problems.

Effectiveness of symbolic solvers. Another important feature of our framework is its integration with external symbolic solvers. To isolate the impact of this component, we evaluate the LLMs’ performance when provided with the target reasoning type selected by the routing module but bypassing the call to the external solver, relying entirely on LLMs’ inference to derive the answer. These ablations directly correspond to the ‘w/ rout-

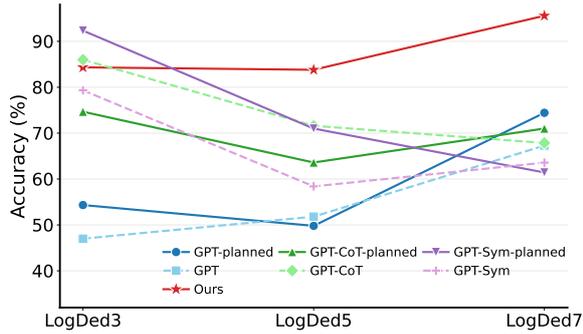


Figure 2: Model performance on Logic Deduction tasks of varying difficulty.

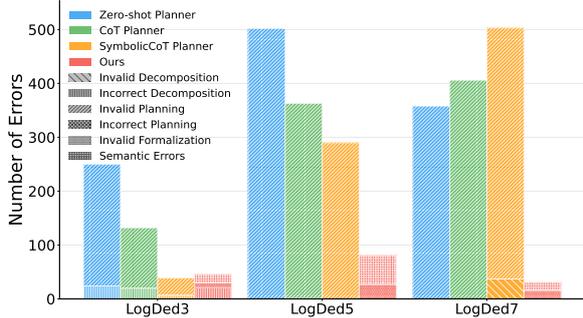


Figure 3: Distribution of error types of our framework for Logic Deduction tasks.

ing” columns in Tables 1 and 2: while “Zero-shot” and “CoT” reason directly over natural language, CoT_{sym} reasons over the formalized code without external execution. The results demonstrate that using a symbolic solver outperforms direct LLM inference in most cases across various settings. For instance, our framework surpasses the strongest alternative by 17% and 6% for GPT-4o and DeepSeek-V3.1 on the mixed dataset, respectively. This gap demonstrates that even when the correct reasoning paradigm is identified, the external solver remains an essential consistency anchor for executing deterministic logic with a level of rigor that surpasses current LLMs’ capabilities.

3.5 Performance Across Varying Difficulty Levels

To investigate how the framework performs across problems of varying difficulty levels, we selected the Logical Deduction subset from BigBench (Srivastava et al., 2023), which requires deducing the order of a sequence of objects from a minimal set of conditions and contains an increasing number of variables (i.e. 3, 5, and 7, denoted as LogDed3, LogDed5, and LogDed7, respectively). These subsets form an evaluation suite with increasing levels

Methods	Overall Acc.	
	w/o routing	w/ routing
GPT-Zero-Shot	24.5	<u>27.0</u>
GPT-CoT	27.3	<u>28.0</u>
GPT-CoT_{sym}	21.0	<u>38.2</u>
Ours	-	54.4

Table 4: Performance on the sequential multi-question reasoning task. Bold indicates the best result for each dataset. Underline indicates improvement after adding routing information.

of reasoning complexity. Given GPT-4o’s strong capabilities in formalization and logical understanding, we choose it as the backbone model to examine how different methods perform under varying logical deduction difficulty.

Figure 2 highlights our framework’s superior robustness over pure LLM methods. As complexity scales from 3 to 7 objects, Zero-shot and CoT baselines fluctuate significantly, whereas our framework maintains a stable accuracy profile by offloading the combinatorial search to the symbolic engine. The slightly lower performance on LogDed3 compared to LogDed7 is elucidated by the error analysis in Figure 3: shorter problem descriptions provide fewer structural cues for the Router, resulting in higher *Incorrect Decomposition* rates. Conversely, the explicit constraints in LogDed7 provide stronger pattern-matching signals. This demonstrates that once the reasoning paradigm is correctly identified, the symbolic engine ensures performance remains anchored by formal logic, effectively improving robustness with increased problem difficulty.

3.6 Performance in Multi-Question Scenarios

Real-world scenarios often present composite queries with sub-problems requiring different reasoning paradigms. In this section, we evaluate our proposed framework in a more complex, compositional setting, where models are required to sequentially solve multiple reasoning problems. Inspired by the experimental setup in Reasoning Evaluation through Simultaneous Testing (REST) (Pan et al., 2025), we extend our evaluation to require the solution of multiple problems in a single pass.

Specifically, similarly to Pan et al. (2025), we randomly shuffle the Mixed dataset and group the data into batches of three examples, concatenating each group into a single input prompt for joint inference. We define *Overall Accuracy* to measure

global-level reasoning consistency, which requires all three answers to be correct. Details on the construction of this evaluation setting are described on Appendix I. GPT-4o is used as the reference model for this evaluation given its performance on the mixed dataset.

Table 4 shows the performance of all settings for the multi-question task, where we find that:

Explicit problem decomposition is essential for composite reasoning tasks. Our framework outperforms the second-best method by 16.2% in *Overall Accuracy*, demonstrating the critical value of systematic parsing and planning for composite reasoning tasks.

Routing integration provides universal benefits for cross-task consistency. Consistent *Overall Accuracy* improvements across all Router-integrated methods (2.5% for GPT-Zero-shot, 0.7% for GPT-CoT, and 17.2% for GPT-CoT_{sym}) reveals the impact of joint decomposition and coordination across different reasoning types, which prevent reasoning drift across tasks. This suggests that maintaining logical coherence across composite tasks requires explicit orchestration lacking in monolithic approaches.

Adaptive neuro-symbolic integration mitigates reasoning drift in composite tasks. Table 4 reveals that pure LLMs suffer from “reasoning drift,” where local errors accumulate and degrade global coherence, resulting in low overall accuracy (~27%) on multiple reasoning problems. In contrast, our adaptive neuro-symbolic framework more than doubles this performance (54.4%) by providing a symbolic anchor. By offloading sub-problems to formal solvers, we prevent the semantic interference inherent in monolithic autoregressive chains, ensuring a level of cross-task consistency that pure LLM reasoning cannot maintain.

3.7 Error Analysis

For each input problem, our framework proceeds through five stages: (1) problem decomposition, (2) routing, (3) autoformalization, (4) symbolic reasoning, and (5) output translation. Since reasoning relies on deterministic symbolic reasoners with formal rules, we observe errors primarily under the generative-dominated stages (1), (2), (3) and (5).

