

SLR: Automated Synthesis for Scalable Logical Reasoning

Lukas Helff^{1,2}, Ahmad Omar¹, Felix Friedrich^{3*}, Antonia Wüst¹, Hikaru Shindo¹, Tim Woydt¹, Rupert Mitchell^{1,2}, Patrick Schramowski^{1,2,4,5}, Wolfgang Stammer^{1,2,6}, Kristian Kersting^{1,2,4,6,7}

¹TU Darmstadt ²hessian.AI ³Meta FAIR ⁴DFKI ⁵CERTAIN, Germany

⁶Lab1141 ⁷Centre for Cognitive Science, Darmstadt

Code: <https://github.com/ml-research/ScalableLogicalReasoning>

Data: <https://huggingface.co/datasets/AIML-TUDA/SLR-Bench>

Abstract

We introduce SLR, an end-to-end framework for systematic evaluation and training of Large Language Models (LLMs) via Scalable Logical Reasoning. Given a user’s task specification, SLR automatically synthesizes (i) an instruction prompt for an inductive reasoning task, (ii) a validation program, executable on model outputs to provide verifiable rewards, and (iii) the latent ground-truth rule. This process is fully automated, scalable, requires no human annotations, and offers precise control over task difficulty. Using SLR, we create SLR-BENCH, a benchmark comprising 19k prompts organized into 20 curriculum levels that progressively increase in relational, arithmetic, and recursive complexity. Large-scale evaluation reveals that contemporary LLMs readily produce syntactically valid rules, yet often fail at correct logical inference. Recent reasoning LLMs demonstrate improved performance but incur very high test-time computation, with costs exceeding \$300 for just 1,000 prompts. Finally, curriculum learning via SLR doubles Llama-3-8B accuracy on SLR-BENCH, achieving parity with Gemini-Flash-Thinking at a fraction of computational cost. Moreover, these reasoning capabilities generalize to a wide range of established benchmarks, underscoring the effectiveness of SLR for downstream reasoning.

1 Introduction

Logical reasoning is a fundamental aspect of intelligence, yet state-of-the-art AI systems still struggle with tasks that require robust reasoning and systematic generalization (Delfosse et al., 2025; Kostikova et al., 2025; Woydt et al., 2025; Wüst et al., 2025; Helff et al., 2025; Sinha et al., 2019). Existing benchmarks intended to evaluate reasoning capabilities, however, primarily emphasize *deductive* reasoning, where conclusions necessarily follow from given premises. This includes tasks such as math

word problems (Hendrycks et al., 2021) and logic puzzles (Lin et al., 2025; Xie et al., 2025; Liu et al., 2025). *Inductive* reasoning, by contrast, involves inferring general rules or patterns from specific examples, which remains particularly challenging and underexplored in large language models (Luo et al., 2024; Xie et al., 2024) (see also Tab. 1).

Current evaluation frameworks commonly employ constrained formats (e.g., multiple-choice) or rely on other LLMs as judges (Patel et al., 2024a; Lin, 2024; Lin et al., 2024), making it difficult to assess whether models genuinely understand logical structure or are merely exploiting superficial patterns in the data. Moreover, as training sets grow, benchmark items or their paraphrases increasingly overlap with pre-training data, making apparent reasoning abilities potentially just memorization (Shojaee et al., 2025; Xie et al., 2024).

To tackle these challenges, we introduce **SLR** (Scalable Logical Reasoning), an open-source framework for evaluating and training models in inductive logical reasoning. Given a user-defined logic task (Fig. 1, left), the task synthesizer (center) automatically generates novel reasoning tasks of controllable complexity. Each task comes with (i) a latent ground-truth rule, (ii) an executable validation program, and (iii) an instruction prompt. The ground-truth rule serves as the reference answer, while the validation program deterministically evaluates any candidate hypothesis. SLR supports both systematic model evaluation (Fig. 1, top right) and downstream model training, via supervised finetuning or reinforcement learning with rewards provided by the integrated symbolic judge (Fig. 1, bottom right). SLR’s fully symbolic and automated pipeline removes the need for human annotation and prevents dataset overlap.

Leveraging SLR, we present **SLR-BENCH** (Fig. 3), a “19k task benchmark” that forms a twenty-level curriculum of increasing logical complexity. These levels are further organized into four

*work done while at TU Darmstadt/hessian.AI/Lab1141

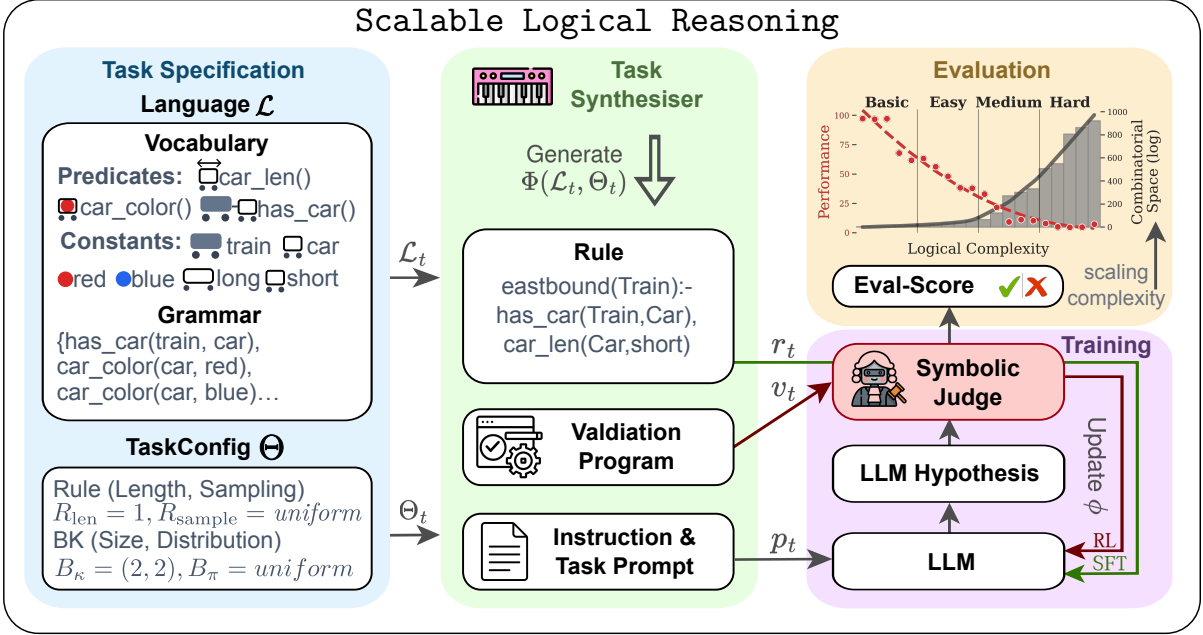


Figure 1: Overview of the **SLR** Framework, including task specification, automated task synthesis, training, and evaluation. **Left (blue)**: Language defines vocabulary and grammar, Task Config specifies configuration parameters for the synthesis. **Middle (green)**: The task synthesizer automatically generates ground-truth rules, validation programs, and instruction prompts. **Right (purple)**: Training LLMs on logic tasks via SFT (cross-entropy) or RL (symbolic judge feedback). **Right (orange)**: Evaluates LLMs using feedback provided by the symbolic judge. Arrows denote data and control flow through synthesis, prompting, evaluation, and downstream training loops.

curriculum tiers: *basic*, *easy*, *medium*, and *hard*. Each task is unique, with a precise and systematic assessment of inductive logical reasoning skills. In evaluations, we find that while LLMs are generally well-versed in generating valid rules, robust logical reasoning remains challenging. Performance declines sharply as task complexity increases. Scaling model size brings only marginal improvements, while scaling test-time compute boosts reasoning, but returns diminish as complexity rises.

Beyond benchmarking, SLR enables curriculum learning, boosting reasoning both in-domain and across established reasoning benchmarks. SLR-tuned models not only surpasses conventional LLMs on SLR-BENCH, but also outperforms reasoning LLMs, such as Gemini-2.0-flash-thinking, while using fewer inference tokens. Notably, these enhanced reasoning capabilities generalize downstream (e.g., GPQA (Rein et al., 2024), and CLUTRR (Sinha et al., 2019)).

In sum, our contributions are: (i) SLR, an open framework for automated synthesis and symbolic evaluation of logical reasoning in LLMs; (ii) SLR-BENCH, a 19k-task benchmark organized as a 20-level curriculum of increasing logical complexity, enabling both training and evaluation across a con-

trolled reasoning spectrum; (iii) a large-scale evaluation of LLMs on SLR-BENCH, revealing key insights and trade-offs in model performance; (iv) curriculum learning with SLR substantially improves both in-domain and downstream reasoning.

2 Related Work

Evaluating LLMs’ Logical Reasoning. Tab. 1 provides an overview of existing logical reasoning benchmarks in terms of inference types, dataset origins, and evaluation formats. Notable datasets include LogiQA/2.0 (Liu et al., 2020, 2023), FOLIO (Han et al., 2022) (deductive reasoning), AbductionRules (Young et al., 2022) (abductive reasoning), bAbI (Weston et al., 2015), and CLUTRR (Sinha et al., 2019) (synthetic QA with inductive reasoning). Aggregate testbeds such as BIG-Bench (Kazemi et al., 2025; Suzgun et al., 2023), HLE (et al., 2025), FineLogic (Zhou et al., 2025), and LogiGLUE (Luo et al., 2024) span a range of tasks and inference styles. Other benchmarks, such as Proofwriter, PrOntoQA, FLD, Multi-LogiEval, SynLogic, ZebraLogic, and the K&K Sandbox (Oyvind et al., 2021; Saparov and He, 2023; Morishita et al., 2024; Patel et al., 2024b; Liu et al., 2025; Lin et al., 2025; Xie et al., 2025,

Dataset	Reasoning Type	Data Creation	Evaluation Methodology	Task Synthesis	Custom Tasks	Curriculum Learning	Scalable Complexity
LogiQA (Liu et al., 2020)	Deduction	Human	MC	✗	✗	✗	✗
LogiQA 2.0 (Liu et al., 2023)	Mix	Human	MCQA/Auto	✗	✗	✗	✗
FOLIO (Han et al., 2022)	Deduction	Human	EM	✗	✗	✗	✗
AbductionRules (Young et al., 2022)	Abduction	Synthetic/Human	EM	✗	✗	✗	✗
FineLogic (Zhou et al., 2025)	Deduction	DS collection	Symbolic/LLM	✗	✗	✗	✗
HLE (et.al., 2025)	Mix	DS collection	LLM	✗	✗	✗	✗
Big-Bench (et.al., 2025)	Mix	DS collection	Mix	✗	✗	✗	✗
LogiGLUE (Luo et al., 2024)	Mix	DS collection	MCQA/Auto	✗	✗	✗	✗
Multi-LogiEval (Patel et al., 2024b)	Deduction	Synthetic	Symbolic	✗	✗	✗	✗
CLUTRR (Sinha et al., 2019)	Induction	Synthetic	EM	✗	✗	✗	✗
KOR-Bench (Ma et al., 2024)	Mix	Synthetic/Human	EM/Symbolic	(✓)	✗	✗	✗
PrOntoQA (Saparov and He, 2023)	Deduction	Synthetic	EM	✓	✗	✗	✗
bAbI (Weston et al., 2015)	Mix	Synthetic	EM	✓	✗	✗	✗
SynLogic (Liu et al., 2025)	Deduction	Synthetic	EM	✓	✗	✗	✗
FLD (Morishita et al., 2024)	Deduction	Synthetic	Symbolic	✓	✓	✗	(✓)
K&K (Xie et al., 2025, 2024)	Deduction	Synthetic	EM	✓	✗	(✓)	(✓)
ZebraLogic (Lin et al., 2025)	Deduction	Synthetic	EM	✓	✗	(✓)	(✓)
SLR (ours)	Induction	Synthetic	Symbolic	✓	✓	✓	✓

Table 1: Comparison of **logic reasoning benchmarks**. **Reasoning Type**: Logical inference type (deduction, induction, abduction). **Creation**: Dataset origin (synthetic, human-annotation, DS collection). **Evaluation**: Output scoring (symbolic execution, multiple choice (MC), LLM, exact match (EM)). **Task Synthesis**: Supports for tasks generation. **Custom Tasks**: User-defined task creation (via language, grammar, or setup). **Curriculum Learning**: Curriculum-based progression of difficulty. **Scalable Complexity**: Supports arbitrarily scaling task complexity. (✓): fully supported, ✗: not supported, (✓): partially/limited

2024) generate tasks from fixed ontologies, often with exact-match evaluation. Classic ILP datasets (e.g. Mutagenesis (Debnath et al., 1991)) lack scalability and natural-language integration. In contrast, SLR introduces scalable, curriculum-based synthesis with controllable difficulty and verifiable evaluation, addressing key gaps in prior works.

Limits and Promises of Reasoning LLMs. LLMs like GPT-4 (Team, 2024b), Llama-3 (Team, 2024a), and Qwen (Bai et al., 2023) can handle basic reasoning and coding tasks but often struggle with true abstraction (Shojaee et al., 2025; Xie et al., 2024). Recent *reasoning LLMs* attempt to bridge this gap by scaling *test-time* compute. Systems like OpenAI’s *o1/o3* (OpenAI, 2025) or DeepSeek-R1 (Team, 2025) generate and re-rank thousands of reasoning traces per query, achieving state-of-the-art results on, e.g., math or coding (Quan et al., 2025; Hendrycks et al., 2021; Rein et al., 2024; Gao et al., 2024a). However, these gains come at a steep cost (Fan et al., 2025; Kim et al., 2025). Some studies question whether such models truly learn logical structure or merely exploit surface-level patterns (Fan et al., 2025; Shojaee et al., 2025; Xie et al., 2024). Curriculum learning has been shown to enhance training robustness and generalization (Bengio et al., 2009; Bursztytn et al., 2022),

yet no prior framework offers a flexible framework for task synthesis for automatic curriculum generation with symbolic evaluation for reasoning at scale. SLR addresses this gap.

