

LogicSkills: A Structured Benchmark for Formal Reasoning in Large Language Models

Anonymous ACL submission

Abstract

Large language models have demonstrated notable performance across various logical reasoning benchmarks. However, it remains unclear which core *logical skills* they truly master. To address this, we introduce LOGICSKILLS, a unified benchmark designed to isolate three fundamental skills in formal reasoning: (i) *formal symbolization*—translating premises into first-order logic; (ii) *countermodel construction*—formulating a finite structure in which all premises are true while the conclusion is false; and (iii) *validity assessment*—deciding whether a conclusion follows from a given set of premises. Items are drawn from the two-variable fragment of first-order logic (without identity) and are presented in both natural English and a Carroll-style language with nonce words. All examples are verified for correctness and non-triviality using the SMT solver Z3. Across leading models, performance is high on validity but substantially lower on symbolization and countermodel construction, suggesting reliance on surface-level patterns rather than genuine symbolic or rule-based reasoning.

1 Introduction

Logic, “the science of reasoning,” is not psychology’s project of describing how we *do* reason (Byrne et al., 1993; Rips, 1994) but of how we *ought* to reason (Frege, 1897/1979). It holds that rational agents should infer what their beliefs entail and avoid inconsistency (Harman, 1986; Field, 2009). In artificial intelligence, reasoning is typically evaluated *behaviorally*—by whether systems succeed on tasks humans typically solve using explicit reasoning—without assuming any particular internal method. In this setting, logic exercises serve as external *proxies* for deductive competence: success should reflect sensitivity to logical structure rather than to surface cues or world knowledge. However, standard benchmarks for logical competence (Tafjord et al., 2021; Parmar et al., 2024)




Figure 1: Accuracy of closed models across three skills.

often conflate distinct logical subskills, making it difficult to determine which aspects of reasoning models genuinely master.

For instance, a model might correctly carry out a formal proof when given symbolic premises (as tested in PROOFWRITER (Tafjord et al., 2021) or LOGICBENCH (Parmar et al., 2024)), yet struggle to translate natural-language sentences into the symbolic premises. Likewise, a model might apply modus ponens or other inference rules correctly but falter when asked to find a countermodel that demonstrates the invalidity of a step, revealing weaknesses in its model-theoretic understanding.

To address this issue, we introduce LOGICSKILLS, a benchmark for formal reasoning that disentangles fundamental logical skills and exposes failure modes that composite tasks may hide. Drawing on formal logic pedagogy, LOGICSKILLS isolates three fundamental logical skills: (i) *formal symbolization* (mapping sentences to logical form), (ii) *countermodel construction* (demonstrating invalidity via model-theoretic falsification), and (iii) *validity assessment* (recognizing what follows from what). LOGICSKILLS is solver-verified (via Z3 (de Moura and Bjørner, 2008)) and includes a bilingual task design (English and Carrollian), with the Carrollian nonce-word condition isolating reasoning from prior semantic knowledge.

Systematic evaluations on LOGICSKILLS across frontier open-weight and closed-weight large language models (LLMs) reveal a consistent profile: strong performance on validity assessment, but sub-

stantial weaknesses in symbolization and countermodel construction (see Figure 1). Together, these results suggest that various LLMs achieve high apparent reasoning accuracy through pattern-based approximations that fall short of the symbolic or rule-based procedures traditionally used in logic.

2 Related Work

Prior evaluations of LLM reasoning divide broadly into two traditions. *Puzzle-based benchmarks* (e.g., knights-and-knaves, logic-grid, and minesweeper tasks) probe combinatorial reasoning often revealing brittleness under search and small perturbations (Mondorf and Plank, 2024; Lin et al., 2025; Giadikiaroglou et al., 2024; Li et al., 2024). *Rule- or proof-oriented suites* such as PROOFWRITER, LOGICBENCH, and LOGICASKER assess inference templates across propositional and first-order logic (Tafjord et al., 2021; Parmar et al., 2024; Han et al., 2024; Wan et al., 2024), but do not disentangle the subskills they presuppose.

LOGICSKILLS complements both traditions by offering a structured, skill-level evaluation framework rather than a single composite reasoning task. The following section presents an overview of the benchmark and details each task type.