Specifically, we categorize the failure cases into the following six types: 1. **Invalid Decomposition:** Invalid problem decomposition or reasoning

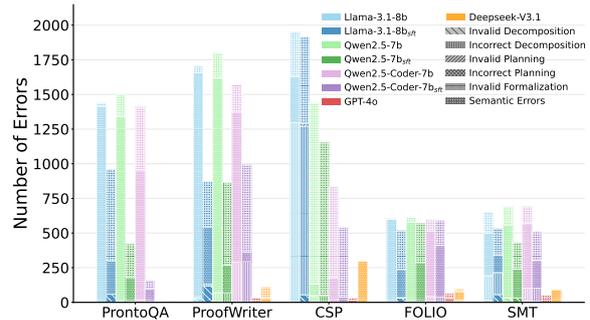


Figure 4: Distribution of error types of our framework. Each bar represents the number of incorrect predictions made by a model on a given dataset. Within each bar, different fill patterns indicate different error types.

type classification. 2. **Incorrect Decomposition:** Misclassification of the problem type. 3. **Invalid routing:** Invalid or empty output generated by the Router, causing failure in composing solvers. 4. **Incorrect routing:** Using inappropriate solver for solution. 5. **Invalid Formalization:** Formal code that does not conform to the target solver’s syntax, causing the reasoner to fail. 6. **Semantic Errors:** All remaining incorrect predictions that do not fall into the above five categories. Semantic errors typically arise from two sources: (i) the LLM generates syntactically valid but semantically incorrect formal code (e.g., wrong variable assignments or misaligned rules), leading the solver to an incorrect conclusion; or (ii) the LLM mistranslates the solver’s output during the final decoding step, mapping it to the wrong label.

Error taxonomy reveals model-specific failure regimes: small models failed early in autoformalization but can be improved via fine-tuning, GPT-4o fails post-reasoning. We report error type distributions of different models in our framework in Figure 4. As noted in Section 3.2, open-source models demonstrate limited formalization capabilities, consequently degrading system performance. Figure 4 reveals that *Invalid Formalization* constitutes the predominant error type for open-source models (including DeepSeek-V3), suggesting systematic failures in producing syntactically valid symbolic representations, particularly for domain-specific formal languages. However, such errors are significantly reduced via fine-tuning. In contrast, GPT-4o’s errors are primarily categorized as *Semantic Errors*, with only a few cases of *Invalid Formalization* observed on the FOLIO and CSP datasets.

4 Related Work

Neuro-symbolic methods pair LLMs with symbolic solvers to provide deterministic inference. Early strategies like LogicLLaMA (Yang et al., 2023a) and Sat-LM (Ye et al., 2023) focus on FOL or SAT translation, while Ishay et al. (Ishay et al., 2023) and [LLM]+ASP (Yang et al., 2023b) target Answer Set Programming. To mitigate formalization errors, LINC (Olausson et al., 2023) utilizes K -way majority voting. Despite their success, these methods are typically restricted to a single, predetermined paradigm, limiting their flexibility in heterogeneous real-world scenarios where reasoning requirements are unknown.

Building on these foundations, modular architectures have emerged to support a broader spectrum of reasoning tasks. Logic-LM (Pan et al., 2023) empowers LLMs with a suite of solvers and introduces self-refinement to fix syntax errors based on solver feedback. VERUS-LM (Callewaert et al., 2025) adopts the Knowledge Base Paradigm, separating domain knowledge from queries to facilitate knowledge reuse for tasks like optimization and explanation. However, these systems still largely rely on static solver integration that requires the target paradigm to be predetermined or explicitly specified by the user, which is addressed through adaptive solver orchestration in this paper.

Beyond external integration, some approaches embed symbolic structure directly within the reasoning trace. SymbCoT (Xu et al., 2024) performs reasoning internally via step-by-step symbolic CoT, while CLOVER (Ryu et al., 2025) and $\forall\text{uto}\exists\text{val}$ (Karia et al., 2024) prioritize atomic decomposition and formalization fidelity. Refinement has shifted from syntactic corrections to semantic validation; for instance, VERUS-LM (Callewaert et al., 2025) introduces satisfiability checks to identify logical contradictions. Unlike internal traces that can suffer from semantic drift in complex sequences, our approach utilizes symbolic solvers as logical anchors to ensure verified consistency across composite reasoning challenges.

5 Discussion

Our research focuses on the development of an adaptive agentic reasoning framework centered on three critical pillars: reasoning type identification, dynamic solver routing, and auto-formalization. The experimental results suggest that through structured pattern-matching prompting, the ma-

jority of current LLMs – including smaller open-source models – can effectively categorize reasoning paradigms and identify the appropriate solvers with high precision. This capability to navigate the reasoning space provides a robust foundation for the construction of adaptive agentic systems, as it demonstrates that even modest-scale models can serve as reliable orchestrators when provided with clear categorical priors. While these findings validate the feasibility of our orchestration logic, we acknowledge that the current benchmarks remain relatively controlled; subsequent research will seek to evaluate this framework in more heterogeneous and practical scenarios where reasoning requirements are more implicit and complex.

On the other hand, our findings clearly indicate that the primary factor hindering overall reasoning performance is the auto-formalization stage. While models are adept at selecting the correct tool, they frequently struggle to generate the precise, error-free syntax required by specialized formal engines such as Prover9 or MiniZinc. This “formalization gap” is particularly evident in smaller models, where the majority of failures stem from syntactical invalidity rather than logical misidentification. However, our experiments offer an optimistic takeaway: this limitation is not an inherent deficit in reasoning capacity but rather a modular challenge that can be significantly mitigated through targeted fine-tuning. For future researchers, this suggests that the path toward more capable neuro-symbolic agents lies in enhancing the reliability of formalization interfaces – perhaps through intermediate representations or iterative feedback loops – to ensure that the material-formal integration remains as robust as the underlying paradigm selection.

6 Conclusion

In this paper, we present an adaptive and extensible framework for end-to-end LLM-supported symbolic reasoning. By decomposing the reasoning process into three explicit stages – problem decomposition, routing, and solver-based reasoning, our framework provides a unified interface for handling a diverse range of reasoning tasks. Its agent-based design allows for a dynamic and extensible interface of reasoning functions. Auto-formalization cycles, closely connected to a set of symbolic solvers ensure that reasoning is formally grounded, enabling both interpretability and epistemic trust.

Limitations

While our framework demonstrates promising results for adaptive LLM-symbolic reasoning, several limitations warrant discussion:

Model Scale Dependencies. Our experiments reveal a strong dependency on model scale for effective symbolic integration. Pre-trained smaller models (7B-8B parameters) show significant performance degradation when integrated with external solvers, achieving only 8.0% accuracy on mixed datasets compared to 92.1% with GPT-4o. This limitation suggests that effective LLM-symbolic reasoning may be constrained to frontier models or require substantial fine-tuning investments for smaller models as we have explored in this paper.

Formalization Bottleneck. The framework’s performance is fundamentally limited by the quality of auto-formalization. Invalid formalization constitutes the primary error source for smaller models, indicating that the translation from natural language to formal representations remains a critical bottleneck. While our results show that fine-tuning can significantly improve formalization capabilities (improving performance from 8.0% to 41.0% on Llama-3.1-8b), this dependency on formalization quality may still limit the framework’s applicability to domains with complex or non-standard logical structures. Few-shot learning approaches or domain-specific fine-tuning strategies are ways to address this limitation.

Solver Coverage. While our framework supports four solver types (LP, FOL, CSP, SMT), it does not cover the full spectrum of formal reasoning paradigms. Mathematical reasoning, probabilistic inference, temporal logic, and other specialized reasoning domains are not addressed. However, the proposed model is extensible to additional solvers as new reasoning agents can be added to the portfolio without modifying the core routing framework. This extensibility represents a key design principle that should enable future expansion to other reasoning domains as new solvers become available.