3 SLR: Automatic Benchmark Synthesis

SLR is a scalable methodology for systematically generating, evaluating, and training LLMs on inductive reasoning tasks. Its goal is to automate the creation of diverse, challenging logical reasoning benchmarks, embedded as natural language prompts, with model outputs that can be efficiently verified via symbolic execution of Inductive Logic Programming (ILP) programs (Muggleton and de Raedt, 1994; Cropper and Dumančić, 2022). The overall pipeline (Fig. 1) has three main stages: *task specification*, *synthesis*, and *evaluation/training*.

Task Notation. SLR follows the Learning-from-Entailment (LFE) paradigm (Raedt, 1997) in ILP, where each task is defined as $\mathcal{I} = (B, E^+, E^-)$ consisting of background knowledge B , positive examples E^+ , and negative examples E^- . A candidate hypothesis H solves the task if and only if $B \cup H \models E^+$ and $B \cup H \not\models E^-$. Logical entailment (\models) is evaluated via Prolog execution under the closed-world assumption (details see App. A).

Motivating Example. Consider a simple train do-

Algorithm 1 Task Synthesizer

Require: \mathcal{L} , B_π , κ_{pos} , κ_{neg} , R_{sample} , R_{len}

- 1: $B \leftarrow \emptyset$, $E^+ \leftarrow \emptyset$, $E^- \leftarrow \emptyset$
- # Rule Synthesis
- 2: $R^* \leftarrow \text{RULEGENERATOR}(\mathcal{L}, R_{\text{len}}, R_{\text{sample}})$
- 3: **while** $|E^+| < \kappa_{\text{pos}}$ **or** $|E^-| < \kappa_{\text{neg}}$ **do**
- # Background Synthesis
- 4: $b \leftarrow \text{BACKGROUNDGENERATOR}(\mathcal{L}, B_\pi)$
- 5: $(y, q) \leftarrow \text{ASSIGNLABEL}(R^*, b)$
- # Stratified Rejection Sampling
- 6: **if** $y = 1$ **and** $|E^+| < \kappa_{\text{pos}}$ **then** \triangleright accept positive
- 7: $B \leftarrow B \cup \{b\}$; $E^+ \leftarrow E^+ \cup \{q\}$
- 8: **else if** $y = 0$ **and** $|E^-| < \kappa_{\text{neg}}$ **then** \triangleright accept negative
- 9: $B \leftarrow B \cup \{b\}$; $E^- \leftarrow E^- \cup \{q\}$
- 10: **else**
- 11: **continue** \triangleright reject sample
- 12: **end if**
- 13: **end while**
- 14: $\text{program} \leftarrow \text{VALIDATIONPROGRAM}(B, E^+, E^-)$
- 15: $\text{prompt} \leftarrow \text{PROMPTGENERATOR}(B, E^+, E^-)$
- 16: **return** $(R^*, \text{program}, \text{prompt})$

main where each train consists of cars described by attributes like color, length, or roof type. The learning goal is to induce a rule, e.g., “the train has a yellow car.” Here, the *background knowledge* (B) contains all known facts about each train, e.g., which cars it includes and their attributes. The *positive examples* (E^+) all have yellow cars, while the *negative examples* (E^-) do not have any yellow cars. During synthesis, SLR automatically generates such tasks by varying both the background knowledge and the ground-truth rule, effectively scaling the complexity of the reasoning problem.

3.1 Task Specification (Input)

The SLR synthesizer is controlled by the Task Language \mathcal{L} , which defines the logical vocabulary and grammar, and the Task Configuration Θ , which controls the generation process (see Fig. 1, left).

Language Specification (\mathcal{L}): We define a language $\mathcal{L} = (\mathcal{V}, \mathcal{G})$ that specifies the building blocks for task generation. The Vocabulary \mathcal{V} comprises a set of constant, function, and predicate symbols that form the syntax for generating rules, examples, and background knowledge. The vocabulary induces the Herbrand base $\text{HB}(\mathcal{V})$, which is the set of all syntactically valid ground atoms (facts) (Lloyd, 2012). The Grammar \mathcal{G} is formed by a set of semantic rules that filter the Herbrand base to include only meaningful atoms, $\text{HB}_G(\mathcal{V})$. For instance, a color can be assigned to a car ($\text{car_color}(\text{car}, \text{red})$) but not to semantically incompatible objects.

Task Configuration (Θ): The configuration parameters $\Theta = \langle R_{\text{sample}}, R_{\text{len}}, B_\pi, \kappa \rangle$ give control over the synthesis process. The (i) Rule Sampling Policy

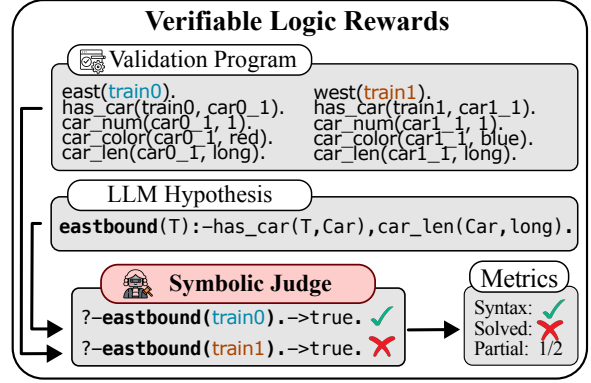


Figure 2: **Verifiable Logic Rewards:** A candidate hypothesis is evaluated by executing it against the validation program. It outputs three metrics: syntactic validity (binary), perfect task completion (binary), and a partial score for the fraction of correctly classified examples.

(R_{sample}) controls the synthesis of the ground truth rule R^* , which can either be sampled randomly (Uniform Sampling) or generated via an LLM (LLM-Guided Generation). To ensure the LLM produces diverse and challenging logic rules, we leverage an exhaustive prompt (see App. B.5) that covers a wide array of logical structures and Prolog features (containing arithmetics, recursions, variables, cuts, or comparison operators, etc.). The (ii) Rule Length (R_{len}) specifies the number of literals in the body of the ground-truth rule R^* . Structural predicates such as $\text{has_car}(T, C)$, which merely link entities and impose no constraints, are not counted. The (iii) Background Sampling Policy (B_π) defines a probability mass function that assigns a selection probability to each ground atom in the grammar-filtered Herbrand base $\text{HB}_G(\mathcal{V})$, enabling designers to encode priors on the data distribution (e.g., *uniform*). We further introduce mirror sampling, where backgrounds for (E^+ , E^-) are identical except for atoms relevant to R^* . For each positive example, a corresponding “mirror” negative is created by randomly altering the rule-specific atoms. The (vi) Problem Size ($\kappa = (\kappa_{\text{pos}}, \kappa_{\text{neg}})$) specifies the target number of positive ($|E^+|$) and negative ($|E^-|$) examples. This directly controls the size and class balance of each generated task.

3.2 Task Synthesis (Generation)

The task synthesizer (Fig. 1, center) is an automated process detailed in Alg. 1. Given a task specification, the synthesizer generates complete and solvable ILP problems. It comprises two main phases: rule synthesis and background synthesis.

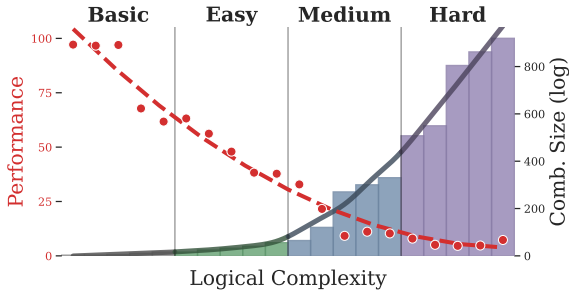


Figure 3: Overview of **SLR-BENCH**: The benchmark curriculum spans from basic to hard tasks with increasing logical and combinatorial complexity (bars, right y-axis). As logical complexity increases, Model performance (red, left y-axis) declines, highlighting current LLMs’ limitations on more challenging reasoning tasks.

Rule Synthesis (Alg. 1, line 2). The process begins with the **RULEGENERATOR** creating a latent, ground-truth rule R^* . This rule represents the underlying logical pattern that a model is expected to induce. The generation is guided by pre-defined parameters ($R_{\text{len}}, R_{\text{sample}}$) that control the length and generation policy for the rule. The resulting rule is a syntactically valid definite clause of the form $h : -b_1, \dots, b_{R_{\text{len}}}$.

Background Synthesis (Alg. 1, lines 3-13). Once R^* is fixed, the synthesizer enters a loop to construct the background knowledge B and the label sets E^+ and E^- . This loop executes three steps until the desired number of positive and negative examples is generated:

(i) **Sample Background:** The **BACKGROUND-GENERATOR** samples a set of ground atoms specifying the properties and relationships of the background instance. Ground atoms are drawn from the probability mass function B_π over $\text{HB}_G(\mathcal{V})$.

(ii) **Assign Label:** The function determines whether a query atom q (that is the ground atom of the target predicate h) is logically entailed by the sampled background b and the ground-truth rule R^* (i.e., whether $b \cup R^* \models q$ holds). This produces a label (positive or negative) for the query. We denote the labeling function as: $\text{ASSIGNLABEL}(R^*, b) = (1, q)$ if $b \cup R^* \models q$, and negative otherwise, $(0, q)$.

(iii) **Accept/Reject Sample:** To ensure the desired class balance, a stratified rejection sampling strategy is used to populate the example sets. The generated background b and query q are accepted only if the corresponding example set (E^+ or E^-) is not yet full, as specified by the task size parameter (κ). If accepted, b is added to the task’s main background knowledge B , and q is added to the

appropriate example set. Otherwise, it is discarded.

Synthesizer Outputs (Alg. 1, lines 14–17). For each task, the synthesizer generates three outputs: (1) the *latent ground-truth rule* R^* ; (2) a *validation program*, an executable logic program encoding (B, E^+, E^-) for automatic evaluation; and (3) an *instruction prompt* presenting the task in natural language or Prolog, ready for LLM input. See App. Fig. 6 for an example synthesis run.

3.3 Training and Evaluation

The final stage, shown in Fig. 1 (right), uses the synthesized task to evaluate and train models.

VERIFIABLE LOGIC REWARDS. In SLR verifiable rewards are implemented via the Symbolic judge (see Fig. 2), which is used for both training and evaluation. It deterministically assesses candidate hypotheses for logic tasks by executing them against a validation program and providing rewards signals. Specifically, it checks whether all positive examples (E^+) are entailed and all negative examples (E^-) are not entailed (see App. B.4).

Model Evaluation. SLR streamlines the creation of logical reasoning benchmark datasets for systematic model evaluation. By specifying various combinations of task language and configuration, users can automatically generate diverse tasks that span a broad range of domains, prompt styles, and reasoning complexities. Each synthesized task comprises a natural language prompt, a ground-truth rule, and an executable validation program.

Model Training. SLR enables automated training loops, with two types of model feedback. In supervised fine-tuning (SFT), the ground-truth rule R^* acts as the training target, enabling gradient-based updates from predicted rules. In reinforcement learning, the verifiable logic rewards provide a reward signal (RLVR) that guides policy updates. This cohesive pipeline enables scalable generation, evaluation, and training of LLMs, driving systematic advances in logical reasoning capabilities.

Novelty and Systematic Generalization. The generator enforces a strict separation between training and test splits, ensuring that ground-truth rules (R^*) and background knowledge (B) are entirely distinct. As SLR is fully synthetic, overlap with existing data is statistically negligible, making it robust against data leakage and memorization. This enables assessing whether the model is capable of systematically generalizing to entirely new rules.

Instruction & Task Prompt

You are a train classifier who is observing trains that are traveling either east- or westbound. Each train is composed of one or more cars, and each car is characterized by a set of properties, represented as ground atoms over a fixed set of predicates. The direction (eastbound or westbound) of a train is to be determined from its composition. To describe the trains we define a set of predicates and grounding domains:

```
'has_car(Train, Car)': Specifies that 'Car' is part of the train 'Train'.
'car_num(Car, CarNumber)': Specifies the position of the car within its train. 'CarNumber' is a positive integer.
'car_color(Car, Color)': Specifies the color of the car. 'Color' can be 'red', 'blue', 'green', 'yellow', or 'white'.
'car_len(Car, Length)': Specifies the length of the car. 'Length' can be either 'short' or 'long'.
'has_wall(Car, WallType)': Specifies the wall type of a car. 'WallType' can be either 'full' or a 'railing'.
```

You are provided with positive and negative examples in the form of eastbound(*t*) or westbound(*t*) for each train *t*, together with background knowledge consisting of ground facts over the above predicates which describe its composition.

```
eastbound(train0).                westbound(train1).
has_car(train0, car0_1).          has_car(train1, car1_1).
car_num(car0_1, 1).              car_num(car1_1, 1).
car_color(car0_1, red).          car_color(car1_1, red).
car_len(car0_1, long).           car_len(car1_1, short).
has_wall(car0_1, railing).       has_wall(car1_1, railing).
```

Your task is to formulate a hypothesis, i.e. a prolog rule of the form 'eastbound(*Train*) :- Body.' that correctly distinguishes eastbound from westbound trains. The hypothesis must be true for all positive examples (i.e., eastbound trains) and false for all negative examples (i.e., westbound trains). Aim to find the shortest correct rule, that is, one that uses the fewest possible body literals subject to the prior constraints. Your rule must use only predicates defined above and must perfectly separate eastbound from westbound trains.

Ground-Truth Rule

```
eastbound(Train) :- has_car(Train,Car), car_len(Car,long)
```

Figure 4: Illustrative **prompt** and **ground-truth rule** generated by SLR (Level 1, SLR-BENCH). Language (\mathcal{L}): 5 predicates, 1 car variable per train. Task configuration (Θ): $\kappa = (1, 1)$ (one positive and one negative example); $B_\pi = \text{mirror}$; $R_{\text{len}} = 1$; $R_{\text{sample}} = \text{uniform}$. The prompt provides the full ILP instance, including background B , positive/negative examples (E^+ , E^-), and natural-language instructions for the learning task.