3 Benchmark Overview

We build a compositional generator that produces grammatical natural-language sentences paired with formulas in the two-variable fragment of first-order logic (FO^2 , without identity). We choose FO^2 since this logic retains decidability and supports finite satisfiability while being expressive enough to capture nontrivial reasoning (Scott, 1962; Mortimer, 1975). Every formula is realized twice: (i) controlled English and (ii) a Carroll-style nonsense lexicon (Carroll, 1871). The latter provides a structurally identical surface form that removes real-world associations and isolates pure logical competence. For example, both the following language examples, i.e., “A human chased a donkey that chased it” and “A borogove snicker-snacked a tove that snicker-snacked it” map to the formula:

$$\exists x \exists y ((Nx \wedge My) \wedge (Rxy \wedge Ryx)) \quad (1)$$

By varying quantifier structure, predicate choice, name assignments, connectives, and recursive clausal patterns, the generator can produce indefinitely many grammatical sentences in both English and Carrollian.

Because every sentence is paired with an FO^2 formula, we leverage SMT solvers (Davis and Putnam, 1960)—in particular the state-of-the-art theorem prover Z3 (de Moura and Bjørner, 2008)—to compute core semantic properties: (i) whether two formulas are equivalent, (ii) whether a set of formulas is jointly satisfiable, (iii) whether one formula follows from others, and (iv) whether a proposed finite model satisfies (or falsifies) a given argument. These capabilities enable large-scale, correctness-guaranteed construction and verification of all deductive tasks in the benchmark. In this way, Z3 is essential in generating arguments. To construct arguments, we sample sets of 3–5 jointly satisfiable premises and use Z3 to search for semantically relevant, non-trivial conclusions that follow from them (valid). For each valid conclusion, we additionally generate 5 structurally similar distractor conclusions—by matching sentence type and symbol set—that do not follow (invalid). Appendix B provides further details about the dataset and task construction pipeline, including sentence generation, argument filtering, and task instantiation.

3.1 Task types

We instantiate three task types, each aligned with a canonical logical skill (Kalish and Montague, 1964). Exact task prompts can be found in appendix E.

Formal symbolization Given a sentence in English/Carroll and a fixed symbol key, models must output a single well-formed formula capturing the sentence using only the provided predicates, constants, and standard logical operators (i.e., $\neg, \wedge, \vee, \rightarrow, \leftrightarrow, \forall, \exists$).

Input:

Sentence: “All raths are uffish.”
Key: F : “ x is a rath”, G : “ x is uffish”.

Output: $\forall x (Fx \rightarrow Gx)$

Countermodel construction For an invalid argument, models must provide a finite structure (on a fixed small domain) that makes all premises true and the conclusion false.

Input:

Invalid Argument: $\forall x (Fx \rightarrow Gx), Ga \models Fa$.
Domain: $\{0, 1\}$

Output: $F = \{0\}, G = \{0, 1\}, a = 1$.

Argument validity Given premises and six candidate conclusions in English/Carroll, models must decide which conclusion(s), if any, *must* follow.

Input:

Premises: “All raths are uffish”; “Zindle is a rath.”
Candidate Conclusions:
1. Zindle is not uffish.
2. If Zindle is uffish, then Zindle is a rath.
3. Zindle is uffish.

Output: 3

4 Evaluation Setup

LOGICSKILLS comprises a fixed evaluation set of 1500 problems randomly sampled from generated sentences and arguments: 600 *symbolization* problems (300 English, 300 Carroll), 600 *validity* problems (300 English, 300 Carroll), and 300 *countermodel* problems. Validity assessment and formal symbolization are evaluated bilingually (English and Carroll), whereas countermodel construction is language-neutral.

We evaluate a range of open- and closed-weight LLMs from multiple vendors and across various sizes: *Llama-3.1-8B*, *Llama-3.1-70B*, *Qwen2.5-Math-72B*, *Qwen3-32B*, *Claude-3.7-Sonnet*, *GPT-4o*, *Gemini-2.5-Flash*, *DeepSeek-Chat*, and *Phi-4*. API-based models (OpenAI, Anthropic, Google, DeepSeek) were accessed between July and August 2025. Model responses are processed in two stages. First, an extractor LLM (GPT-4o) cleans and normalizes raw completions.¹ Second, task-specific verification procedures check correctness:

Symbolization Normalized formulas are first checked for an exact match; if necessary, they are parsed and repaired. Cosmetic variants (e.g., \wedge as $\&$, \neg as \sim , \rightarrow as \supset , or $R(x, y)$ vs. Rxy) are handled via the extractor LLM. Final parses are validated for logical equivalence using Z3.