Acknowledgments

This work was partially funded by the Swiss National Science Foundation (SNSF) projects RATIONAL and M-RATIONAL.

References

- BIG bench authors. 2023. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *Transactions on Machine Learning Research*.
- Benjamin Callewaert, Simon Vandeveld, and Joost Vennekens. 2025. Verus-lm: a versatile framework for combining llms with symbolic reasoning. *arXiv preprint arXiv:2501.14540*.
- Leonardo de Moura and Nikolaj Bjørner. 2008. Z3: An efficient smt solver. In *Tools and Algorithms for the Construction and Analysis of Systems*, pages 337–340, Berlin, Heidelberg. Springer Berlin Heidelberg.
- DeepSeek-AI. 2024. [Deepseek-v3 technical report](#). *Preprint*, arXiv:2412.19437.
- Bruce Frederiksen. 2008. Applying expert system technology to code reuse with pyke. *PyCon: Chicago*.
- Franciszek Górski, Oskar Wysocki, Marco Valentino, and Andre Freitas. 2025. Integrating expert knowledge into logical programs via llms. *arXiv preprint arXiv:2502.12275*.
- Daya Guo, Dejian Yang, Haowei Zhang, Ruoyu Zhang Junxiao Song, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, and et al. Xiao Bi. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *Preprint*, arXiv:2501.12948.
- Simeng Han, Hailey Schoelkopf, Yilun Zhao, Zhenyuan Qi, Martin Riddell, Wenfei Zhou, James Coady, David Peng, Yujie Qiao, Luke Benson, Lucy Sun, Alexander Wardle-Solano, Hannah Szabó, Ekaterina Zubova, Matthew Burtell, Jonathan Fan, Yixin Liu, Brian Wong, Malcolm Sailor, and 16 others. 2024. [FOLIO: Natural language reasoning with first-order logic](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 22017–22031, Miami, Florida, USA. Association for Computational Linguistics.
- Adam Ishay, Zhun Yang, and Joohyung Lee. 2023. [Leveraging large language models to generate answer set programs](#). In *Proceedings of the 20th International Conference on Principles of Knowledge Representation and Reasoning*.
- Rushang Karia, Daniel Bramblett, Daksh Dobhal, Pulkit Verma, and Siddharth Srivastava. 2024. [AutoEval: Autonomous assessment of llms in formal synthesis and interpretation tasks](#). *Preprint*, arXiv:2403.18327.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *arXiv preprint arXiv: 2205.11916*.
- Llama team. 2024. The llama 3 herd of models. *arXiv preprint arXiv: 2407.21783*.

- Tharindu Madusanka, Marco Valentino, Iqra Zahid, Ian Pratt-Hartmann, and Riza Theresa Batista-Navarro. 2025. Unravelling the logic: Investigating the generalisation of transformers in numerical satisfiability problems. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 25155–25168.
- W. McCune. 2005–2010. Prover9 and mace4. <http://www.cs.unm.edu/~mccune/mace4/>.
- Nicholas Nethercote, Peter J Stuckey, Ralph Becket, Sebastian Brand, Gregory J Duck, and Guido Tack. 2007. Minizinc: Towards a standard cp modelling language. In *International conference on principles and practice of constraint programming*, pages 529–543. Springer.
- Theo Olausson, Alex Gu, Ben Lipkin, Cedegao Zhang, Armando Solar-Lezama, Joshua Tenenbaum, and Roger Levy. 2023. Linc: A neurosymbolic approach for logical reasoning by combining language models with first-order logic provers. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5153–5176.
- OpenAI. 2024. Gpt-4o system card. *arXiv preprint arXiv: 2410.21276*.
- Liangming Pan, Alon Albalak, Xinyi Wang, and William Wang. 2023. Logic-lm: Empowering large language models with symbolic solvers for faithful logical reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3806–3824.
- Zhuoshi Pan, Qizhi Pei, Yu Li, Qiyao Sun, Zinan Tang, H. Vicky Zhao, Conghui He, and Lijun Wu. 2025. Rest: Stress testing large reasoning models by asking multiple problems at once. *Preprint*, arXiv:2507.10541.
- Xin Quan, Marco Valentino, Danilo Carvalho, Dhairya Dalal, and Andre Freitas. 2025a. PEIRCE: Unifying material and formal reasoning via LLM-driven neuro-symbolic refinement. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 11–21, Vienna, Austria. Association for Computational Linguistics.
- Xin Quan, Marco Valentino, Louise Dennis, and Andre Freitas. 2024a. Enhancing ethical explanations of large language models through iterative symbolic refinement. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1–22, St. Julian’s, Malta. Association for Computational Linguistics.
- Xin Quan, Marco Valentino, Louise A. Dennis, and Andre Freitas. 2024b. Verification and refinement of natural language explanations through LLM-symbolic theorem proving. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 2933–2958, Miami, Florida, USA. Association for Computational Linguistics.
- Xin Quan, Marco Valentino, Louise A. Dennis, and Andre Freitas. 2025b. Faithful and robust LLM-driven theorem proving for NLI explanations. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 17734–17755, Vienna, Austria. Association for Computational Linguistics.
- Qwen Team. 2024. Qwen2.5 technical report. *arXiv preprint arXiv: 2412.15115*.
- Hyun Ryu, Gyeongman Kim, Hyemin S Lee, and Eunho Yang. 2025. Divide and translate: Compositional first-order logic translation and verification for complex logical reasoning. In *The Thirteenth International Conference on Learning Representations*.
- Abulhair Saparov and He He. 2023. Language models are greedy reasoners: A systematic formal analysis of chain-of-thought. In *The Eleventh International Conference on Learning Representations*.
- Ian Soboroff. 2021. Overview of trec 2021. In *TREC*.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R. Brown, Adam Santoro, Aditya Gupta, Adria Garriga-Alonso, Agnieszka Kluska, Aitor Lewkowycz, Akshat Agarwal, Alethea Power, Alex Ray, Alex Warstadt, Alexander W. Kocurek, Ali Safaya, Ali Tazarv, and 431 others. 2023. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *Trans. Mach. Learn. Res.*, 2023.
- Oyvind Tafjord, Bhavana Dalvi, and Peter Clark. 2021. ProofWriter: Generating implications, proofs, and abductive statements over natural language. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3621–3634, Online. Association for Computational Linguistics.
- Yuhuai Wu, Albert Qiaochu Jiang, Wenda Li, Markus Rabe, Charles Staats, Mateja Jamnik, and Christian Szegedy. 2022. Autoformalization with large language models. *Advances in neural information processing systems*, 35:32353–32368.
- Jundong Xu, Hao Fei, Liangming Pan, Qian Liu, Mong-Li Lee, and Wynne Hsu. 2024. Faithful logical reasoning via symbolic chain-of-thought. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13326–13365, Bangkok, Thailand. Association for Computational Linguistics.
- Yuan Yang, Siheng Xiong, Ali Payani, Ehsan Shareghi, and Faramarz Fekri. 2023a. Harnessing the power of large language models for natural language to first-order logic translation. *arXiv preprint arXiv:2305.15541*.
- Zhun Yang, Adam Ishay, and Joohyung Lee. 2023b. Coupling large language models with logic programming for robust and general reasoning from text. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 5186–5219.