4 SLR-BENCH: Instantiating SLR

With SLR-BENCH, we instantiate SLR as a 20-level curriculum of logical reasoning tasks with increasing complexity (see Fig. 3). Each level specifies its own language \mathcal{L} and configuration θ , producing a total of 19k generated reasoning tasks. Each level contains 1k train¹, 10 eval, and 50 test samples. Each task comes with (i) a generated latent ground-truth rule, (ii) the corresponding validation program, and associated instruction prompt for the task. An illustrative example for prompts and ground-truth rules can be found in Fig. 4.

Design rationale. The logic task is inspired by the V-LoL *trains* domain (Helff et al., 2025; Michalski, 1980; Mitchell, 1997), chosen for three main reasons. First, its hierarchical structure (trains \rightarrow cars \rightarrow attributes) naturally yields first-order rules richer than simple lookups yet more tractable than full theorem proving. Second, its small, discrete attribute domains enable precise control over task complexity. Third, it offers a clean proof of concept for scalable logic synthesis, being fully synthetic and easily extendable to new domains.

Languages. Each curriculum level is parameterized by level-specific language \mathcal{L} , detailed in App. B.2. The vocabulary includes mutually exclusive class labels *eastbound* and *westbound*, which serve as the targets for classification tasks

(E^+ , E^-). Background knowledge B is sampled from a vocabulary of five predicates (*has_car*, *car_num*, *car_color*, *car_len*, and *has_wall*) with their respective grounding domains defined in App. B.2. As curriculum levels increase, the vocabulary expands monotonically by introducing new predicates and grounding domains selected from a predefined set. Semantic coherence is ensured by constraining predicate groundings (e.g., only colors as arguments for *car_color*), and by enforcing mutually exclusive constraints across predicates (e.g., passenger cars cannot carry payloads).

Task configs. Each curriculum level is parameterized by level-specific settings of θ , summarized in App. Tab. 4 and supplied directly to the synthesizer (Alg. 1). Problem size (κ) increases steadily across levels, maintaining an equal balance of positive and negative samples. Levels 1–5 use a *mirror* sampling policy for background knowledge, generating simple, nearly identical east- and westbound trains that differ only in ground atoms relevant to R^* . From level 6 onward, the background is sampled uniformly from the filtered Herbrand base, increasing diversity. Rule generation is uniform for the basic levels; from level 6, 30% of rules are LLM-guided, introducing greater variety in variables, arithmetic, recursion, and more.

SLR-Curriculum. SLR-BENCH comprises 20 levels across four tiers: *basic*, *easy*, *medium*, and *hard*. Each level systematically increases complexity by

¹The trainsets for levels 1-3 are smaller (26, 234, and 793)

	Curriculum Learning (SLR-BENCH)					
	LRL (\uparrow 0-20)	Syntax(\uparrow %)	Basic(\uparrow %)	Easy(\uparrow %)	Medium(\uparrow %)	Hard(\uparrow %)
Llama3.1-8b-it	3.8	80	70	7	0	0
Llama3.1-8b-it-SLR	8.7 (+4.9)	100 (+20)	96 (+26)	54 (+47)	18 (+18)	5 (+5)
	Downstream Reasoning Performance					
	SLR-Homes (\uparrow %)	MMLU (\uparrow %)	MMLU-Stats (\uparrow %)	CLUTRR (\uparrow %)	LogiQA (\uparrow %)	LogiQA2 (\uparrow %)
Llama3.1-8b-it	10.3	63.3	42.6	25.5	30.1	34.3
Llama3.1-8b-it-SLR	42.3 (+32)	66.1 (+2.8)	59.7 (+17)	32.1 (+6.6)	31.0 (+0.9)	39.4 (+5.2)
	GPQA (\uparrow %)	GPQA-Ext. (\uparrow %)	GPQA-Dia. (\uparrow %)	ARC-Easy (\uparrow %)	ARC (\uparrow %)	HellaSwag (\uparrow %)
Llama3.1-8b-it	31.7	26.9	21.7	81.4	52.7	57.4
Llama3.1-8b-it-SLR	32.8 (+1.1)	33.0 (+6.1)	28.3 (+6.6)	82.8 (+1.4)	54.6 (+1.9)	58.9 (+1.5)

Table 2: **Curriculum Learning and Generalization.** Benchmark scores (\uparrow) for base and SLR-tuned models on SLR-BENCH and downstream benchmarks; LRL measuring cumulative curriculum progress. The tuned model surpasses the baseline across all curriculum stages, while generalizing to other downstream reasoning tasks.

expanding task size (κ), adding new car constants and predicates, lengthening rules, and varying both the background knowledge and rule sampling policy; see App. Sec. B, Tab. 4. As a result, the combinatorial space of possible tasks grows exponentially, and later levels become progressively harder and require deeper reasoning beyond surface cues. **SLR-Homes.** Alongside the main curriculum, we introduce an OOD dataset that instantiates SLR in a new symbolic domain. It uses a distinct vocabulary, relational structure, and target concept to define house classification tasks based on attributes. SLR-Homes enables the assessment of cross-domain generalization and comprises 500 tasks across five difficulty levels, increasing in complexity.

Intended Use. SLR-BENCH is designed for two complementary purposes. (1) As a *static* benchmark, it enables fine-grained evaluation of an LLM’s reasoning abilities across tasks of increasing logical complexity. It is also easily extensible to accommodate future improvements in model capabilities. (2) As a *dynamic* curriculum, it serves as a training backbone, supplying structured reasoning tasks and feedback to enhance reasoning in both conventional and reasoning LLMs.

5 LLMs Can’t Do Induction at Scale

We benchmark and train LLMs on SLR-BENCH, assessing reasoning, syntactic correctness, and computational efficiency across the four difficulty tiers: *basic*, *easy*, *medium*, and *hard*. Our analysis highlights key trends, common failure modes, and the effectiveness of curriculum-based logic-tuning.

Training Setup. We investigate how LLMs benefit from curriculum training on SLR-BENCH with SFT (for more details see App. Sec. D). To prevent data leakage, we ensure that no prompts or rules

from the test set are included in the training set.

Evaluation Setup. All models are evaluated in a zero-shot setting using SLR-BENCH prompts, with a single attempt per task (pass@1). We report the following metrics: (i) *Logical Reasoning Level (LRL)*: A cumulative score over all curriculum levels L , where $\#solved_\ell$ and $\#tasks_\ell$ denote the number of solved and total tasks at level ℓ : $LRL = \sum_{\ell=1}^L \frac{\#solved_\ell}{\#tasks_\ell}$ (ii) *Syntax Score*: The proportion of predicted logic rules that are syntactically valid. (iii) *Logical-Reasoning Accuracy*: The fraction of correct solutions per complexity tier. (iv) *Compute*: The aggregate completion tokens and computational cost for each the models. For further details on downstream evaluations, see App. Sec. D, and for pricing, refer to App. Tab. 5.

5.1 Analysis and Key Findings

In this section, we analyze the key effects of SLR training. We first report downstream reasoning gains for models trained with SLR, namely Llama-3.1-8B-IT (Tab. 2) and Qwen3-8B (App. Tab. 10). We then benchmark SOTA conventional and reasoning LLMs on SLR-BENCH (Tab. 3), with an extended leaderboard provided in App. Tab. 15.

SLR Boosts Downstream Reasoning. Curriculum learning on SLR-BENCH yields substantial gains in both in-domain and downstream reasoning (Tab. 2). SLR-tuned models consistently outperform conventional LLMs on SLR-BENCH (cf. Tab. 3) and even surpass dedicated reasoning models such as Gemini-2.0-flash-thinking, while requiring significantly fewer inference tokens and lower compute. Performance on SLR-Homes further demonstrates robust zero-shot transfer to previously unseen domains, suggesting that SLR learns domain-agnostic reasoning capabilities rather than

Model	LRL	Syntax	Logical-Reasoning Acc. (%) [↑]				Total Compute	
	([↑] 0-20)	Score ([↑] %)	Basic	Easy	Medium	Hard	Tokens (↓M)	Costs (↓\$)
o3	15.5	80	99	93	74	45	4.30	207.24
gpt-5	15.4	98	100	90	72	46	16.40	103.13
gpt-5-mini-high	14.2	0.94	99	83.0	63.0	38.0	13.12	27.98
gpt-5-mini	12.0	95	99	82	41	20	4.90	11.54
o1	11.9	68	92	89	41	15	5.19	364.72
gpt-5-mini-low	9.7	0.88	97	71.0	20.0	7.0	1.17	4.07
R1-Llama-70B ²	8.8	75	98	67	8	4	11.61	5.33
gpt-5-nano	8.5	99	97	61	10	2	6.17	2.81
Gemini-thinking ¹	8.6	83	93	65	13	1	—	—
gpt-4.5-prev	7.3	100	94	47	5	1	0.37	576.40
gpt-4o	6.2	100	93	29	2	0	0.26	20.03
Llama 3.3-70B	5.6	100	90	22	1	0	0.49	0.82
gpt-4-turbo	5.4	100	89	18	2	0	0.41	81.30
Llama 3.1-8B	3.8	80	70	7	0	0	0.20	0.14

¹Gemini-2.0-flash-thinking-exp-01-21 ²DeepSeek-R1-Distill-Llama-70B — information not available

Table 3: **SLR-BENCH Leaderboard.** We report the models’ Logical Reasoning Level (LRL), syntax score, stage-specific logical reasoning accuracy (basic, easy, medium, hard), total completion tokens, and inference cost. Higher LRL and accuracy indicate superior logical reasoning; lower compute, greater efficiency. Performance drops as complexity increases, while Reasoning LLMs (orange) consistently outperform conventional LLMs (blue).

dataset-specific heuristics. Beyond SLR-BENCH, these gains transfer to a broad set of established reasoning benchmarks, including MMLU High School Statistics (+17) (Hendrycks et al., 2021), CLUTRR (Sinha et al., 2019), LogicQA (Liu et al., 2020), LogicQA2 (Liu et al., 2023), ARC (Clark et al., 2018), HellaSwag (Zellers et al., 2019), and GPQA variants (Rein et al., 2024). We hypothesize that SLR training introduces an inductive bias toward systematic rule induction over structured relational contexts. Models are encouraged to internalize abstractions for identifying relational structure, composing and applying reasoning patterns. These abstractions generalize to downstream deductive and knowledge-intensive tasks, where performance likewise depends on structured reasoning.

Curriculum Order Matters. An ablation study (App. B.3) examines how curriculum order affects training performance by comparing ordered, random, and reverse progressions. Training with the ordered curriculum achieves the best results, confirming that structured progression supports smoother and more stable learning.

Curriculum Levels Modulate Task Complexity: LLMs Break Down as Complexity Increases. SLR-BENCH creates a controlled gradient in logical complexity as model performance steadily declines throughout the curriculum levels (cf. Fig. 3,

Tab. 3). Most models readily solve *basic* levels. Base LLMs already falter on *easy* ones, solving fewer than half. Reasoning LLMs perform better, yet their accuracy drops sharply at the *medium* tier, and none succeed on *hard*. This pattern is also reflected in the LRL score, empirically indicating how far each model can progress before performance collapses. An ablation study (App. E.1) further shows how increasing rule length or problem size, and switching from “uniform” to “LLM-guided” or “mirror” to “uniform” sampling, influences task complexity, confirming that SLR provides fine-grained control over logical complexity.

Test-time Scaling Improves Reasoning, but Returns Diminish and Costs Escalate. Reasoning LLMs clearly outperform the base models; not even the best base model is able to match any of the reasoning LLMs (cf. Tab. 3). CoT prompting provides modest gains ($\approx 5\%$, App. F.4). Further, scaling test-time computation comes at a steep cost as moving from GPT-4o to o3 doubles accuracy, but considerably increases the number of completion tokens (1777%) and thus the computational costs (1034%). Moreover, scaling test-time compute on the same model also boosts overall performance (see App. E.3), but does not guarantee higher accuracy across all tasks. E.g., while o4-mini-high typically outperforms o4-mini (LRL: 12.8 vs. 12.3), it

underperforms on the medium tier (40% vs. 52%). This plateau effect demonstrates that, beyond a certain threshold, additional compute may yield diminishing or even negative returns.

Scaling Model Parameters Brings Limited Gains in Logical Reasoning. Increasing model size yields only marginal gains in logical reasoning. While larger models like Llama-3-70B and GPT-4.5-prev generally outperform their smaller counterparts, returns diminish as improvements are increasingly modest (see Tab. 3). As even the largest base models still fall short of the reasoning models, a capability gap remains that suggests that scaling model parameters alone does not guarantee substantial advances in logical reasoning capabilities.

Common Error Modes. Base LLMs reliably produce syntactically valid rules, as reflected by their high syntax scores (Tab. 3). Reasoning models achieve slightly lower syntax scores, particularly on tasks of higher complexity. However, the dominant performance bottleneck is semantic rather than syntactic, as evidenced by the persistent gap between syntax and LRL. An error analysis (App. G) identifies five recurring failure modes: (i) over-generalized rules, (ii) under-generalized rules, (iii) failure of abstraction, and (iv) degenerate reasoning loops. Further, accuracy decreases with increasing output length (App. Tab. 14), suggesting that extended reasoning traces often lead to derailment instead of deeper reasoning.

6 Conclusion and Future Direction

In this work, we introduced SLR for scalable and fully automated synthesis for logical reasoning benchmarks with verifiable rewards provided via logic programs. Our instantiation, SLR-BENCH, offers a 20-level curriculum spanning 19k tasks with increasing logical complexity.

Our evaluations reveal that while current LLMs readily produce syntactically valid logic rules, robust logical reasoning remains elusive for conventional LLMs, especially as task complexity scales. Scaling parameters yield limited gains. Reasoning LLMs, aided by increased test-time compute, close part of this gap, albeit at significant computational costs. Notably, curriculum learning on SLR-BENCH substantially boosts in-domain and downstream reasoning. Specifically, our SLR-tuned Llama3-8B outperforms most conventional LLMs on SLR-BENCH and surpasses several SOTA reasoning LLMs at a fraction of their inference costs.