Countermodels Candidate models are normalized into domain, constants, and predicate extensions. The system enforces *symbol completeness* and validates type/domain consistency (integer domains, correct arities). Candidate models are then translated into SMT-LIB, merged with the negated argument, and checked for satisfiability with Z3.

Validity Extracted labels are compared directly against the gold annotations.

¹A meta-evaluation shows 98–99% extraction accuracy; see Appendix D for more details.

5 Results

Figure 2 illustrates the models’ accuracy across the three different task types described in Section 3. The dominant pattern is a strong *validity* performance contrasted with weaker *symbolization* and especially *countermodel* performance. For instance, GPT-4o achieves 87.0% on validity but only 10.0% on countermodels; Gemini-2.5-Flash reaches 96.0% vs. 13.0%, respectively. The main exception is *Qwen3-32B*, which performs strongly across all three skills (97.0% validity, 85.0% symbolization, 89.0% countermodels).

This striking deviation appears correlated with Qwen3-32B’s inference strategy. Whereas most models rely on generic “reasoning rhetoric” (e.g., “Let’s break this down step-by-step”), Qwen3-32B consistently engages in spontaneous self-scaffolding, marked by explicit internal delimiters (e.g., `<think>`). Its chain-of-thought traces are, on average, $7.5\times$ longer than those of its peers, coinciding with its superior performance across all three tasks.² Per-model, per-task accuracies ($n=300$ per condition) are shown in Table 3 in the Appendix.

Across top-tier models, *validity* accuracy clusters near ceiling, suggesting that premise–conclusion classification is largely mastered under the given conditions. In contrast, *symbolization* remains a consistent bottleneck, reflecting persistent difficulty in mapping natural-language statements to formal representations even when inference competence is high. The *countermodel* task is the most discriminative, revealing substantial gaps in model-theoretic reasoning—with all but Qwen3 failing to generate consistent falsifying models.

5.1 Bilingual evaluation

Perhaps surprisingly, bilingual evaluation shows only minor differences: averaging across models, English and Carroll results are closely aligned (validity 81.1% vs. 82.8%; symbolization 68.1% vs. 66.0%), with a maximum difference of 7 percentage points across models. This null result implies that the models are not relying on semantic heuristics to process the syntactic structure. For an overview of accuracy values per model, task and language, we refer to Table 3 in the Appendix.

²Notably, this scaffolding emerges uniformly across all task types. See (Yang et al., 2025) for details on the model’s “thinking mode” architecture and large-scale reasoning training regime.




Figure 2: **Skill profile across models** (accuracy %; $n=300$ per task/language). Grouped bars show *Validity*, *Symbolization*, and *Countermodel*. Most models score highly on validity but lag on subskills, especially countermodels.

5.2 Error analysis

Error analyses across tasks reveal distinct but converging weaknesses in model reasoning. In the *symbolization* task, nearly half of the 1,775 incorrect outputs (46.9%) are semantically wrong despite syntactic well-formedness, while the remainder show structural issues such as predicate-arity mismatches or ill-formed parentheses. In the *countermodel* task, most erroneous outputs (72.7% of 2,010) are structurally well-formed but fail to falsify the argument—typically because one or more premises are false—while the rest (27.3%) reflect missing interpretations or malformed predicate definitions. In the *validity* task, a majority of the 973 incorrect outputs (57.6%) include the correct answer alongside additional invalid options, suggesting difficulty in eliminating plausible distractors. Further details about the error analysis procedure can be found in Appendix C.

5.3 Subskill Transfer via PEFT

In principle, skill at symbolization and countermodel construction should reinforce validity assessment: to judge validity, a model must

Train Data	Symb. Δ	CM Δ	Val. Δ
Symbolization (100k)	+80.7	0.0	-4.3
Countermodel (100k)	-3.0	+15.7	-6.4
Symb + Counter (200k)	+84.0	+25.6	-5.7

Table 1: Transfer deltas (points vs. base).

map sentences to their logical form and evaluate whether any counterexample structure can satisfy the premises while falsifying the conclusion. To assess potential cross-skill transfer, we conduct a fine-tuning study. We fine-tune *Llama-3.2-3B-Instruct* with LoRA adapters on 100k *symbolization* and 100k *countermodel* examples (targets generated via tpg; Schwarz, 2025), using a single epoch with AdamW and standard regularization. After fine-tuning, we observe large *in-task* gains but no measurable transfer to the validity assessment task. Yet, joint fine-tuning on symbolization and countermodel data yields larger countermodel gains than countermodel-only training (+25.6 vs. +15.7), indicating task interaction. Table 1 summarizes the corresponding transfer matrix.