Xi Ye, Qiaochu Chen, Isil Dillig, and Greg Durrett. 2023. Satlm: Satisfiability-aided language models using declarative prompting. In *Proceedings of NeurIPS*.

Lan Zhang, Marco Valentino, and Andre Freitas. 2025. Autoformalization in the wild: Assessing llms on real-world mathematical definitions. *Preprint*, arXiv:2502.12065.

Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lema Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi. 2023. Siren’s song in the ai ocean: A survey on hallucination in large language models. *arXiv preprint arXiv: 2309.01219*.

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

A Solvers Used in Our Framework

Our framework includes a set of specialized, instantiated agents for formal inference. All formal agents are derived from a common base class, `LLMToolAgent`, which is instantiated with a specific backend solver and associated methods for *formalizing*, *reasoning*, and *converting*. The *formalizing* method translates the natural language problem into a formal representation compatible with the syntax of the embedded backend reasoning engine. The solver then performs *reasoning* over this formalized input using the backend engine. Finally, the *conversion* method maps the reasoning output back to the corresponding option label among the candidate answers provided in the original problem. Our solvers fall into two main categories: solvers for logical deduction and solvers for constraint satisfiability.

A.1 Solvers for Logical Deduction

Our framework includes a set of pre-defined instantiated solvers to support logical reasoning tasks. Using the controlled autoformalization method described in PEIRCE (Quan et al., 2025a), each solver employs LLM-based auto-formalization to convert natural language problems into formal representations, with iterative refinement capabilities to automatically correct syntax errors when the generated code returns an error message. Two back-end solvers are implemented in our framework: *Pyke*

(Frederiksen, 2008) and *Prover9* (McCune, 2005–2010) as backend-solvers, allowing **Logic Programming (LP)** and **First-order Logics (FOL)** respectively, a set which can be expanded by instantiating the `LLMToolAgent` interface.

A.2 Solvers for Constraint Satisfiability

Two other solver types are implemented: **Satisfiability Modulo Theories (SMT)** and **Constraint Satisfaction Problems (CSP)** solvers. The SMT Solver uses the *Z3* theorem, auto-formalizing natural language constraints into the SMT-LIB format, then querying *Z3* to determine satisfiability. This approach is particularly effective for matching problems involving integer constraints, such as determining patient eligibility for clinical trials based on age, disease stage, and other medical criteria. The CSP Solver employs the *MiniZinc* constraint modeling language (Nethercote et al., 2007) and the *Gecode* solver to handle Constraint Satisfaction Problems, also auto-formalizing natural language problem descriptions into *MiniZinc* models, then finding all valid solutions that satisfy the given constraints. This solver type is well-suited for scheduling and assignment problems, such as matching research papers to conference sessions while respecting topic compatibility and capacity constraints.

B Execution Semantics

In this section, we provide a detailed description of the workflow execution process. Formally, given the question components Q and predicted reasoning types \mathcal{T} , the Router dynamically instantiates and composes an executable workflow π . It is part of the Router to assess the required solvers associated with the task, in case there are no suitable solver to solve the task, the Router will instantiate a neural solver modules with LLM. Besides, we define Memory as a shared cache for intermediate results among modules during execution. Each module first retrieves the required inputs from Memory and write down the running results. By doing this, each module relies solely on Memory for its input parameters, which unifies the interface across various modules.

The execution workflow outputs (i) the set of required modules as execution nodes, and (ii) a set of directed edges (source module \rightarrow target module) as follows:

$$\pi = (V, E, \prec), \quad (1)$$

which is specified by:

- $V \subseteq \mathcal{S} \cup \mathcal{Q} \cup \widehat{\mathcal{S}}$: a finite set of execution nodes. In this framework, we instantiate each question in \mathcal{Q} as a static module, so as to write the problem data into Memory. $\mathcal{S} = \{S_{LP}, S_{FOL}, S_{CSP}, S_{SMT}\}$ denotes the symbolic solver for different reasoning types. $\widehat{\mathcal{S}} = \{\widehat{S}_1, \dots\}$ is an unbounded pool of dynamic solver modules that may be instantiated on-the-fly to solve the problem.
- $E \subseteq V \times V$: a set of directed edges. $(u, v) \in E$ means “pass the output of u as (part of) the input of v ”.
- \prec : a strict total order on V that yields the execution sequence $u_1 \prec u_2 \prec \dots \prec u_\ell$, $u_i \in V$.

For each module $u \in V$, we assign a unique module id and denote its output as $o(u)$. Let $in(v) = \{u \mid (u, v) \in E\}$ denote the set of predecessors of v . To ensure that module v receives the correct outputs from its predecessors, each module u stores its output in memory with its agent id, for example, “result_[ID]”. Before invoking module v , the framework gathers the outputs of all module in $in(v)$ with corresponding module ids and assembles them as input to v :

$$i(v) = \{o(u) \mid (u, v) \in in(v)\}. \quad (2)$$

The invocation $v(i(v)) \rightarrow o(v)$ is side-effect free except for potential updates to v ’s private MEMORY.

To better organize the execution results, we introduce two special virtual modules: <START> and <END>. We connect <START> to all source nodes (i.e., nodes with no incoming edges), and direct all terminal nodes (i.e., nodes with no outgoing edges) to <END>. The <END> module aggregates the outputs from all execution paths, ensuring that the overall output is complete. For example, when solving multiple subproblems, the Planner invokes different agents to handle each subproblem, and <END> collects all the predicted results into a list and returns them.

C Details of Evaluating Datasets

Below we provide the detailed information of the dataset used in this paper.

1. **PrOntoQA** (Saparov and He, 2023) is a synthetic dataset designed to evaluate the deductive reasoning ability of language models. Following the Logic-LM (Pan et al., 2023) setup, we adopt the most challenging subset from PrOntoQA, which involves five-hop reasoning over fictional characters. This subset consists of 500 examples, each requiring the model to determine the truth value of a derived fact.
2. **ProofWriter** (Tafjord et al., 2021) is a widely used dataset for testing deductive reasoning, where the questions are expressed in natural language. To reduce computational cost, we use the 600-example subset sampled by Logic-LM (Pan et al., 2023) from the depth-5 partition under the open-world assumption. This subset includes balanced labels from PROVED, DISPROVED, UNKNOWN, with each instance represented as a (premise set, query) pair.
3. **FOLIO** (Han et al., 2024) is a challenging dataset for first-order logic reasoning, composed of expert-written problems grounded in real-world knowledge. The examples involve complex semantic structures and require multi-step inference. We evaluate on the full test set, which contains 204 examples.
4. **LogDed7** (bench authors, 2023) is constructed using a subset of the LogicalDeduction dataset from BigBench (bench authors, 2023). This task challenges models to determine the sequential ordering of objects based on a minimal set of logical constraints. From the original dataset, we selected the most challenging problems that involve 7 objects, which includes 700 samples.
5. **TREC_{trials}** (Soboroff, 2021) data including 300 trial-patient pairs, which is built upon the 2021 Clinical Trials Track from TREC to create a trial-patient matching challenge. Our dataset comprises 75 patients, where each patient is paired with both two compatible clinical trials and two incompatible trials. The ground truth matching labels are based on expert annotations and should be considered approximate rather than definitive, since patient descriptions typically do not cover all the variables specified in the trial criteria. As a result, matching decisions often involve informed

guesses and therefore depend on experts’ personal judgments.