Looking ahead, SLR enables diverse extensions, including reinforcement learning, richer logical domains, benchmarking neuro-symbolic and interactive reasoning systems, and advancing toward higher-order tasks such as causal inference. While SLR-BENCH intentionally focuses on clean, single-rule ILP tasks with a perfect symbolic judge to isolate reasoning ability, future work can incorporate controlled forms of real-world noise, ambiguity, and multi-rule interactions to systematically study robustness alongside reasoning capabilities. Overall, SLR provides a flexible and extensible testbed for probing and improving LLM reasoning.

7 Limitations

While SLR and SLR-BENCH provide a scalable testbed for logical reasoning, there remain many opportunities for further enrichment. Although SLR-BENCH currently applies SLR to the train domain with a single rule, the framework is readily extensible to multiple more complex, multi-rule reasoning scenarios and to entirely different domains. Our current focus on first-order, function-free Horn clauses enables systematic benchmark creation and evaluation; future instantiations could expand towards higher-order logic or probabilistic reasoning. While synthetic task generation comes with many benefits, such as ensuring novelty and precise control, it makes it difficult to incorporate real-world diversity and ambiguity. Our symbolic judge provides deterministic, discrete scoring and could potentially be enhanced to also recognize partial solutions, syntactically invalid rules, or natural language formulations. Overall, these points highlight the flexibility of our framework and outline promising directions for broadening its reach and impact.

Acknowledgments

We acknowledge support of the hessian.AI Innovation Lab (funded by the Federal Ministry of Research, Technology and Space, BMFTR, grant no. 16IS22091), the hessian.AISC Service Center (funded by the Federal Ministry of Education and Research, BMBF, grant No 01IS22091), and the Center for European Research in Trusted AI (CERTAIN). Further, this work benefited from the ICT-48 Network of AI Research Excellence Center “TAILOR” (EU Horizon 2020, GA No 952215), the Hessian research priority program LOEWE within the project “WhiteBox”, the HMWK clus-

ter projects “Adaptive Mind” and “Third Wave of AI”, and from the NHR4CES. This work has also benefited from the BMW project “Sovereign Open Source Foundational Models für European Intelligence (SOOFI),” 13IPC040G, and also from early stages of the Cluster of Excellence “Reasonable AI” funded by the German Research Foundation (DFG) under Germany’s Excellence Strategy—EXC-3057; funding will begin in 2026. This work was supported by the Priority Program (SPP) 2422 in the subproject “Optimization of active surface design of high-speed progressive tools using machine and deep learning algorithms” funded by the German Research Foundation (DFG). Further, this work was funded by the European Union (Grant Agreement no. 101120763 - TANGO) as well as the AlephAlpha Collaboration lab 1141. This work was supported in part by OpenAI Research Credits.

Broader Impact

SLR and SLR-BENCH provide a scalable, reproducible foundation for evaluating and advancing logical reasoning in AI without relying on human annotation. By enabling robust measurement and targeted training, our framework supports progress in areas such as scientific discovery, program synthesis, and trustworthy AI. However, as LLMs acquire deeper logical competence, the risk of dual-use knowledge increases, enabling beneficial applications but also the potential for misuse, such as generating deceptive arguments or bypassing safety mechanisms. We urge responsible use and active consideration of ethical risks as these capabilities advance.

References

Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*.

Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML ’09*, page 41–48, New York, NY, USA. Association for Computing Machinery.

Victor Bursztyin, David Demeter, Doug Downey, and Larry Birnbaum. 2022. Learning to perform complex tasks through compositional fine-tuning of language

models. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 1676–1686, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *ArXiv*, abs/1803.05457.

Andrew Cropper and Sebastijan Dumančić. 2022. Inductive logic programming at 30: A new introduction. *J. Artif. Int. Res.*, 74.

Asok K. Debnath, Ricardo L. Lopez de Compadre, Goutam Debnath, Alan J. Shusterman, and Corwin Hansch. 1991. Structure–activity relationship of mutagenic aromatic and heteroaromatic nitro compounds: Correlation with molecular orbital energies and hydrophobicity. *Journal of Medicinal Chemistry*, 34(2):786–797.

Quentin Delfosse, Jannis Blüml, Fabian Tatai, Théo Vincent, Bjarne Gregori, Elisabeth Dillies, Jan Peters, Constantin Rothkopf, and Kristian Kersting. 2025. Deep reinforcement learning agents are not even close to human intelligence. *Preprint*, arXiv:2505.21731.

Long Phan et.al. 2025. Humanity’s last exam. *Preprint*, arXiv:2501.14249.

Siqi Fan, Peng Han, Shuo Shang, Yequan Wang, and Aixin Sun. 2025. Cothink: Token-efficient reasoning via instruct models guiding reasoning models. *Preprint*, arXiv:2505.22017.

Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, and 5 others. 2024a. The language model evaluation harness.

Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, and 5 others. 2024b. The language model evaluation harness.

Simeng Han, Hailey Schoelkopf, Yilun Zhao, Zhenting Qi, Martin Riddell, Luke Benson, Lucy Sun, Ekaterina Zubova, Yujie Qiao, Matthew Burtell, David Peng, Jonathan Fan, Yixin Liu, Brian Wong, Malcolm Sailor, Ansong Ni, Linyong Nan, Jungo Kasai, Tao Yu, and 7 others. 2022. Folio: Natural language reasoning with first-order logic. *arXiv preprint arXiv:2209.00840*.

- Lukas Helff, Wolfgang Stammer, Hikaru Shindo, Devendra Singh Dhami, and Kristian Kersting. 2025. [V-lol: A diagnostic dataset for visual logical learning](#). *Journal of Data-centric Machine Learning Research*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Mehran Kazemi, Bahare Fatemi, Hritik Bansal, John Palowitch, Chrysovalantis Anastasiou, Sanket Vaibhav Mehta, Lalit K. Jain, Virginia Aglietti, Disha Jindal, Peter Chen, Nishanth Dikkala, Gladys Tyen, Xin Liu, Uri Shalit, Silvia Chiappa, Kate Olszewska, Yi Tay, Vinh Q. Tran, Quoc V. Le, and Orhan Firat. 2025. [Big-bench extra hard](#). *Preprint*, arXiv:2502.19187.
- Jiin Kim, Byeongjun Shin, Jinha Chung, and Minsoo Rhu. 2025. [The cost of dynamic reasoning: Demystifying ai agents and test-time scaling from an ai infrastructure perspective](#). *Preprint*, arXiv:2506.04301.
- Aida Kostikova, Zhipin Wang, Deidamea Bajri, Ole Putz, Benjamin Paassen, and Steffen Eger. 2025. [Lllms: A data-driven survey of evolving research on limitations of large language models](#).
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Bill Yuchen Lin. 2024. [Zeroeval: A unified framework for evaluating language models](#).
- Bill Yuchen Lin, Ronan Le Bras, Kyle Richardson, Ashish Sabharwal, Radha Poovendran, Peter Clark, and Yejin Choi. 2025. [ZebraLogic: On the scaling limits of llms for logical reasoning](#).
- Bill Yuchen Lin, Yuntian Deng, Khyathi Chandu, Faeze Brahman, Abhilasha Ravichander, Valentina Pyatkin, Nouha Dziri, Ronan Le Bras, and Yejin Choi. 2024. [Wildbench: Benchmarking llms with challenging tasks from real users in the wild](#). *Preprint*, arXiv:2406.04770.
- Hanmeng Liu, Jian Liu, Leyang Cui, Zhiyang Teng, Nan Duan, Ming Zhou, and Yue Zhang. 2023. [Logiqa 2.0—an improved dataset for logical reasoning in natural language understanding](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:2947–2962.
- Jian Liu, Leyang Cui, Hanmeng Liu, Dandan Huang, Yile Wang, and Yue Zhang. 2020. [Logiqa: A challenge dataset for machine reading comprehension with logical reasoning](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 3622–3628. International Joint Conferences on Artificial Intelligence Organization.
- Junteng Liu, Yuanxiang Fan, Zhuo Jiang, Han Ding, Yongyi Hu, Chi Zhang, Yiqi Shi, Shitong Weng, Aili Chen, Shiqi Chen, Yunan Huang, Mozhi Zhang, Pengyu Zhao, Junjie Yan, and Junxian He. 2025. [Synlogic: Synthesizing verifiable reasoning data at scale for learning logical reasoning and beyond](#). *Preprint*, arXiv:2505.19641.
- John W Lloyd. 2012. *Foundations of logic programming*. Springer Berlin, Heidelberg.
- Man Luo, Shrinidhi Kumbhar, Ming shen, Mihir Parmar, Neeraj Varshney, Pratyay Banerjee, Somak Aditya, and Chitta Baral. 2024. [Towards logiglue: A brief survey and a benchmark for analyzing logical reasoning capabilities of language models](#). *Preprint*, arXiv:2310.00836.
- Kaijing Ma, Xinrun Du, Yunran Wang, Haoran Zhang, Zhoufutu Wen, Xingwei Qu, Jian Yang, Jiaheng Liu, Minghao Liu, Xiang Yue, Wenhao Huang, and Ge Zhang. 2024. [Kor-bench: Benchmarking language models on knowledge-orthogonal reasoning tasks](#). *Preprint*, arXiv:2410.06526.
- Ryszard S. Michalski. 1980. [Pattern recognition as rule-guided inductive inference](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-2(4):349–361.
- Tom Michael Mitchell. 1997. [Machine learning, international edition](#). In *McGraw-Hill Series in Computer Science*.
- Terufumi Morishita, Gaku Morio, Atsuki Yamaguchi, and Yasuhiro Sogawa. 2024. Enhancing reasoning capabilities of llms via principled synthetic logic corpus. In *Annual Conference on Neural Information Processing Systems*.
- Stephen Muggleton and Luc de Raedt. 1994. Inductive logic programming: Theory and methods. *The Journal of Logic Programming*, 19-20:629–679.
- OpenAI. 2025. [Openai o3 and o4-mini system card](#). <https://cdn.openai.com/pdf/2221c875-02dc-4789-800b-e7758f3722c1/o3-and-o4-mini-system-card.pdf>.
- Tafjord Oyvind, Dalvi Bhavana, and Clark Peter. 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. Association for Computational Linguistics.
- Bhrij Patel, Souradip Chakraborty, Wesley A. Suttle, Mengdi Wang, Amrit Singh Bedi, and Dinesh Manocha. 2024a. [Aime: Ai system optimization via multiple llm evaluators](#). *Preprint*, arXiv:2410.03131.

- Nisarg Patel, Mohith Kulkarni, Mihir Parmar, Aashna Budhiraja, Mutsumi Nakamura, Neeraj Varshney, and Chitta Baral. 2024b. [Multi-logieval: Towards evaluating multi-step logical reasoning ability of large language models](#). *Preprint*, arXiv:2406.17169.
- Shanghaoran Quan, Jiayi Yang, Bowen Yu, Bo Zheng, Dayiheng Liu, An Yang, Xuancheng Ren, Bofei Gao, Yibo Miao, Yunlong Feng, and 1 others. 2025. Codeelo: Benchmarking competition-level code generation of llms with human-comparable elo ratings. *arXiv preprint arXiv:2501.01257*.
- Luc De Raedt. 1997. [Logical settings for concept-learning](#). *Artif. Intell.*, 95:187–201.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. 2024. [GPQA: A graduate-level google-proof q&a benchmark](#). In *First Conference on Language Modeling*.
- 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*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y.K. Li, Y. Wu, and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#).
- Parshin Shojaei, 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](#).
- Koustuv Sinha, Shagun Sodhani, Jin Dong, Joelle Pineau, and William L. Hamilton. 2019. [CLUTRR: A diagnostic benchmark for inductive reasoning from text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4506–4515, Hong Kong, China. Association for Computational Linguistics.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V. Le, Ed H. Chi, Denny Zhou, and Jason Wei. 2023. [Challenging big-bench tasks and whether chain-of-thought can solve them](#). In *ACL (Findings)*, pages 13003–13051.
- DeepSeek-AI Team. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *Preprint*, arXiv:2501.12948.
- Meta Team. 2024a. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- OpenAI Team. 2024b. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Jason Weston, Antoine Bordes, Sumit Chopra, and Tomas Mikolov. 2015. [Towards ai-complete question answering: A set of prerequisite toy tasks](#). *arXiv: Artificial Intelligence*.
- Tim Woydt, Moritz Willig, Antonia Wüst, Lukas Helff, Wolfgang Stammer, Constantin A. Rothkopf, and Kristian Kersting. 2025. [Fodor and pylyshyn’s legacy – still no human-like systematic compositionality in neural networks](#).
- Antonia Wüst, Tim Tobiasch, Lukas Helff, Inga Ibs, Wolfgang Stammer, Devendra S Dhami, Constantin A Rothkopf, and Kristian Kersting. 2025. [Bongard in wonderland: Visual puzzles that still make ai go mad?](#) In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*.
- Chulin Xie, Yangsibo Huang, Chiyuan Zhang, Da Yu, Xinyun Chen, Bill Yuchen Lin, Bo Li, Badih Ghazi, and Ravi Kumar. 2024. [On memorization of large language models in logical reasoning](#). *Preprint*, arXiv:2410.23123.
- Tian Xie, Zitian Gao, Qingnan Ren, Haoming Luo, Yuqian Hong, Bryan Dai, Joey Zhou, Kai Qiu, Zhirong Wu, and Chong Luo. 2025. [Logic-rl: Unleashing llm reasoning with rule-based reinforcement learning](#). *Preprint*, arXiv:2502.14768.
- Nathan Young, Qiming Bao, Joshua Bensemann, and Michael Witbrock. 2022. [AbductionRules: Training transformers to explain unexpected inputs](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 218–227, Dublin, Ireland. Association for Computational Linguistics.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [Hellaswag: Can a machine really finish your sentence?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- 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.
- Yujun Zhou, Jiayi Ye, Zipeng Ling, Yufei Han, Yue Huang, Haomin Zhuang, Zhenwen Liang, Kehan Guo, Taicheng Guo, Xiangqi Wang, and Xiangliang Zhang. 2025. [Dissecting logical reasoning in llms: A fine-grained evaluation and supervision study](#). *Preprint*, arXiv:2506.04810.