6 Conclusion

From Leibniz’s (1685) call to “let us calculate” to Frege’s *Begriffsschrift* (1879), the project of logic has rested on the conviction that reasoning becomes reliable only when translated into a formal system—where errors can be detected and validity established by explicit rules. Taken together, the results presented in this study suggest that LLMs approximate reasoning success without the explicit symbolic and model-based procedures that underpin the standard methods of logic. Arguably, a system without such explicit control structures cannot satisfy the normative standards that make reasoning accountable and corrigible.

299 Limitations

300 While our study introduces a novel benchmark
301 for formal reasoning that disentangles core logi-
302 cal skills and exposes failure modes that composite
303 tasks may hide, there are several limitations that
304 future work can address.

305 **Logical and linguistic scope** LOGICSKILLS
306 evaluates reasoning within a deliberately restricted
307 logical and linguistic scope. All tasks are drawn
308 from the two-variable fragment of first-order logic
309 without identity, excluding richer logical resources
310 such as full first-order logic, identity, or modal oper-
311 ators. Consequently, results may not generalize to
312 reasoning in more expressive logical systems. The
313 benchmark also relies on a controlled fragment of
314 English and a synthetically generated Carrollian
315 language. While this design isolates logical struc-
316 ture, it limits linguistic diversity and does not cap-
317 ture many sources of difficulty present in naturally
318 occurring language.

319 **Evaluation procedure** Model performance is
320 additionally sensitive to prompting and evaluation
321 protocols. In particular, the countermodel and sym-
322 bolization tasks require fully specified structured
323 outputs, so measured performance may reflect
324 both reasoning ability and success in satisfying
325 interface and formatting constraints. We partially
326 mitigate this by using liberal extraction and
327 normalization procedures that permit notational
328 variants and apply automated repairs, though
329 some sensitivity remains. Finally, LOGICSKILLS
330 provides a behavioral evaluation only. While
331 it reveals systematic differences across logical
332 subskills, it does not directly address the internal
333 representations or mechanisms by which models
334 arrive at their answers (Chalmers, 2025; Stolfo
335 et al., 2023).

336
337 Within the scope of this paper, these limitations
338 are largely intentional design choices, reflecting a
339 tradeoff between expressivity and reliable, solver-
340 verifiable evaluation.

341 References

342 Ruth MJ Byrne, Jonathan St BT Evans, and Stephen E
343 Newstead. 1993. *Human reasoning: The psychology*
344 *of deduction*. Psychology Press.

345 Lewis Carroll. 1871. *Through the Looking-Glass, and*
346 *What Alice Found There*. Macmillan and Co., Lon-
347 don. Chapter 1, pp. 21–24.

David J. Chalmers. 2025. [Propositional interpretability in artificial intelligence](#). *Preprint*, arXiv:2501.15740. 348 349

Martin Davis and Hilary Putnam. 1960. [A computing procedure for quantification theory](#). *J. ACM*, 7(3):201–215. 350 351 352

Leonardo de Moura and Nikolaj Bjørner. 2008. Z3: An efficient smt solver. In *Tools and Algorithms for the Construction and Analysis of Systems*, pages 337–340, Berlin, Heidelberg. Springer Berlin Heidelberg. 353 354 355 356

Hartry Field. 2009. [What is the normative role of logic?](#) *Aristotelian Society Supplementary Volume*, 83(1):251–268. 357 358 359

Gottlob Frege. 1879. *Begriffsschrift: eine der arithmetischen nachgebildete Formelsprache des reinen Denkens*. Verlag von Louis Nebert, Halle. English translation in Jean van Heijenoort (ed.), *From Frege to Gödel: A Source Book in Mathematical Logic, 1879–1931*, Harvard University Press, 1967, pp. 1–82. 360 361 362 363 364 365 366