D Evaluation and Fine-tuning Details

D.1 Selection criteria on Foundational models

GPT-4o selection: GPT-4o serves as our primary large-scale baseline due to its state-of-the-art performance across diverse tasks and widespread industry adoption. Its consistent API availability ensures reproducible experimental conditions.

DeepSeek selection: We selected DeepSeek-V3.1 as the representative large-scale open-source baseline due to its state-of-the-art performance for reasoning.

7B-8B parameter range: We selected models in the 7B-8B parameter range as they represent the optimal balance between computational feasibility and performance capability. This size range enables deployment on single high-end GPUs while maintaining sophisticated reasoning abilities, making them practical alternatives to larger models.

Model diversity: Our selection encompasses proprietary (GPT-4o) versus open-source models, general-purpose versus code-specialized variants, and different institutional approaches (DeepSeek-AI Team, OpenAI Team, Meta Team, and Qwen Team). The 7B-8B range enables direct comparison between open-source alternatives (Llama-3.1-8B vs. Qwen-2.5-7B) and evaluation of specialization effects (Qwen-2.5-7B vs. Qwen-2.5-Coder-7B).

Table 5 presents the specifications and characteristics of each model.

D.2 Evaluation Details

Here we provide the LLM settings compared in this paper.

1. **Zero-shot:** In the zero-shot setting, the LLM is asked to directly provide the final answer.
2. **Chain-of-Thought (CoT):** In the CoT setting, the LLM is instructed to output both the answer and a step-by-step reasoning process.
3. **Symbolic CoT (CoT_{sym}):** In the Symbolic CoT setting, the LLM is first prompted to translate the problem into a formal language representation, and then perform reasoning with LLM based on the formalization.

Additionally, to validate the effectiveness of the routing mechanism within our framework, we integrate the Router’s outputs as auxiliary prompts into the aforementioned baseline settings. Specifically, we prepend the Router’s output to the input prompt as an auxiliary hint before the original question, allowing the LLM to perform reasoning with this additional guidance. In particular, for CoT_{sym}, after the Router invokes the corresponding solver, we directly provide the resulting formalization to the LLM, instead of requiring the model to generate it on its own. This setup enables us to observe how the Router influences model performance across different prompting strategies.

For GPT-4o, we use the default API parameters to obtain responses. For open-source models, we set the maximum number of new tokens to 4096 and the temperature to 0.01. Besides, all models are required to return their predictions in JSON format using the keyword “ans”. For the CoT and Symbolic CoT settings, models are also required to include their reasoning process with the keyword “reasoning”.

To manage inference costs, we evaluate GPT-4o and DeepSeek-V3.1 using pass@1 metric, which measures accuracy based on a single inference run. For each small model, since different settings use the same pretrained version for both routing and parsing, we perform routing and parsing once, and share the results across all inference settings. We conduct experiments three times with different random seeds and report the average performance across the three runs.

D.3 Fine-tuning details

Due to the poor understanding capabilities of open-source models for non-mainstream programming languages, we constructed formalization datasets for relevant reasoning engines and fine-tuned these models to enhance their performance within the our framework framework.

Our training data is primarily constructed according to the tasks used in our experiments. We begin by selecting problems for formalization across different reasoning engines:

- For the **LP solver**, we use problems from the training portion of ProofWriter provided by Logic-LM and one-hop problems from Pron-toQA.
- For the **FOL solver**, we select the training portion of FOLIO provided by Logic-LM.

Model	Parameters	Context Length	Key Features/Specialties	License/Access
GPT-4o	-	128K tokens	Multimodal, advanced reasoning	API-based (gpt-4o)
DeepSeek-V3.1	671B	128K tokens	Hybrid inference, Improved tool / agent capability	MIT License
Llama-3.1-8B -Instruct	8.03B	128K tokens	Training: 15T tokens	Llama 3.1 Community License
Qwen-2.5-7B -Instruct	7.07B	32K tokens (expandable to 128K)	Multilingual, reasoning tasks	Apache 2.0
Qwen-2.5-Coder -7B-Instruct	7.07B	32K tokens	Code generation, 40+ programming languages	Apache 2.0

Table 5: Model Specifications and Characteristics

Dataset	Solver Type	# Examples
ProntoQA	LP Solver	1,390
ProofWriter	LP Solver	1,398
FOLIO	FOL Solver	776
Logical Deduction	CSP Solver	1,185
TREC	SMT Solver	376
Total	All Solvers	5,125

Table 6: Training Data Statistics for Each Reasoning Engine

- For the **CSP solver**, we use the training portion of LogDed provided by Logic-LM.
- For the **SMT solver**, we use the TREC 2022 Clinical Trials Track.

For each problem, we utilize GPT-4o to obtain the corresponding formalization code for each reasoning engine. To enhance data quality, we execute the formalization with the corresponding reasoning engines and retain only the code that runs successfully. The final data quantities obtained for each component are shown in Table 6.

Using the Llamafactory (Zheng et al., 2024) training framework, we employed LoRA fine-tuning for all models to prevent overfitting given the limited training samples. The LoRA configuration was set with rank = 16 and applied to all modules. For all models, we set the sequence_length to 3072, per_device_train_batch_size to 2, and gradient_accumulation_steps to 8. The learning_rate was selected from the set $\{1e - 3, 1e - 4\}$ based on optimal performance on the validation set. All models were trained for three epochs. All training experiments are conducted on two H100 GPUs.

E Complete Results for Small Models Evaluation

For small-scale model evaluation in Section 3.3, we conduct experiments with 3 random seeds and report averaged performance. Here we present the complete results, including means and standard deviations in Table 7.

F Prompt for Text Parsing

We use the prompt in Figure 5 and Figure 6 to decompose the input text.

G Prompt for Routing

The Router uses prompt in Figure 7 to select task-relevant agents and construct an execution plan.

H Prompt Templates for Textual Instance Construction

To support end-to-end evaluation where models does not have the prior knowledge of task-specific information, we convert each data instance into the natural language format using dataset-specific prompt templates. These templates integrate the key components of each problem (e.g., premises, hypotheses, questions) into a single input text.

Specifically, for PrOntoQA, ProofWriter, and FOLIO, we apply the following templates to convert structured fields into natural language text.

Prompt for Text Parsing: Part 1

SYSTEM: You are a logician and reasoning systems expert specializing in symbolic reasoning frameworks. Given a text that may contain one or multiple logical reasoning problems, identify each problem, determine its type, and decompose the text accordingly. Return the result strictly as a JSON object with "result" containing an array of problem objects.

Specifically, your task is to:

1. First, analyze the input text to identify how many distinct reasoning problems it contains.
2. For each identified problem, determine its type from the following categories:
 - Logic Programming (LP): Problems where conclusions are typically deduced step by step from a set of known facts and rules. These problems often involve applying simple logical rules repeatedly to infer new facts until the goal is reached.
 - First-order Logic (FOL): Problems that require more expressive reasoning, such as statements like "for all" or "there exists", and complex relationships among multiple entities.
 - Constraint Satisfaction Problem (CSP): Problems that involve finding assignments of values to variables within finite domains such that all explicit or implicit constraints are satisfied. These often include tasks like ordering, allocation, or scheduling.
 - Boolean Satisfiability (SAT): These problems involve determining whether all logical constraints in a system are simultaneously satisfied. In the context of reasoning tasks, SAT typically focuses on checking if a particular configuration or entity conforms to a set of conditions expressed as logical formulas. Unlike CSP, which searches for value assignments over finite domains, SAT emphasizes verifying logical consistency in given assignments and is often applied to analytical reasoning questions.