A Extended Notation

We revisit essential definitions of first-order logic that we follow in this paper. An FOL *Language* \mathcal{L} is a tuple $(\mathcal{P}, \mathcal{A}, \mathcal{F}, \mathcal{V})$, where \mathcal{P} is a set of predicates, \mathcal{A} is a set of constants, \mathcal{F} is a set of function symbols (functors), and \mathcal{V} is a set of variables. A *term* is a constant, a variable, or a term that consists of a functor. A *ground term* is a term with no variables. We denote n -ary predicate p by p/n . An *atom* is a formula $p(t_1, \dots, t_n)$, where p is an n -ary predicate symbol and t_1, \dots, t_n are terms. A *ground atom* or simply a *fact* is an atom with no variables. A *literal* is an atom or its negation. A *positive literal* is just an atom. A *negative literal* is the negation of an atom. A *clause* is a finite disjunction (\vee) of literals. A *definite clause* is a clause with exactly one positive literal. If A, B_1, \dots, B_n are atoms, then $A \vee \neg B_1 \vee \dots \vee \neg B_n$ is a definite clause. We write definite clauses in the form of $A :- B_1, \dots, B_n$. Atom A is called the *head*, and set of negative atoms $\{B_1, \dots, B_n\}$ is called the *body*. We call definite clauses by *rules* for simplicity in this paper. An atom is an *atomic formula*. For formula F and G , $\neg F$, $F \wedge G$, and $F \vee G$ are also formulas. *Interpretation* of language \mathcal{L} is a tuple $(\mathcal{D}, \mathcal{I}_A, \mathcal{I}_F, \mathcal{I}_P)$, where \mathcal{D} is the domain, \mathcal{I}_A is the assignments of an element in \mathcal{D} for each constant $a \in \mathcal{A}$, \mathcal{I}_F is the assignments of a function from \mathcal{D}^n to \mathcal{D} for each n -ary function symbol $f \in \mathcal{F}$, and \mathcal{I}_P is the assignments of a function from \mathcal{D}^n to $\{\top, \perp\}$ for each n -ary predicate $p \in \mathcal{P}$. For language \mathcal{L} and formula X , an interpretation \mathcal{I} is a *model* if the truth value of X w.r.t \mathcal{I} is true. Formula X is a *logical consequence* or *logical entailment* of a set of formulas \mathcal{H} , denoted $\mathcal{H} \models X$, if \mathcal{I} is a model for \mathcal{H} implies that \mathcal{I} is a model for X for every interpretation \mathcal{I} of \mathcal{L} .

B SLR-BENCH Details

B.1 Task Specification

Each task in SLR-BENCH is precisely governed by a combination of language features and task configuration parameters, enabling fine-grained control over complexity and diversity. A task specification comprises two main components: (i) the logical language, which determines the set of predicates and argument types available, and (ii) the task configuration, which defines structural aspects of the task such as problem size, background knowledge sampling, rule length, sampling strategy, and the combinatorial space of realizable tasks. Tab. 4

details the curriculum’s level-wise specifications, showing how both language elements and the task config to create the individual levels.

B.2 Language

Predicates and Types. SLR defines a flexible, extensible vocabulary to support the systematic generation and evaluation of logical reasoning tasks. The primary predicate signatures and their argument types used in SLR-BENCH are: eastbound(TRAIN), westbound(TRAIN), has_car(TRAIN,CAR), car_num(CAR,NUM), car_color(CAR,COLOR), car_len(CAR,LEN), has_wall(CAR,WALL), has_roof(CAR,ROOF), has_payload(CAR,LOADS), load_num(CAR,NPAY), has_wheel(CAR,WHEELS), has_window(CAR,WINDOW), car_type(CAR,CTYPE),

Grounding Domains. Each argument type is grounded in a finite set of discrete constants:

```

NUM ::= [[0-9]]+
CAR ::= [[0-9]]+
COLOR ::= red | blue | green |
         yellow | white
LEN ::= short | long
WALL ::= full | railing
ROOF ::= roof_foundation |
        solid_roof | braced_roof |
        peaked_roof | none
WHEELS ::= 2 | 3
LOADS ::= blue_box | golden_vase |
         barrel | diamond | metal_pot |
         oval_vase | none
NPAY ::= 0 | 1 | 2 | 3
WINDOW ::= full | half | none
CTYPE ::= passenger | freight | mixed
NPAX ::= [[0-9]]

```

Grammar Constraints. Predicates are only instantiated with semantically compatible constant types. For example, `car_color(·,·)` only takes car objects and color constants as arguments; ill-typed facts are excluded during synthesis.

Task Synthesis Algorithm. Algorithm 1 details the synthesis pipeline, which first generates a gov-

Table 4: **SLR-BENCH Learning Curriculum**: The table details how both language and task configuration are systematically increased throughout the curriculum levels. Notably, higher levels involve richer problems with more constants and predicates, larger problems, longer rules, and a transition from mirror to uniform and LLM-guided sampling. The final column reports the approximate combinatorial size of unique tasks available at each level.

Stage	Basic					Easy					Medium					Hard				
Levels	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Language																				
#Const.	1	1	1	2	2	2	2	2-3	2-3	2-3	2-4	2-4	4-6	4-6	4-6	5-6	5-6	5-6	5-6	5-6
#Pred.	5	5	5	5	5	5	6	6	6	7	7	9	9	9	9	10	10	12	12	12
Task Config.																				
κ	2	2	4	4	6	6	6	8	10	12	14	16	18	20	22	24	26	28	30	32
B_π	\mathcal{M}	\mathcal{M}	\mathcal{M}	\mathcal{M}	\mathcal{M}	\mathcal{U}	\mathcal{U}	\mathcal{U}	\mathcal{U}	\mathcal{U}	\mathcal{U}	\mathcal{U}	\mathcal{U}	\mathcal{U}	\mathcal{U}	\mathcal{U}	\mathcal{U}	\mathcal{U}	\mathcal{U}	\mathcal{U}
R_{len}	1	1-2	1-2	1-2	1-2	1-2	1-2	1-2	2-3	2-3	2-3	3-4	3-4	4-5	4-5	4-5	4-5	4-5	4-5	5
R_{sample}	\mathcal{U}	\mathcal{U}	\mathcal{U}	\mathcal{U}	\mathcal{U}	\mathcal{U}/L	\mathcal{U}/L	\mathcal{U}/L	\mathcal{U}/L	\mathcal{U}/L	\mathcal{U}/L	\mathcal{U}/L	\mathcal{U}/L	\mathcal{U}/L	\mathcal{U}/L	\mathcal{U}/L	\mathcal{U}/L	\mathcal{U}/L	\mathcal{U}/L	\mathcal{U}/L
Comb. Size	10^3	10^3	10^5	10^{10}	10^{16}	10^{16}	10^{24}	10^{32}	10^{40}	10^{55}	10^{65}	10^{120}	10^{271}	10^{300}	10^{330}	10^{507}	10^{549}	10^{805}	10^{861}	10^{919}

\mathcal{M} : mirror sampling \mathcal{U} : uniform sampling L : LLM-guided generation

erning rule before iteratively sampling background contexts and queries. These candidates undergo stratified rejection sampling to curate balanced positive and negative sets of samples that compose the final task prompt and validation program.

B.3 Learning Curriculum in Detail

SLR-BENCH presents a 20-level curriculum that progresses systematically in difficulty and is organized into four overarching complexity tiers: *basic*, *easy*, *medium*, and *hard*. For a detailed overview of the level-wise task specifications, see App. Tab. 4. Each level increases in difficulty through systematic modifications to both the task language and configuration, including: (i) increasing the overall size of the task (κ), (ii) expanding the vocabulary by adding new car constants per train, and (iii) broadening the set of predicates and grounding domains. Additionally, (iv) we adapt the policy used for synthesizing background knowledge, (v) alter the length R_{len} of the ground truth rule, and (vi) adjust the rule sampling policy. Finally, (vii) the combinatorial space of tasks, quantified as the approximate number of distinct tasks that can be generated per level, also increases across the curriculum. The compounding effect of larger problems, expanded vocabulary (constants and predicates), and more complex rules not only yields exponential growth in the combinatorial space but also effectively increases the complexity of the tasks. Consequently, while models identifying surface-level patterns might solve early levels, reasoning in later levels becomes more sophisticated, demanding more advanced problem-solving capabilities to progress.

B.4 Verifiable Logic Rewards

The SYMBOLICJUDGE computes verifiable logic rewards for a candidate hypothesis H with respect to a background knowledge base B and sets of positive (E^+) and negative (E^-) examples. These rewards are used both for evaluation and for training, and they follow standard practice in Inductive Logic Programming (ILP). Concretely, the judge deterministically tests whether a hypothesis satisfies classical ILP coverage criteria via Prolog execution, which provides the operational semantics for entailment and refutation. A key advantage of this setup is that it naturally supports *non-unique solutions*: although tasks in SLR-BENCH are generated from a single ground-truth rule, the learned hypothesis is not required to match it syntactically. Multiple distinct hypotheses may be semantically equivalent in that they entail all positives and reject all negatives under B , and any such hypothesis is considered correct. This allows evaluation to focus purely on semantic correctness rather than surface form. We compute three distinct metrics:

Syntax Validity Score: This binary score (0 or 1) indicates whether the candidate hypothesis H is a syntactically and semantically valid Prolog rule. This serves as a prerequisite check for further evaluation; if H is invalid, other scores are typically assigned 0.

OVERALLSCORE: This metric provides a binary indication of perfect task completion. It is 1 if and only if the hypothesis H , in conjunction with the background knowledge B , correctly entails all positive examples (E^+) and correctly refutes all

negative examples (E^-). Otherwise, the score is 0.

$$\text{OVERALLSCORE}_{B,E^+,E^-}(H) = \llbracket \forall q \in E^+ : (B \cup H) \models q \wedge \forall q \in E^- : (B \cup H) \not\models q \rrbracket \in \{0,1\}$$

Where $\llbracket \cdot \rrbracket$ is the Iverson bracket, evaluating to 1 if the condition inside is true, and 0 otherwise.

PARTIALSCORE: This metric reflects the fraction of examples (from both E^+ and E^-) that are correctly classified by the candidate hypothesis H when combined with the background knowledge B . This provides a continuous signal of progress, even when the overall task is not perfectly completed.

$$\text{PARTIALSCORE}_{B,E^+,E^-}(H) = \frac{\sum_{q \in E^+} \llbracket (B \cup H) \models q \rrbracket + \sum_{q \in E^-} \llbracket (B \cup H) \not\models q \rrbracket}{|E^+ \cup E^-|}$$

Here, the numerator sums the count of correctly entailed positive examples and correctly not entailed negative examples. The denominator is the total number of examples, ensuring the score is normalized between 0 and 1.

These metrics provide rich feedback for both discrete evaluation (e.g., for filtering valid rules) and continuous optimization (e.g., for guiding reinforcement learning agents), allowing for robust assessment of learned logical hypotheses.

B.5 LLM-Guided Rule Generation

This section provides the prompt used for LLM-guided rule generation. The prompt was carefully designed to be both diverse and comprehensive, including a wide range of logical structures and Prolog features such as conjunction, disjunction, negation, recursion, aggregation, and pattern matching. By presenting the model with these varied and complex examples, we encourage the generation of challenging and realistic logic rules. This diversity is crucial for robust model generation of new rules, as it ensures that the LLM is exposed to representative samples of possible rule types encountered in real-world logic programming tasks.