Gottlob Frege. 1897/1979. Logic. In Hans Hermes, Friedrich Kambartel, and Friedrich Kaulbach, editors, *Posthumous Writings*, pages 126–151. University of Chicago Press, Chicago. Originally written in 1897. 367 368 369 370

Panagiotis Giadikiaroglou, Maria Lymperaïou, Giorgos Filandrianos, and Giorgos Stamou. 2024. [Puzzle solving using reasoning of large language models: A survey](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11574–11591, Miami, Florida, USA. Association for Computational Linguistics. 371 372 373 374 375 376 377

Simeng Han, Hailey Schoelkopf, Yilun Zhao, Zhen-ting Qi, Martin Riddell, Wenfei Zhou, James Coady, David Peng, Yujie Qiao, Luke Benson, Lucy Sun, Alexander Wardle-Solano, Hannah Szabó, Ekaterina Zubova, Matthew Burtell, Jonathan Fan, Yixin Liu, Brian Wong, Malcolm Sailor, and 16 others. 2024. [FOLIO: Natural language reasoning with first-order logic](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 22017–22031, Miami, Florida, USA. Association for Computational Linguistics. 378 379 380 381 382 383 384 385 386 387 388

Gilbert Harman. 1986. *Change in View: Principles of Reasoning*. MIT Press, Cambridge, MA, USA. 389 390

Donald Kalish and Richard Montague. 1964. *Logic: Techniques of Formal Reasoning*. Oxford University Press. 391 392 393

Gottfried Wilhelm Leibniz. 1685. *Generales inquisitiones de analysi notionum et veritatum*. English translation in Leroy E. Loemker (ed.), *Philosophical Papers and Letters*, 2nd ed., D. Reidel, 1969, pp. 321–325. 394 395 396 397 398

Yinghao Li, Haorui Wang, and Chao Zhang. 2024. [Assessing logical puzzle solving in large language models: Insights from a minesweeper case study](#). In *Proceedings of the 2024 Conference of the North* 399 400 401 402

510	and constants), preferentially match its syntactic	used to distinguish incorrect selections from <i>super-</i>	558
511	type, and be solver-verified as non-entailed. Ar-	<i>set</i> errors in which the correct answer was included	559
512	arguments are stored in a database with fields for	alongside additional invalid options.	560
513	premise identifiers, conclusion identifier, validity	All error analyses were performed programmat-	561
514	label, and language, and form the large shared pool	ically over the full evaluation set, and the result-	562
515	from which validity and countermodel tasks are	ing error-type distributions correspond to those re-	563
516	derived.	ported in the main text.	564
517	Task instantiation. Three task families are de-	D Meta-Evaluation of Answer Extraction	565
518	derived from the sentence and argument tables. For-	Raw model responses were post-processed by an	566
519	mal symbolization tasks are constructed from indi-	extractor LLM to normalize outputs into task-	567
520	vidual sentences paired with their SOA, with the	specific formats prior to scoring. To assess the	568
521	associated FO ² formula as the target. Although	reliability of this extraction step independently of	569
522	the fixed lexicons permit a larger vocabulary in	task correctness, we conducted a meta-evaluation.	570
523	principle, in practice sentences used to instantiate	We randomly sampled 100 evaluation instances,	571
524	tasks involve only 3–4 predicates and 1–2 constants	balanced across all task types (symbolization, valid-	572
525	(4–6 symbols total). Argument validity tasks are	ity, and countermodel construction) and languages	573
526	constructed from solver-certified valid arguments,	where applicable, and including both correct and	574
527	each paired with five invalid alternatives sharing	incorrect model responses. For each instance, the	575
528	the same premises. Countermodel construction	evaluator was given the original task prompt, the	576
529	tasks are constructed from solver-certified invalid	model’s raw response, and the extracted answer	577
530	arguments and require a finite model that satisfies	produced by the extractor.	578
531	all premises while falsifying the conclusion.	The meta-evaluator (GPT-4o) rendered a bi-	579
532	C Error Analysis Methodology	nary judgment—faithful or unfaithful—indicating	580
533	Our evaluation pipeline produced structured out-	whether the extracted answer accurately reflected	581
534	puts beyond binary correctness labels, enabling	the content of the raw response, independent of	582
535	systematic error analysis across all tasks. For each	correctness. Five cases were flagged as potentially	583
536	instance, we retained the raw model response, the	unfaithful.	584
537	extracted or parsed answer, the gold-standard an-	All flagged cases, together with a random sam-	585
538	swer, and task-specific diagnostic metadata from	ple of 10 unflagged instances, were subsequently	586
539	the evaluator.	reviewed by a human evaluator (a meta-meta-	587
540	Formal symbolization For <i>symbolization</i> , each	evaluation). Of the five flagged cases, one consti-	588
541	response included an assessment label (e.g., <i>Not</i>	tuted a clear extraction error, in which the model re-	589
542	<i>logically equivalent</i> , <i>Failed to parse</i>). Incorrect	vised a predicate assignment mid-response and the	590
543	outputs were filtered and categorized by matching	extractor failed to capture the update. One case in-	591
544	these labels, allowing semantic errors to be distin-	volved a normalization failure, where a malformed	592
545	guished from syntactic or structural failures.	but correctly extracted formula was not repaired.	593
546	Countermodel construction For <i>countermodel</i>	The remaining three cases involved internally in-	594
547	<i>construction</i> , the model checker returned a list	consistent raw responses (e.g., conflicting predicate	595
548	of diagnostic error messages (e.g., <i>Not a counter-</i>	arities or self-contradictory conclusions); in these	596
549	<i>model</i> , <i>Missing interpretation for: X</i> , <i>Binary pred-</i>	cases, the extractor’s output was judged reasonable	597
550	<i>icate ‘P’ must be a list of pairs</i> , etc.). Failures	given the ambiguity of the source response.	598
551	were categorized by matching error-message pat-	After human adjudication, only one instance was	599
552	terns, with additional checks performed by evaluat-	deemed a clear extraction failure, with one addi-	600
553	ing premises and conclusions against the proposed	tional case judged ambiguous, corresponding to	601
554	structure when needed.	an effective extraction accuracy of approximately	602
555	Argument validity For <i>validity</i> , both the ex-	98–99%.	603
556	tracted and correct answers were represented as	E Task Prompts	604
557	sets of statement indices. Set comparisons were	Figures 3 to 5 illustrate the prompts used for each	605
		task type described in Section 3.1.	606