To guide your classification:

- Consider whether the problem leans more towards reasoning from facts and rules (often LP or FOL) or checking constraints and conditions (often CSP or SAT).
 - If the focus is on assigning values or arranging elements under constraints, it is more typical of CSP.
 - If the focus is on verifying whether one given description satisfies the logical requirements of another, it is more typical of SAT. Analytical reasoning problems are often classified as SAT.
 - Between LP and FOL, use LP if the reasoning relies on simple rules and chaining; use FOL if it involves richer logical expressions with quantifiers or complex entity relationships.
3. For each problem, create a JSON object with the following structure:
 - "problem_id" (str): A unique identifier following the pattern "ques_1", "ques_2", etc., based on the order of appearance.
 - "problem_type" (str): The type classification. The value must be one of LP, FOL, CSP, SAT.
 - Based on the problem type, include the appropriate fields:
 - If problem_type == "LP" or "FOL":
 - "premise" (str): the given premise.
 - "hypothesis" (str): the hypothesis to be evaluated for truth.
 - "options" (list): the provided answer options.
 - If problem_type == "CSP":
 - "context" (str): background description.
 - "question" (str): the specific question being asked.
 - "options" (list): the provided answer options.
 - If problem_type == "SAT":
 - "trial_description" (str): description of the trial.
 - "sample_description" (str): description of the sample.
 - "options" (list): the provided answer options.

Preserve any existing option labels (e.g., "A", "B"). If options have no labels, assign labels 'A', 'B', 'C', ... automatically.

(Part 2 of the prompt will follow.)

Figure 5: Prompt for text parsing.

Prompt for Text Parsing: Part 2

(Following Part 1 of the prompt for text parsing.)

4. Extract or analyze the overall goal of the input text:

- FIRST, try to extract any explicitly stated overall goal or instruction from the text (e.g., “Answer the above questions one by one”, “Solve all problems to find the final answer”, etc.)
- If no explicit goal is found, analyze the relationship between problems and write a brief description:
- Multiple independent problems: “Solve multiple independent reasoning problems”
- Subproblems contributing to main problem: “Solve subproblems to address the main complex problem”
- Sequential dependent problems: “Solve problems in sequence with dependencies”
- Single problem: “Solve the reasoning problem”

Return a JSON object with two keys:

- “result”: an array containing all identified problems
- “overall_goal”: the extracted goal text or a brief analysis-based description

Example output format:

```
“result”: [
  “problem_id”: “ques_1”,
  “problem_type”: “CSP”,
  “context”: “...”,
  “question”: “...”,
  “options”: [“A) ...”, “B) ...”]
],
“overall_goal”: “Answer the above questions one by one”
USER:
Problem Statement:
problem
```

Figure 6: Prompt for text parsing.

Prompt for Routing

Design a plan that uses the minimal number of agents necessary to achieve the goals. You may be given one or multiple problems to solve, each with a unique problem_id. Select agents only from the provided list. Output a JSON object describing the plan. Requirements:

1. Use the MINIMAL number of agents needed to complete all tasks.
2. The output MUST be a JSON object with exactly two keys:
 - ‘agents’: an array of selected agent names (strings), including any problem_ids that serve as starting points
 - ‘edges’: an array of [source, target] pairs (both strings) showing execution order.
3. For multi-problem scenarios, start the execution flow from the respective problem_ids.
4. The plan MUST end with the special control marker <END>.
5. IMPORTANT: If you need to use the same agent type for multiple different problems, distinguish them by adding ‘:’ followed by a sequence number. For example, if you need two CSP solvers for different problems, use ‘csp_solver:1’ and ‘csp_solver:2’.
6. You can use the same agent for multiple problems if appropriate, but ensure proper sequencing.

Example for multiple problems:
If given ‘QID [ques_1]: CSP scheduling problem’, ‘QID [ques_2]: FOL reasoning task’ and ‘QID [ques_3]: CSP allocation problem’, your output might include agents: [‘ques_1’, ‘ques_2’, ‘ques_3’, ‘csp_solver:1’, ‘fol_solver:1’, ‘csp_solver:2’, ‘<END>’] with edges: [[‘ques_1’, ‘csp_solver:1’], [‘ques_2’, ‘fol_solver:1’], [‘ques_3’, ‘csp_solver:2’], [‘csp_solver:1’, ‘<END>’], [‘fol_solver:1’, ‘<END>’], [‘csp_solver:2’, ‘<END>’]].

7. Respond with **only valid JSON**, without any explanations, markdown formatting, or code fences. Do not wrap the output in “‘json or any other delimiters. Return pure JSON.

Here is the goal: [GOAL]
Problems to solve: [PROBLEMS]
Current agents:
[PORTFOLIO]

Figure 7: Prompt for routing.

Methods	PrOntoQA		ProofWriter		FOLIO		LogDed7		TREC _{trials}		Mixed		Routing
	w/o routing	w/ routing	w/o routing	w/ routing	w/o routing	w/ routing							
Llama-3.1-8b													
Zero-shot	65.2±0.5	58.6±0.3	41.8±0.2	36.3±0.1	46.7±0.2	42.6±0.8	37.0±0.1	25.3±0.2	63.2±0.2	39.7±0.9	48.7±0.1	38.8±0.3	76.8±2.7
CoT	57.9±0.7	54.1±0.7	45.3±0.8	35.4±1.2	38.2±0.4	42.3±0.2	19.5±0.6	24.0±1.4	63.6±0.6	48.3±0.5	42.0±0.5	38.3±0.5	
CoT _{sym}	10.1±1.0	0.0±0.0	10.8±0.6	0.2±0.2	12.1±1.7	0.2±0.2	3.9±0.2	0.6±0.1	2.6±0.4	3.1±0.4	7.6±0.3	0.7±0.1	
Ours	-	3.9±1.4	-	5.3±1.1	-	0.8±0.5	-	7.0±0.2	-	27.6±1.7	-	8.0±0.9	
Ours _{sft}	-	75.7±13.9	-	51.5±11.8	-	31.5±5.0	-	10.1±1.4	-	40.9±2.8	-	41.0±7.3	
Qwen-2.5-7b													
Zero-shot	57.2±0.0	40.9±0.2	37.8±0.1	29.8±0.2	32.4±0.0	40.0±0.2	42.9±0.1	38.5±0.4	59.3±0.0	53.2±0.2	45.9±0.0	38.8±0.2	97.7±0.1
CoT	95.3±0.4	72.8±0.2	59.8±0.7	59.1±0.1	48.0±1.2	47.7±1.7	46.8±1.2	20.6±0.7	47.0±0.8	25.3±0.3	60.8±0.3	45.0±0.1	
CoT _{sym}	55.1±0.1	46.1±0.2	41.7±0.2	29.7±0.6	59.2±0.2	52.3±0.5	8.0±0.6	49.5±0.8	60.2±0.3	61.9±1.1	38.3±0.2	45.5±0.2	
Ours	-	0.0±0.0	-	0.0±0.0	-	0.0±0.0	-	31.5±0.3	-	23.1±0.6	-	12.6±0.2	
Ours _{sft}	-	71.8±8.2	-	51.9±5.6	-	6.4±1.2	-	45.0±0.4	-	51.9±2.2	-	50.1±3.2	
Qwen-2.5-Coder-7b													
Zero-shot	71.5±0.1	67.8±0.0	53.6±0.2	56.5±0.2	45.3±0.2	25.7±0.5	45.8±0.2	44.1±0.1	66.2±0.2	66.7±0.3	56.0±0.1	53.8±0.1	92.9±0.1
CoT	70.5±0.6	64.8±0.3	51.3±0.1	53.3±0.3	45.4±0.2	24.8±0.6	40.5±0.4	40.4±0.3	63.0±0.3	65.2±0.6	53.2±0.1	50.9±0.2	
CoT _{sym}	78.2±14.6	29.8±0.3	46.8±9.6	20.6±1.0	52.1±9.5	13.2±0.7	14.0±2.7	11.3±1.0	38.0±8.3	29.3±2.6	43.0±8.4	20.3±0.3	
Ours	-	6.1±0.2	-	12.8±0.5	-	2.8±0.2	-	60.2±0.4	-	23.2±1.0	-	26.2±0.1	
Ours _{sft}	-	89.4±0.2	-	44.9±0.5	-	3.1±0.5	-	74.2±0.2	-	43.3±1.2	-	59.6±0.1	