1. Conjunction with Existential Quantification: There exists a red short car There is at least one car that is both short and red.

```
1 eastbound(Train) :-
2   has_car(Train, Car),
3   car_color(Car, red),
4   car_len(Car, short).
```

2. Disjunction: Some car is white or yellow. At least one car is either white or yellow.

```
1 eastbound(Train) :-
2   has_car(Train, Car),
3   (car_color(Car, white) ; car_color(Car, yellow)).
```

3. Negation: The train does not contain any red cars No car on the train is red.

```
1 eastbound(Train) :-
2   \+ (has_car(Train, Car), car_color(Car, red)).
```

4. Inequality/Distinctness: Two cars must have different colors There are at least two cars on the train with different colors.

```
1 eastbound(Train) :-
2   has_car(Train, CarA),
3   has_car(Train, CarB),
4   CarA \= CarB,
5   car_color(CarA, Color1),
6   car_color(CarB, Color2),
7   Color1 \= Color2.
```

5. Aggregation/Counting: There are more green cars than yellow cars The train contains more green cars than yellow cars.

```
1 eastbound(Train) :-
2   findall(Car, (has_car(Train, Car), car_color(Car, green)), Greens),
3   findall(Car, (has_car(Train, Car), car_color(Car, yellow)), Yellows),
4   length(Greens, G),
5   length(Yellows, Y),
6   G > Y.
```

6. Mutual Exclusion: Only one car is yellow; all others are not yellow There is exactly one yellow car; all others are not yellow.

```
1 eastbound(Train) :-
2   findall(Car, (has_car(Train, Car), car_color(Car, yellow)), [YellowCar]),
3   forall(
4     (has_car(Train, Car), Car \= YellowCar),
5     (car_color(Car, NotYellow), NotYellow \= yellow)
6   ).
```

7. Uniqueness: No two cars have the same color All cars have unique colors.

```
1 eastbound(Train) :-
2   findall(Color, (has_car(Train, Car), car_color(Car, Color)), Colors),
3   sort(Colors, UniqueColors),
4   length(Colors, N),
5   length(UniqueColors, N).
```

8. No-Other/Uniqueness: Only two cars in the train Only two cars are present in the train.

```
1 eastbound(Train) :-
2   findall(Car, has_car(Train, Car), Cars),
3   length(Cars, 2).
```

9. Universal Quantification: Every full-wall car is long All cars with a full wall must be long.

```
1 eastbound(Train) :-
2   forall(
3     (has_car(Train, Car), has_wall(
4       Car, full)),
5     car_len(Car, long)
6   ).
```

10. Conditional Implication: All long cars are either red or blue Every long car is either red or blue.

```
1 eastbound(Train) :-
2   forall(
3     (has_car(Train, Car), car_len(
4       Car, long)),
5     (car_color(Car, Color), (Color =
6       red ; Color = blue))
7   ).
```

11. Conditional Aggregation: All long cars are either red or blue Every long car is either red or blue.

```
1 eastbound(Train) :-
2   forall(
3     (has_car(Train, Car), car_len(
4       Car, long)),
5     (car_color(Car, Color), (Color =
6       red ; Color = blue))
7   ).
```

12. Pattern Matching: All full-wall cars are white Every full-wall car is white.

```
1 eastbound(Train) :-
2   forall(
3     (has_car(Train, Car), has_wall(Car
4       , full)),
5     car_color(Car, white)
6   ).
```

13. Symmetry: Two cars are neighbors with same color CarA and CarB are neighbors on the train and have the same color.

```
1 eastbound(Train) :-
2   has_car(Train, CarA),
3   has_car(Train, CarB),
4   CarA \= CarB,
5   car_num(CarA, N1),
6   car_num(CarB, N2),
7   (N2 == N1 + 1 ; N2 == N1 - 1),
8   car_color(CarA, Color),
9   car_color(CarB, Color).
```

14. Combinatorial Group: Exactly two short yellow cars There are exactly two yellow cars, and both are short.

```
1 eastbound(Train) :-
2   findall(Car, (has_car(Train, Car),
3     car_color(Car, yellow), car_len(Car,
4     short)), L),
5   length(L, 2).
```

15. Recursion: At least one long car in the train The train has at least one long car.

```
1 eastbound([Car|Cars]) :-
2   car_len(Car, long)
3   ;
4   eastbound(Cars).
```

16. Existence of a Structure (Sublist Pattern Matching) Exists three cars in sequence: Num, Num+1, Num+2, matching pattern.

```
1 eastbound(Train) :-
2   has_car(Train, Car1), car_num(Car1,
3     N),
4   car_len(Car1, short),
5   N2 is N+1, N3 is N+2,
6   has_car(Train, Car2), car_num(Car2,
7     N2), car_len(Car2, long),
8   has_car(Train, Car3), car_num(Car3,
9     N3), car_len(Car3, short).
```

17. Min/Max and Extremal Values A short car followed by a long car followed by a short car, anywhere in the train.

```
1 eastbound(Train) :-
2   findall(N, (has_car(Train, Car),
3     car_num(Car, N)), Numbers),
4   max_list(Numbers, Max),
5   has_car(Train, LastCar),
6   car_num(LastCar, Max),
7   car_color(LastCar, white).
```

18. Subset/Superset Constraints All full-wall cars are among the first three cars.

```
1 eastbound(Train) :-
2   forall(
3     (has_car(Train, Car), has_wall(Car
4       , full)),
5     (car_num(Car, N), N =< 3)
6   ).
```

19. Projection/Aggregation Over Multiple Properties All pairs of cars have different (color, length) tuples.

```
1 eastbound(Train) :-
2   has_car(Train, CarA), has_car(Train,
3     CarB), CarA \= CarB,
4   car_color(CarA, ColA), car_len(CarA,
5     LenA),
6   car_color(CarB, ColB), car_len(CarB,
7     LenB),
8   (ColA \= ColB ; LenA \= LenB).
```

20. All-Different on Multiple Attributes Enforce all car colors are different, AND all car numbers are different (car numbers are unique by assumption, but see structure).

```
1 eastbound(Train) :-
2   findall(Color, (has_car(Train, Car),
3     car_color(Car, Color)), Colors),
4   sort(Colors, UniqueColors),
5   length(UniqueColors, N).
```

B.6 Data Leakage Guarantees.

We verify that no ground-truth rules (R^*) from the training split reappear in the test split, next to the uniqueness of the background. While ensuring complete semantic independence is inherently challenging, surface-level similarity measures (e.g., predicate co-occurrence or template hashing) are insufficient to capture true logical overlap. Semantically distinct rules may share identical predicates and constants (e.g., “there exists a red car” vs. “all cars are red”).

B.7 SLR-BENCH Task Synthesis

App. Fig. 6 illustrates a full synthesis run of SLR-BENCH according to Alg. 1. Starting from a specified language and configuration (rule length, problem size), the synthesizer first generates a latent ground-truth rule, then iteratively samples background facts until the desired number of positive and negative examples is met. The resulting task instance $\mathcal{I} = (B, E^+, E^-)$ forms a self-contained reasoning problem, which can be expressed either symbolically (Prolog format) or in natural language. This process demonstrates how SLR produces logically consistent, verifiable tasks with controllable difficulty and interpretable representations.

C Compute Costs and Model Pricing

Compute Costs. Compute costs are reported as the total USD cost to run all prompts, based on publicly listed API prices as of 01.05.2025, with Qwen3 and GPT5 model pricing added on 01.12.2025. Pricing ignores server-side token caching, as actual cache hit counts are unavailable. Table 5 summarizes per-model cost rates and API sources.

D Training and Evaluation Details

D.1 Training Setup

For our experiments on SLR-BENCH, we fine-tune Llama-3.1-8B-Instruct using supervised fine-tuning (SFT) with LoRA adapters using LLaMA-Factory (Zheng et al., 2024). Training is performed stage-wise, presenting basic, easy, and medium problems sequentially, reflecting a curriculum of increasing logical complexity. Training is distributed across 8 GPUs using DeepSpeed with ZeRO Stage 3 optimization, taking 4 hours. Both optimizer states and model parameters are offloaded to CPU with pinned memory to maximize GPU memory

efficiency. The AdamW optimizer is used in conjunction with a Warmup Cosine learning rate scheduler. Mixed-precision training is employed, with both bfloat16 and fp16 enabled in automatic mode. Communication overlap and contiguous gradients are activated to improve throughput, and model weights are saved in 16-bit precision at each checkpoint. Due to memory limitations, input sequences are truncated to a maximum length of 6k tokens using the Llama3 template, restricting training to `slr_basic_train`, `slr_easy_train`, and `slr_medium_train` splits. Optimization is performed using cross-entropy loss over the ground truth rule R^* , with a per-device batch size of 5 and gradient accumulation over 2 steps, resulting in an effective batch size of 80 samples per step across 8 GPUs. The learning rate is set to 2×10^{-4} , scheduled with a cosine scheduler and a warmup ratio of 0.03. All relevant hyperparameters and training scripts are included in the codebase for full reproducibility.

D.2 Evaluation Setup

To evaluate the generality of reasoning skills induced by SLR, we benchmark models on a diverse set of widely used reasoning datasets. These benchmarks span multiple reasoning modalities, including multi-domain analytical reasoning, relational and multi-hop inference, formal deductive reasoning, symbolic abstraction, scientific reasoning, and grounded commonsense inference. Together, they provide broad coverage of the reasoning capabilities relevant for modern large language models. An overview of all downstream datasets, including their scope, reasoning focus, and references, is provided in App. Tab. 6. MMLU and its subsets assess analytical reasoning across factual knowledge, mathematics, statistics, and computer science (Hendrycks et al., 2021). CLUTRR evaluates relational and multi-hop inference over structured family graphs (Sinha et al., 2019). LogiQA and LogiQA2 focus on formal deductive reasoning and logical consistency in natural language arguments (Liu et al., 2020, 2023). GPQA and its variants (Extended and Diamond) measure graduate-level scientific reasoning (Rein et al., 2024). ARC-Easy and ARC-Challenge test symbolic abstraction and pattern-based reasoning (Clark et al., 2018), while HellaSwag evaluates grounded commonsense inference and contextual continuity (Zellers et al., 2019). For downstream evaluation, we use the Language Model Evaluation Harness (Gao

Model	Model Tag	Input	Input (Cached)	Output	API
OpenAI Models					
gpt-5	gpt-5	1.25	1.25	10.00	OpenAI
gpt-5-mini	gpt-5-mini	0.25	0.25	2.00	OpenAI
gpt-5-nano	gpt-5-nano	0.05	0.005	0.40	OpenAI
gpt-4.5-preview	gpt-4.5-preview-2025-02-27	75.00	37.50	150.00	OpenAI
gpt-4o	gpt-4o-2024-08-06	2.50	1.25	10.00	OpenAI
o1	o1-2024-12-17	15.00	7.50	60.00	OpenAI
o3	o3-2025-04-16	10.00	2.50	40.00	OpenAI
o4-mini	o4-mini-2025-04-16	1.10	0.275	4.40	OpenAI
o3-mini	o3-mini-2025-01-31	1.10	0.55	4.40	OpenAI
o1-mini	o1-mini-2024-09-12	1.10	0.55	4.40	OpenAI
Open-Source Models					
Qwen3-0.6B	Qwen3-0.6B	0.015		0.025	OpenRouter
Qwen3-1.7B	Qwen3-1.7B	0.015		0.025	OpenRouter
Qwen3-4B	Qwen3-4B	0.015		0.025	OpenRouter
Qwen3-8B	Qwen3-8B	0.035		0.138	OpenRouter
Qwen3-14B	Qwen3-14B	0.05		0.22	OpenRouter
Qwen3-32B	Qwen3-32B	0.05		0.20	OpenRouter
Qwen3-235B	Qwen3-235B	0.18		0.54	OpenRouter
Qwen3-30B-A3B	Qwen3-30B-A3B	0.06		0.22	OpenRouter
Qwen3-Coder-30B-A3B	Qwen3-Coder-30B-A3B	0.06		0.25	OpenRouter
Qwen3-Coder-408B-A35B	Qwen3-Coder-408B-A35B	0.38		1.53	OpenRouter
Qwen3-Next-80B-A3B	Qwen3-Next-80B-A3B-Instruct	0.10		0.80	OpenRouter
	Qwen3-Next-80B-A3B-Thinking	0.15		1.20	OpenRouter
Open-Source VLMs					
Qwen3-VL-2B	Qwen3-VL-2B	0.00		0.00	OpenRouter
Qwen3-VL-4B	Qwen3-VL-4B	0.00		0.00	OpenRouter
Qwen3-VL-8B	Qwen3-VL-8B	0.008		0.50	OpenRouter
Qwen3-VL-30B	Qwen3-VL-30B	0.15		0.60	OpenRouter
Qwen3-VL-32B	Qwen3-VL-32B	0.35		1.10	OpenRouter
Qwen3-VL-235B	Qwen3-VL-235B	0.22		0.88	OpenRouter

Table 5: Model Pricing (\$ per 1M tokens). API rates as of 01.05.2025, with Qwen3 and GPT5 model pricing added on 01.12.2025

et al., 2024b) with the default protocol of each benchmark, enabling few-shot as multiturn prompting to support multi-turn contexts where applicable, including the official pass rate (typically pass@1) and whether evaluation is performed in zero-shot or few-shot mode. CLUTRR, which is not supported by evalharness, follows a similar setup with accuracy averaged across all test sets. All evaluations are conducted on 8 GPUs using vLLM (Kwon et al., 2023) for efficient batch inference. Reported scores reflect accuracy for each model and benchmark, and all results are based on the official evaluation splits and standardized prompt formatting consistent with the SLR curriculum.

E Ablation Studies

E.1 Ablation across synthesis parameters

To assess how synthesis configurations affect task difficulty, we conducted controlled ablations across the rule sampling policy, problem size, rule length, and background sampling policy, averaging results across all open-source LLaMA base models. As summarized in App. Tab. 7, these experiments confirm that SLR yields interpretable and controllable gradients in logical complexity. For each factor, we selected representative levels of the dataset where the targeted parameter changes.

(1) **Rule sampling policy.** To measure the influence of rule generation, we compare subsets within

Dataset	Scope & Task	Reasoning Specialty	Reference
MMLU	Multi-domain analytical reasoning	Factual knowledge, mathematics, statistics, computer science	(Hendrycks et al., 2021)
CLUTRR	Relational and multi-hop inference	Structured family graphs	(Sinha et al., 2019)
LogiQA (2.0)	Formal deductive reasoning	Argument consistency	(Liu et al., 2020, 2023)
GPQA	High-level scientific reasoning	Graduate-level difficulty	(Rein et al., 2024)
ARC	Symbolic abstraction	Pattern-based problem solving	(Clark et al., 2018)
HellaSwag	Grounded commonsense inference	Contextual continuity	(Zellers et al., 2019)

Table 6: **Downstream reasoning datasets.** Overview of downstream evaluation datasets and what they evaluate for.

Ablation Factor	Configuration	Accuracy (%)
Rule Sampling	Uniform	11.0
	LLM-guided	2.5 (\downarrow 8.6)
Problem Size (κ)	$\kappa = 4$	35.0
	$\kappa = 6$	30.0 (\downarrow 5.0)
(R_{len})	1–2 literals	7.2
	2–3 literals	2.6 (\downarrow 4.6)
Background Sampling	Mirror	30.0
	Uniform	20.5 (\downarrow 9.5)

Table 7: **Ablations Across Synthesis Parameters.** Accuracy (%) averaged across the LLaMA base models for different synthesis parameters. Each row corresponds to one ablation factor, comparing paired configurations. Arrows in red indicate decreases relative to the first configuration in each pair.

the easy-difficulty tier (levels 6–10) containing either uniformly sampled versus LLM-guided rules. As shown in Table 7, accuracy drops sharply from 11.0 % \rightarrow 2.5 %, demonstrating that LLM-guided synthesis introduces more structured and semantically complex rules that substantially increase reasoning difficulty.

(2) Problem size (κ). To isolate the effect of larger reasoning tasks, we compare level 4 to level 5, where κ increases from 4 to 6. This change leads to a moderate decrease in accuracy (35.0 % \rightarrow 30.0 %), showing that models begin to struggle as the relational scope and combinatorial complexity expand.

(3) Rule length (R_{len}). To examine the effect of rule depth, we compare level 8 (rules with one–two literals) and level 9 (two–three literals). Accuracy drops markedly from 7.2 % \rightarrow 2.6 %, confirming that even modestly longer reasoning chains substantially raise logical complexity and error rates.

(4) Background sampling policy. Finally, we compare level 5, which employs mirror sampling, to level 6, which uses uniform sampling over the background facts. Accuracy decreases from 30.0

Curriculum	LRL	Basic	Easy	Medium	Hard
	(\uparrow 0-20)	(\uparrow %)	(\uparrow %)	(\uparrow %)	(\uparrow %)
Ordered	6.01	88.7	28.4	3.3	0.2
Random	5.99	88.1	28.5	3.6	0.05
Reverse	5.93	86.8	27.7	4.5	0.05

Table 8: **Curriculum Learning Order.** Comparison of curriculum learning orders (ordered, random, and reverse) under matched training budgets (levels 5–7) using llama3.1-8b-it. The ordered curriculum achieves the highest overall LRL.

% \rightarrow 20.5 %, indicating that the mirror trains are significantly easier to solve than the uniformly sampled trains.

Together, these ablations highlight that the synthesis parameters in SLR modulate task difficulty, confirming that SLR-BENCH provides a precise and controllable testbed for analyzing logical reasoning performance across model families.

E.2 Curriculum Learning Order

To evaluate the effect of curriculum ordering on reasoning acquisition, we conducted a controlled ablation comparing standard curriculum learning, random order, and reverse curriculum training. All settings were matched for dataset, sample size, and total token budget, using the same difficulty levels (5–7). As shown in Tab. 8, the curriculum-trained model achieves the highest overall Logical Reasoning Level (LRL) and more balanced performance across tiers, while random and reverse orders yield slightly lower LRL and less stable behavior across difficulty levels. These results indicate that a structured progression of reasoning difficulty facilitates smoother learning and generalization in SLR.

E.3 Scaling Test-time compute with *o4-mini*

Table 9 analyzes the impact of increasing test-time compute within the *o4-mini* family. The *high* variant achieves the highest overall LRL, outperforming both *o4-mini* and *o4-mini-low*, but incurs a cost

Model Variant	LRL	Basic	Easy	Medium	Hard	Cost (\$)
<i>o4-mini-low</i>	10.3	91	81	25	9	7.26
<i>o4-mini</i>	12.3	93	88	52	13	21.43
<i>o4-mini-high</i>	12.8	98	96	40	21	24.24

Table 9: **Scaling Test-time Compute within the *o4-mini* Family.** Increasing compute improves reasoning ability but raises inference costs and does not yield uniform accuracy gains across all tiers.

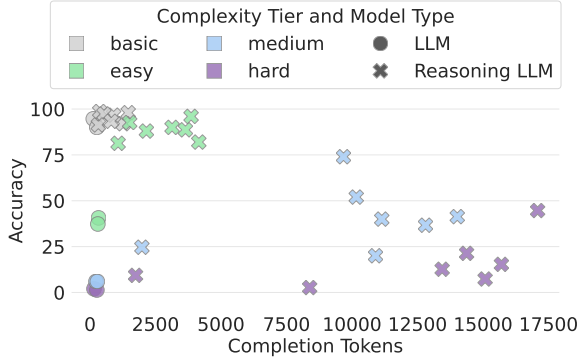


Figure 5: **Compute-Performance Trade-Off.** Reasoning LLMs achieve higher accuracy than base LLMs but require more compute. While more complex tasks typically demand more completion tokens, increased compute does not always translate to higher accuracy.

increase of over 13%. Notably, it performs worse on the *medium* tier (40% vs. 52%), indicating diminishing returns from additional compute. This supports the main text finding that higher inference cost does not guarantee proportional reasoning gains.

E.4 Compute Increases with Task Complexity

As reasoning tasks grow in complexity, models generally require more completion tokens to solve them, leading to higher inference costs (see Fig. 5). This trend serves as a sanity check confirming that SLR-BENCH’s task difficulty correlates with increased reasoning effort rather than being arbitrary. However, since larger models inherently consume more compute per token, these results should not be interpreted as direct efficiency comparisons across models.

F Additional Experiments

F.1 Extended Leaderboard

We extend our evaluation to the Qwen3 and GPT-5 model families and consolidate all new model results into a single unified extended leaderboard (Tab. 15). The leaderboard compares reasoning-oriented models (orange) with non-reasoning mod-

els (blue) across model families and scales, providing a comprehensive overview of current reasoning performance on SLR-Bench. The more recent models show improved LRL. Despite these advances, GPT-5 solves only 46% of the tasks in the Hard tier, highlighting that SLR-Bench remains challenging even for the strongest contemporary models, particularly at higher levels of inductive complexity.

F.2 Qwen3-SLR

We fine-tune a LoRA adapter on Qwen3-8B using a similar setup as the llama model, but with SLR training data shuffled. Subsequently, we evaluate the resulting model on SLR-Bench and on downstream reasoning benchmarks. As shown in App. Tab. 10, Qwen-SLR consistently improves performance on SLR tasks and also yields gains across downstream reasoning benchmarks, indicating that the learned reasoning capabilities transfer beyond the training distribution.

F.3 DeepSeek-SLR

To examine whether the SLR curriculum remains effective in RL setups, we applied it using Group Reinforcement Preference Optimization (GRPO) (Shao et al., 2024) on the small-scale reasoning model *deepseek-ai/DeepSeek-R1-Distill-Qwen-1.5B*. All training was performed under the same setup as SFT. The resulting model exhibits a clear improvement in reasoning ability, with the Logical Reasoning Level (LRL) increasing from 0.0 to 5.6. Performance gains are observed across reasoning tiers, especially on basic and easy reasoning tasks, demonstrating that a structured progression of task difficulty benefits even small models trained via reinforcement optimization. However, training was halted at level 7 due to reward collapse in later stages, suggesting that smaller architectures experience instability once the reasoning task complexity surpasses their capacity. Overall, these results indicate that the SLR curriculum not only supports SFT but also effectively enhances reasoning in RL setups (GRPO), reinforcing its flexibility as a general training paradigm.

F.4 Extended Pass@k and CoT Results

To complement the main results, where we report average Pass@1 (LRL) performance, we provide extended Pass@k scores (*averaged for $k < 8$*) and Chain-of-Thought (CoT) results for open-source models in App. Tab. 12. The table summarizes model performance at different test-time sampling

	LRL (\uparrow 0-20)	Syntax(\uparrow %)	Basic(\uparrow %)	Easy(\uparrow %)	Medium(\uparrow %)	Hard(\uparrow %)
Qwen3-8B	4.6	93	79	14	0	0
Qwen3-8B-SLR	7.9 (+3.3)	100 (+7)	96 (+17)	43 (+29)	14 (+14)	6 (+6)
<i>Downstream Reasoning Performance</i>						
	SLR-Homes (\uparrow %)	MMLU (\uparrow %)	MMLU-Stats (\uparrow %)	CLUTRR (\uparrow %)	LogiQA (\uparrow %)	LogiQA2 (\uparrow %)
Qwen3-8B	38.0	49.8	18.1	19.6	25.0	27.7
Qwen3-8B-SLR	40.7 (+2.7)	55.9 (+6.1)	18.1 (+0.0)	20.6 (+1.0)	31.5 (+6.5)	31.0 (+3.3)
	GPQA (\uparrow %)	GPQA-Ext. (\uparrow %)	GPQA-Dia. (\uparrow %)	ARC-Easy (\uparrow %)	ARC (\uparrow %)	HellaSwag (\uparrow %)
Qwen3-8B	23.4	26.4	21.2	69.2	43.3	48.8
Qwen3-8B-SLR	26.8 (+3.4)	27.1 (+0.7)	27.8 (+6.5)	77.3 (+8.1)	49.8 (+6.5)	51.6 (+2.8)

Table 10: **Qwen3-8B-SLR**. Benchmark scores (\uparrow) for base and SLR-tuned models on SLR-BENCH and downstream benchmarks; LRL measuring cumulative curriculum progress. The tuned model surpasses the baseline across all curriculum stages, while generalizing to other downstream reasoning tasks.

	LRL (\uparrow 0-20)	Syntax(\uparrow %)	Basic(\uparrow %)	Easy(\uparrow %)	Medium(\uparrow %)	Hard(\uparrow %)
DeepSeek-R1-1.5B	0	54	0	0	0	0
DeepSeek-1.5B-SLR	5.6 (+5.6)	97(+43)	84 (+84)	27 (+27)	2 (+2)	0 (+0)

Table 11: **DeepSeek-SLR**. GRPO training of DeepSeek-R1-1.5B on SLR-BENCH. Benchmark scores (\uparrow) for base and SLR-tuned models (DeepSeek-R1-1.5B) on SLR-BENCH; LRL measures cumulative curriculum progress. The SLR-tuned model improves syntax and reasoning capabilities on basic, easy, and medium problems.

levels ($k = 1, 4, 8$), together with compute requirements in tokens and cost for a single pass. Across all models, accuracy increases with larger k , confirming that multiple sampled completions raise the probability of generating a correct reasoning trace. Beyond improving performance, Pass@ k also reveals reasoning variance. Smaller models (e.g., 3B–8B) show larger Pass@1–Pass@8 gaps, doubling in accuracy as we move from pass@1 to pass@8, indicating less stable reasoning, while larger models display narrower gaps and more consistent inference. CoT variants generally show modest gains over their non-CoT counterparts of comparable scale, indicating that explicit reasoning traces can improve reasoning, but also come with higher computational costs. Overall, these findings show that both test-time sampling and CoT prompting not only enhance reasoning outcomes but also serve as useful probes into a model’s internal reasoning diversity and reliability.

F.5 CLUTRR Evaluation

We evaluated SLR-tuned and baseline models on the CLUTRR benchmark (Sinha et al., 2019), an inductive reasoning dataset requiring inference over kinship relations. All models were evaluated with greedy decoding, focusing on accuracy since CLUTRR uses a closed classification space where syntax validity does not apply. Across all vari-

ants, SLR-tuning consistently improves over the base model, with an average gain of +6.6pp, and remains competitive with or superior to DeepSeek-R1-8B on most configurations (App. Tab. 13). Notably, Llama-3.1-8B-SLR achieves higher accuracy than DeepSeek-R1-8B on four of six variants while using dramatically fewer resources—only 11k completion tokens vs. 8.1M for DeepSeek (730× less compute). These results provide strong evidence that SLR enhances inductive reasoning beyond SLR-Bench, achieving higher accuracy and far greater computational efficiency on an established external inductive reasoning benchmark.

G Error Mode Analysis.

To better understand model failures on SLR-BENCH, we conducted a comprehensive error analysis of generated rules. As discussed in Section 4, syntax is rarely the limiting factor; most models achieve near-perfect syntactic validity. Consequently, the primary bottleneck lies in semantic rule induction.

G.1 Taxonomy of Error Modes

By inspecting a sample of failed generations across all models, we identified five prominent failure patterns:

1. **Over-generalization:** Rules that are too broad, often consisting of only one or

Model	LRL@1 (\uparrow 0-20)	LRL@4 (\uparrow 0-20)	LRL@8 (\uparrow 0-20)	Tokens ($M \downarrow$)	Cost (\$ \downarrow)
llama-3.1-8b-SLR	8.2	10.5	11.3	0.05	0.14
Llama-3.3-70B-IT-CoT	5.9	7.0	7.5	0.93	0.93
Llama-3.3-70B-IT	5.6	6.8	7.3	0.49	0.82
Llama-3.1-8B	3.8	5.5	6.0	0.20	0.14
Llama-3.1-8B-IT-CoT	3.4	5.3	6.0	1.37	0.18
Llama-3.2-3B-IT-CoT	1.8	3.8	4.5	1.02	0.13
Llama-3.2-3B-IT	1.0	2.6	3.6	0.10	0.11
Llama-3.2-1B-IT-CoT	0.0	0.0	0.0	1.17	0.17
Llama-3.2-1B-IT	0.0	0.0	0.0	0.47	0.12

Table 12: **SLR-BENCH Pass@k and CoT Results.** We report Logical Reasoning Level (LRL) scores at Pass@1, Pass@4, and Pass@8 (%) for open-source models, together with total completion tokens (M) and inference cost (\$). Higher LRL values indicate stronger logical reasoning on SLR-BENCH. Chain-of-Thought (CoT) variants illustrate the effect of explicit reasoning traces, while increasing k reflects test-time sampling improvements.

Table 13: **CLUTRR Evaluation.** Accuracy (%) of base, SLR-tuned, and reasoning models on CLUTRR test sets (greedy decoding). SLR-tuned outperforms both base and reasoning models on CLUTRR.

CLUTRR Variant	Llama-3.1-8B-it		
	Base	w/ SLR	w/ DeepSeek-R1
gen_train23_test2to10	10.2	19.1	24.0
gen_train234_test2to10	9.5	16.4	25.0
rob_train_clean_23_test_all	29.1	35.6	26.0
rob_train_sup_23_test_all	43.2	45.2	23.0
rob_train_irr_23_test_all	31.3	34.5	20.0
rob_train_disc_23_test_all	29.7	41.6	18.0
Average	25.5	32.1	23.0

two predicates (e.g., $eastbound(T) : -has_car(T, C)$).