Table 7: Average performance of small open-source models over three random seeds with standard deviations.

Method	LP (in/out)	FOL (in/out)	CSP (in/out)	SAT (in/out)
<i>GPT-4o</i>				
Zero-shot	208.7 / 10.1	186.7 / 10.8	225.0 / 10.9	609.6 / 10.6
Zero-shot (r)	1705.3 / 209.4	1661.5 / 178.2	1732.0 / 247.5	2308.2 / 406.1
CoT	228.7 / 268.9	206.7 / 302.2	244.0 / 356.4	629.6 / 440.4
CoT (r)	1740.3 / 675.6	1685.6 / 713.1	1778.0 / 841.1	2334.8 / 782.9
SymCoT	1706.7 / 871.6	1684.7 / 450.4	1722.0 / 446.6	2107.6 / 521.7
SymCoT (r)	4601.6 / 980.8	4248.4 / 737.4	3486.0 / 696.5	5184.9 / 1336.9
Ours	3760.6 / 885.6	3968.4 / 548.0	3600.5 / 585.1	4570.7 / 1144.5
<i>Qwen-2.5-Coder-7b</i>				
Zero-shot	213.8 / 14.0	190.5 / 14.1	231.1 / 14.0	644.9 / 14.0
Zero-shot (r)	234.8 / 217.5	211.5 / 256.2	251.1 / 265.9	665.9 / 222.9
CoT	1709.4 / 233.5	1653.7 / 257.5	1747.1 / 265.9	2461.1 / 539.1
CoT (r)	1755.1 / 607.6	1728.0 / 447.0	1794.1 / 677.4	2505.1 / 791.0
SymCoT	1734.8 / 613.5	1711.5 / 611.3	1751.1 / 655.2	2165.9 / 606.1
SymCoT (r)	4177.1 / 866.3	4414.0 / 613.8	3527.5 / 1303.3	5297.7 / 1683.2
Ours	3606.4 / 969.6	3444.0 / 516.3	3853.0 / 591.2	4598.9 / 1626.8

Table 8: Average input and output tokens for each problem, where ‘(r)’ denotes the results with routing strategy.

Prompt
STATEMENT: {premise}
QUESTION: {conclusion}
{options}

For CSP task, we use the following template:

Prompt
You get a trial and a patient and have to say if there is a match:
TRIAL: {trial_description}
PATIENT: {patient_description}
Does the patient match the trial?
A) True
B) False

Prompt
STATEMENT: {input}
Which of the following is true?
{options}

For SMT task, we use the following template:

Each prompt is filled with instance-specific content to form the complete input, which is then passed to the model without revealing dataset boundaries or task types. This ensures the evaluation closely reflects practical use cases where models must perform problem-type identification and reasoning jointly.

Prompt for Merging Questions

Answer the following questions one by one.
 Q1:{question1}
 Q2:{question2}
 Q3:{question3}

Figure 8: Prompt for merging questions.

Trial Description in Natural Language

Inclusion Criteria

- Adult patients (age >18 years) with acute pancreatitis according to the revised Atlanta criteria.
- Informed consent.
- Known time of debut of symptoms.

Exclusion Criteria

- Chronic pancreatitis.
- Pregnancy.
- Known malignant disease.
- More than 72 hours from debut of symptoms to inclusion.

Figure 9: Trial Description in natural language.

I Prompt Templates for Generating Multi-Question Prompts

We use the prompt format in Figure 8 to concatenate multiple questions into a single input prompt.

J Example for Clinical Trial Matching

An example of an obtained z3 formalization of clinical trial-patient pair.

As shown in Figure 9 and Figure 10, formalizing the clinical trial description involves declaring all mentioned variables and the eligibility rules – which are often divided into inclusion criteria and exclusion criteria.

The patient description needs to be formalized accordingly: in such a way that the z3 code can run on the trial-patient pair, as shown in Figure 11 and 12.

Note that the SMT only sets the parameter values that are present in the formalized clinical trials description, as any other values will not play any role in determining if there is a match.

The agent also adds comments to justify the parameter values it sets. This is important, as it will

Trial Description in Formalized Form

```

; Declare variables and their types
(declare-const age_in_years Int)
(declare-const acute_pancreatitis Bool)
(declare-const informed_consent Bool)
(declare-const time_of_debut_of_symptoms Real)
(declare-const chronic_pancreatitis Bool)
(declare-const pregnancy Bool)
(declare-const malignant_disease Bool)
; Assert inclusion criteria
(assert (and (> age_in_years 18)
acute_pancreatitis informed_consent (>
time_of_debut_of_symptoms 0) ))
; Assert exclusion criteria
(assert (not (or chronic_pancreatitis
pregnancy malignant_disease (>
time_of_debut_of_symptoms 72) )))

```

Figure 10: Trial Description in formalized form.

Patient Description in Natural Language

The patient is a 57-year-old man with abdominal pain and vomiting. The pain started gradually about 20 hours ago in the epigastric and periumbilical regions, radiating to his back. He drinks around 60 units of alcohol per week and smokes 22 cigarettes per day. He is otherwise healthy with no history of allergies or medications. His family history is positive for type 2 diabetes (father and sister).
 On examination, the abdomen is tender and soft, bowel sounds are normal, heart rate is 115/min, and blood pressure is 110/75 mmHg. Laboratory results show leukocytosis (19.5), urea of 8.5, high CRP (145), high amylase (1200), and glucose of 15. Cross-sectional imaging was negative for obstructive pancreatitis. The diagnosis of acute pancreatitis is confirmed.

Figure 11: Patient description in natural language.