- Under-generalization:** Rules that are overly specific, often containing excessive literals (more than eight) that fail to capture the concise ground truth (e.g., $eastbound(T) : -has_car(T, C), car_len(C, short), \dots$).
- Failure of Abstraction:** The model lacks semantic understanding of the induction task. Instead of lifting specific examples into variables, the model overfits to the prompt context, replacing general logic with grounded atoms (e.g., hardcoding $eastbound(T) : -car1$). This suggests the model is performing text completion on the prompt rather than logical reasoning.
- Degenerate Reasoning:** The model falls into repetition loops, meaningless enumerations (e.g., repeating “iiii”), or fails to produce structured logic entirely.

Table 14: **Success Rate by Output Length Decile.** Binning generations by relative output length (short \rightarrow long) on DeepSeek-R1-Llama-80B. Longer reasoning traces correlate strongly with lower accuracy, indicating that verbosity often signals derailment rather than deeper reasoning.

Decile (short \rightarrow long)	0	1	2	3	4	5	6	7	8	9
Accuracy (%)\uparrow	96	87	87	80	42	7	2	0	0	0

G.2 Quantitative analysis

Quantitative analysis across all models reveals that 7% of generations were short rules (≤ 2 literals, 38% success), 20% overly long rules (≥ 8 literals, 14% success), 1% had grounded heads (0% success), and 8% showed degenerate reasoning patterns (33% success). Degenerations were identified via simple heuristics (e.g., ≥ 10 repeated characters, ≥ 10 enumerated lines, or $\geq 30\%$ dominance of a single character). While approximate, these checks capture systematic breakdowns. Notably, even degenerate outputs sometimes partially recover, with 25% still solving the task.

G.3 Model-Specific Behaviors

While most failure modes emerge across all models, some failure modes are predominantly assignable to some model classes.

Degenerate Reasoning in Distilled Reasoning Models. In contrast, distilled reasoning models (e.g., DeepSeek-R1-Llama-70B) struggle with *Degenerate Reasoning* (Error Type 4). Unlike smaller models that simply fail to abstract, these larger distilled models attempt complex reasoning but frequently derail. As detailed in Table 14, we observed a strong inverse correlation between reason-

ing length and success. Success rates drop from 96% in the shortest decile to 0% in the longest, indicating verbosity reflects derailment rather than deeper reasoning.

Failure of Abstraction in Small Models. Conversely, smaller distilled models (e.g., Llama-3.2-1B) struggle with fundamental *Failure of Abstraction* (Error Type 3). These models lack the capacity to maintain the distinction between the example data (constants) and the logical rule (variables). They rarely attempt shortcuts or fall into complex reasoning loops; instead, they simply parrot the input context or hallucinate grounded atoms, effectively treating the logic induction task as a text completion exercise.

H Code and Licenses

This work introduces and publicly releases several scientific artifacts, including the SLR framework for scalable logical reasoning with large language models, the SLR-BENCH dataset comprising 19,000 tasks across 20 curriculum levels, and associated training, evaluation, and logic validation scripts. All code and data with the logic reward interface will be made publicly available after publication.

All original software developed as part of this research is distributed under the MIT License, while the datasets are released under the Creative Commons Attribution 4.0 International License (CC BY 4.0), unless specified otherwise in the respective repositories. These licenses permit broad academic and research use, as well as modification and redistribution, provided appropriate credit is given to the original authors.

In addition to the artifacts created in this project, several external resources were utilized, including pretrained language models (e.g., Llama, OpenAI, DeepSeek, Gemini) and open-source Python libraries such as HuggingFace Transformers and PyTorch. All third-party resources were used strictly in accordance with their respective licenses and intended research purposes, and are appropriately cited in this paper and in the code repositories.

We further note that AI-based tools were used during the preparation of this work. Specifically, AI-guided writing assistants (such as ChatGPT) were employed to refine scientific text, and GitHub Copilot was used to support code development and debugging. The use of these tools was limited to improving clarity and efficiency; all research

design, results interpretation, and final manuscript decisions were made by the authors.

The intended use of all released code and data is for research, academic, and educational purposes. Commercial use or deployment in production environments is not permitted without explicit permission or legal review. Any derivatives or extensions of the dataset must comply with the original license terms and the conditions of any incorporated sources. Users are encouraged to consult the individual license files provided in each repository for further details.

I Potential Risks

While this work is primarily intended to advance research in logical reasoning with language models, we recognize several potential risks associated with its development and open release. Enhanced reasoning capabilities in LLMs may be misused, for example, in generating persuasive but misleading arguments, automating manipulation, or circumventing safety mechanisms. The resources and benchmarks we provide, although synthetic and research-focused, could be repurposed for unintended or dual-use applications.

Additionally, while our work does not directly contribute to artificial general intelligence (AGI), we acknowledge broader discussions in the AI community regarding the long-term risks of increasingly capable AI systems. We believe the immediate risks of our work relate to dual-use and misuse as described above, and we encourage responsible use and ongoing monitoring of downstream applications as AI capabilities continue to evolve.

Synthesis Process and Outputs

Task Specification:

(i) Language $\mathcal{L} = (\mathcal{V}, \mathcal{G})$:

- Vocabulary \mathcal{V} : Predicates $\mathcal{P} = \{\text{is_red_train}/1, \text{has_car}/2, \text{car_color}/2, \text{car_len}/2\}$; Constants $\mathcal{C} = \{t1, t2, c1, c2, \text{red}, \text{blue}, \text{short}, \text{long}\}$
- Grammar \mathcal{G} : Restricts predicates to apply to compatible constant types.

(ii) Configuration Θ : Rule length $R_{\text{len}} = 2$; Problem size $\kappa = (\kappa_{\text{pos}} = 1, \kappa_{\text{neg}} = 1)$

Synthesis Steps:

1. Rule Synthesis: The RULEGENERATOR produces a latent ground-truth rule R^* :

```
is_red_train(T) :- has_car(T, C), car_color(C, red).
```

2. Background Synthesis (Loop):

Iteration 1 (finds a positive example):

- Sample Background (b_1): ‘has_car(t1, c1). car_color(c1, red).’
- Assign Label: Query $q_1 = \text{is_red_train}(t1)$. Entailment $b_1 \cup R^* \models q_1$ holds. Result: $(1, q_1)$.
- Accept/Reject: $|E^+| < \kappa_{\text{pos}}$, sample is **accepted**. $B \leftarrow b_1, E^+ \leftarrow \{q_1\}$.

Iteration 2 (finds a negative example):

- Sample Background (b_2): ‘has_car(t2, c2). car_color(c2, blue).’
- Assign Label: Query $q_2 = \text{is_red_train}(t2)$. Entailment $b_2 \cup R^* \models q_2$ fails. Result: $(0, q_2)$.
- Accept/Reject: $|E^-| < \kappa_{\text{neg}}$, sample is **accepted**. $B \leftarrow B \cup b_2, E^- \leftarrow \{q_2\}$.

The loop terminates as both target sizes are met. The final task is $\mathcal{I} = (B, E^+, E^-)$.

Final Synthesizer Outputs:

1. Latent Ground-Truth Rule (R^*):

```
is_red_train(T) :- has_car(T, C), car_color(C, red).
```

2. Validation Program (B, E^+, E^-):

```
has_car(t1, c1).
car_color(c1, red).
has_car(t2, c2).
car_color(c2, blue).
is_red_train(t1).
```

3. Instruction Prompt (example formats):

(a) Prolog-style Prompt:

```
% Given the following background knowledge:
has_car(t1, c1).
car_color(c1, red).
has_car(t2, c2).
car_color(c2, blue).
is_red_train(t1).
% Find a rule "is_red_train(T) :-" that solves the bk.
```

(b) Natural Language Prompt:

```
% Given the following background knowledge:
Train t1 has a car c1. The car c1 is red.
Train t2 has a car c2. The car c2 is blue.
% Find a rule "is_red_train(T) :-" that solves the bk.
```

Figure 6: Step-by-step example of the automatic ILP task synthesis process in SLR. Given a task specification, comprising a language and a task config, the synthesizer generates a ground-truth rule, samples background knowledge, assigns positive and negative example labels, and produces symbolic (Prolog-style) or natural-language prompts. The figure illustrates all intermediate steps and the final output of the synthesizer.

Org	Model	Reasoning	LRL	Syntax	Logical-Reasoning Acc. (%)↑				Total Compute	
		enabled	(↑0-20)	(↑%)	Basic	Easy	Medium	Hard	Tok. (↓M)	Costs (↓\$)
🌀	o3	✓	15.5	80	99	93	74	45	4.30	207.24
🌀	gpt-5	✓	15.4	98	100	90	72	46	16.40	103.13
🌀	gpt-5-mini-high	✓	14.2	94	99	83	63	38	13.12	27.98
🌀	o4-mini-high	✓	12.8	88	98	96	40	21	4.62	24.24
🌀	o4-mini	✓	12.3	86	93	88	52	13	3.98	21.43
🌀	gpt-5-mini	✓	12.0	95	99	82	41	20	4.90	11.54
🌀	o1	✓	11.9	68	92	89	41	15	5.19	364.72
🌀	o3-mini	✓	11.6	75	97	90	37	7	4.73	24.71
🌀	Qwen3-235B-A22B-Thinking-2507	✓	10.4	89	82	76	38	12	26.32	15.53
🌀	o4-mini-low	✓	10.3	91	91	81	25	9	0.77	7.26
🌀	Qwen3-235B-A22B-Instruct-2507	✗	10.3	100	99	76	20	10	5.81	4.45
🌀	o1-mini	✓	10.1	95	97	82	20	3	3.65	19.98
🌀	gpt-5-mini-low	✓	9.7	88	97	71	20	7	1.17	4.07
🌀	Qwen3-Next-80B-A3B-Thinking	✓	9.6	96	93	73	18	8	16.03	20.33
🌀	Qwen3-32B	✓	9.0	100	97	64	14	6	7.73	1.91
🌀	Qwen3-30B-A3B-Thinking-2507	✓	9.0	97	97	67	14	3	17.52	4.29
🌀	Qwen3-VL-30B-A3B-Thinking	✓	8.9	83	94	70	12	3	18.07	11.94
🌀	Qwen3-Next-80B-A3B-Instruct	✗	8.8	100	93	64	15	5	8.68	7.68
🌀	R1-Llama-70B ²	✓	8.8	75	98	67	8	4	11.61	5.33
🌀	Gemini-thinking ¹	✓	8.6	83	93	65	13	1	—	—
🌀	gpt-5-nano	✓	8.5	99	97	61	10	2	6.17	2.81
🌀	Qwen3-VL-32B-Thinking	✓	8.4	99	86	66	14	3	19.33	23.82
🌀	Qwen3-4B-Thinking-2507	✓	8.1	99	98	56	8	2	12.71	0.43
🌀	Qwen3-30B-A3B-Instruct-2507	✗	8.0	100	95	57	5	3	7.46	2.08
🌀	Qwen3-30B-A3B	✓	7.9	100	93	57	7	1	8.62	2.33
🌀	Qwen3-VL-8B-Thinking	✓	7.7	97	87	55	9	2	20.77	10.44
🌀	Qwen3-8B	✓	7.6	100	93	50	7	3	10.17	1.66
🌀	Qwen3-VL-4B-Thinking	✓	7.6	86	91	52	6	2	0.00	0.00
🌀	Qwen3-14B	✓	7.5	100	88	53	8	2	8.99	2.34
🌀	gpt-4.5-prev	✗	7.3	100	94	47	5	1	0.37	576.40
🌀	Qwen3-4B	✓	6.9	99	91	41	5	1	9.77	0.35
🌀	Qwen3-4B-Instruct-2507	✗	6.6	100	73	50	6	2	9.40	0.34
🌀	Qwen3-235B-A22B	✗	6.5	100	87	41	3	0	1.05	1.88
🌀	gpt-4o	✗	6.2	100	93	29	2	0	0.26	20.03
🌀	Qwen3-32B	✗	5.9	100	89	28	2	0	0.63	0.49
🌀	Qwen3-14B	✗	5.6	100	86	25	2	0	0.90	0.56
🌀	Llama 3.3-70B	✗	5.6	100	90	22	1	0	0.49	0.82
🌀	gpt-4-turbo	✗	5.4	100	89	18	2	0	0.41	81.30
🌀	Qwen3-Coder-30B-A3B-Instruct	✗	5.3	100	88	17	2	0	3.32	1.27
🌀	Qwen3-30B-A3B	✗	5.2	100	85	19	1	0	1.31	0.73
🌀	Qwen3-8B	✗	4.6	100	78	15	0	0	1.57	0.47
🌀	Qwen3-4B	✗	4.1	100	70	12	1	0	1.47	0.15
🌀	Llama 3.1-8B	✗	3.8	80	70	7	0	0	0.20	0.14
🌀	Qwen3-VL-30B-A3B-Instruct	✗	3.7	100	71	4	0	0	8.21	1.12
🌀	Qwen3-VL-32B-Instruct	✗	3.5	100	68	2	0	0	8.21	2.62
🌀	Qwen3-VL-8B-Instruct	✗	3.3	99	65	1	0	0	8.21	0.33
🌀	Qwen3-1.7B	✓	1.9	100	35	4	0	0	8.21	0.31
🌀	Qwen3-VL-2B-Instruct	✗	1.8	58	36	0	0	0	0.00	0.00
🌀	Qwen3-0.6B	✗	1.5	80	30	0	0	0	4.45	0.22
🌀	Llama 3.2-3B	✗	1.0	25	19	1	0	0	0.10	0.11
🌀	Qwen3-VL-2B-Thinking	✓	0.4	87	8	0	0	0	0.00	0.00
🌀	Qwen3-1.7B	✗	0.3	95	6	1	0	0	1.59	0.15
🌀	Qwen3-0.6B	✓	0.2	98	4	0	0	0	5.57	0.25
🌀	Llama 3.2-1B	✗	0.0	93	0	0	0	0	0.47	0.12

Table 15: **Extended Leaderboard.** Logical Reasoning Level (LRL), syntax score, and stage-specific logical reasoning accuracy (basic, easy, medium, hard) and computational costs. Models are ranked by LRL. Models with reasoning accuracy mode enabled outperform their base counterparts across reasoning stages, though performance decreases as task complexity rises.