Patient Description in Formalized Form

```

; Assign patient values
(assert (= age_in_years 57))
(assert (= acute_pancreatitis true)) ;
Confirmed based on clinical and lab data
(assert (= informed_consent true)) ;
Assumed given for participation
(assert (= time_of_debut_of_symptoms 20))
(assert (= chronic_pancreatitis false))
(assert (= pregnancy false)) ; Male patient
(assert (= malignant_disease false))
; Check eligibility
(check-sat)

```

Figure 12: Patient in formalized form.

often need to infer likely values that are not mentioned in the patient description. Indeed, in our

Method	LP	FOL	CSP	SAT
<i>GPT-4o</i>				
Zero-shot	1.100	1.086	1.207	1.218
Zero-shot (r)	3.907	5.690	4.336	6.449
CoT	5.994	7.122	9.160	9.348
CoT (r)	10.337	15.158	13.698	14.262
SymCoT	14.198	8.098	9.035	13.230
SymCoT (r)	14.002	13.253	11.842	21.126
Ours	30.360	20.608	19.891	40.504
<i>Qwen-2.5-Coder-7b</i>				
Zero-shot	0.492	0.514	0.545	0.530
Zero-shot (r)	6.946	7.859	8.121	7.159
CoT	7.250	7.913	8.456	16.550
CoT (r)	17.051	12.367	20.249	21.809
SymCoT	17.698	16.974	19.048	16.800
SymCoT (r)	23.390	15.932	36.683	42.660
Ours	53.486	28.034	35.312	84.216

Table 9: Average inference time for each problem (seconds), where ‘(r)’ denotes the results with routing strategy.

Method	LP	FOL	CSP	SAT
Ours (Solver Only)	0.030	0.015	0.427	0.014

Table 10: Average symbolic solver execution time per problem (seconds).

example, the trial requires a confirmed diagnosis of acute pancreatitis according to the revised Atlanta criteria, but the patient description only provides indirect evidence – such as elevated amylase, characteristic pain, and inflammatory markers. The model therefore infers that the patient meets the diagnostic criteria and sets the variable for acute pancreatitis to true. Note that such cases of unmentioned parameters (that thus require abductive inferences) are the norm in the TREC dataset – and one of the reason the truth matching labels are approximate.

K Example for Other Solvers

Besides the SMT solver, here we also present formalization examples for the remaining solvers. Figure 13, Figure 14, and Figure 15 demonstrate the formalization processes for the LP solver, FOL solver, and CSP solver, respectively.

L Computational Cost Analysis

In this section, we provide a detailed analysis of the computational costs of our framework. We select 10 problems from each reasoning type and report the mean costs of GPT-4o and Qwen-2.5-Coder-7b.

L.1 Inference Time and Token Consumption

As shown in Table 8 and Table 9, our method incurs token consumption (Table 8) and higher total inference time (Table 9) due to the multi-stage architecture involving problem decomposition, routing, and auto-formalization. Each stage requires specific LLM calls, which explains the increased latency compared to zero-shot or single-pass CoT methods. This cost is a trade-off for the enhanced logical consistency and interpretability provided by the neuro-symbolic integration.

L.2 Symbolic Solver Execution Time

Notably, as shown in Table 10, the actual execution time of the symbolic solvers (Pyke, Prover9, Z3, and MiniZinc) is minimal, often representing less than 2% of the total inference time. This indicates that the primary computational bottleneck resides in the LLM-based formalization and routing phases, rather than the formal reasoning engine itself.

Formalization of LP Solver

Premise:

Every dumpus is not red. Tumpuses are red. Dumpuses are impuses. Impuses are not feisty. Impuses are yumpuses. Stella is a dumpus.

Hypothesis:

Stella is not red.

Formalization:

Predicates:

dumpus(\$x, bool) ::: Is x a dumpus?
red(\$x, bool) ::: Is x red?
tumpus(\$x, bool) ::: Is x a tumpus?
impus(\$x, bool) ::: Is x an impus?
feisty(\$x, bool) ::: Is x feisty?
yumpus(\$x, bool) ::: Is x a yumpus?
isA(\$x, \$type, bool) ::: Is x of type \$type?

Facts:

dumpus(Stella, True) ::: Stella is a dumpus.

Rules:

dumpus(\$x, True) >> red(\$x, False) ::: Every dumpus is not red.
tumpus(\$x, True) >> red(\$x, True) ::: Tumpuses are red.
dumpus(\$x, True) >> impus(\$x, True) ::: Dumpuses are impuses.
impus(\$x, True) >> feisty(\$x, False) ::: Impuses are not feisty.
impus(\$x, True) >> yumpus(\$x, True) ::: Impuses are yumpuses.

Query:

red(Stella, False) ::: Stella is not red.

Figure 13: Formalization of LP Solver.

Formalization of FOL Solver

Premise:

All birds have wings. No animal with wings is a reptile. Some animals that fly are birds. If something is an iguana, then it is a reptile. John is neither an iguana nor a bird.

Hypothesis:

John is a reptile.

Formalization:

Predicates:

Bird(x) ::: x is a bird.
Wings(x) ::: x has wings.
Animal(x) ::: x is an animal.
Reptile(x) ::: x is a reptile.
Fly(x) ::: x can fly.
Iguana(x) ::: x is an iguana.

Premises:

$\forall x (\text{Bird}(x) \rightarrow \text{Wings}(x))$::: All birds have wings.
 $\forall x (\text{Wings}(x) \rightarrow \neg \text{Reptile}(x))$::: No animal with wings is a reptile.
 $\exists x (\text{Fly}(x) \wedge \text{Bird}(x))$::: Some animals that fly are birds.
 $\forall x (\text{Iguana}(x) \rightarrow \text{Reptile}(x))$::: If something is an iguana, then it is a reptile.
 $\neg \text{Iguana}(\text{john}) \wedge \neg \text{Bird}(\text{john})$::: John is neither an iguana nor a bird.

Conclusions:

Reptile(john) ::: John is a reptile.

Figure 14: Formalization of FOL Solver.

Formalization of CSP Solver

Premise:

The following paragraphs each describe a set of five objects arranged in a fixed order. The statements are logically consistent within each paragraph.

In a golf tournament, there were five golfers: Rob, Ada, Dan, Joe, and Mel. Ada finished above Mel. Mel finished above Dan. Joe finished below Dan. Ada finished second.

Formalization:

Formalization:

```
include "globals.mzn";
```

```
% Define the golfers
```

```
enum GOLFER = Rob, Ada, Dan, Joe, Mel;
```

```
% Each golfer has a unique finishing position from 1 (first) to 5 (last)
```

```
array[GOLFER] of var 1..5: pos;
```

```
% All golfers must finish in distinct positions
```

```
constraint all_different([pos[g] | g in GOLFER]);
```

```
% Ada finished above Mel
```

```
constraint pos[Ada] < pos[Mel];
```

```
% Mel finished above Dan
```

```
constraint pos[Mel] < pos[Dan];
```

```
% Joe finished below Dan
```

```
constraint pos[Joe] > pos[Dan];
```

```
% Ada finished second
```

```
constraint pos[Ada] = 2;
```

```
solve satisfy;
```

Ordering

left to right: (1 = leftmost, 5 = rightmost)

Figure 15: Formalization of CSP Solver.