607

F Additional Results

608

Table 3 reports the complete per-model results underlying §5.

609

610

G Use of AI Assistants

611

We used GitHub Copilot for parts of the project’s source code and ChatGPT to correct minor grammatical errors.

612

613

ARGUMENT VALIDITY PROMPT

Your task is to solve a logical reasoning problem. Use any approach you find effective, but clearly and explicitly state your final answer.

Task

Consider the following situation:

Everything is a tove, a borogove, or a rath (exclusively), and there's at least one of each. Zindle or Bungo will whiffle. Only toves will whiffle. Every rath chortled at Bungo. If Zindle will whiffle, then every rath is mimsy only if every borogove chortled at Bungo.

Which, if any, of the following statements must be true in this situation?

1. Zindle and Bungo will whiffle.
2. Every tove will whiffle.
3. Not all toves will whiffle.
4. A tove will whiffle.
5. No toves will whiffle.
6. Zindle will whiffle.

Figure 3: Illustrative instance of the *Argument Validity* task.

FORMAL SYMBOLIZATION PROMPT

Your task is to translate the provided sentence into formal predicate logic, using the abbreviations provided.

Instructions

- Use only the abbreviations given.
- Return a single well-formed formula in standard predicate logic syntax.
- Use standard logical symbols:
 - Quantifiers: \forall , \exists
 - Connectives: \neg , \wedge , \vee , \rightarrow , \leftrightarrow
 - Do not include any explanation or extra text—just the formula.

Example

Sentence: Every linguist admires Charlie.

Abbreviations:

- L: “[1] is a linguist”
- R: “[1] admires [2]”
- c: “Charlie”

Translation: $\forall x(Lx \rightarrow Rxc)$

Task

Sentence: Not all raths will whiffle and no toves chortled at Bungo.

Abbreviations:

- O: “[1] is a rath”
- F: “[1] will whiffle”
- M: “[1] is a tove”
- P: “[1] chortled at [2]”
- c: “Bungo”

Figure 4: Illustrative instance of the *Formal Symbolization* task.

COUNTERMODEL CONSTRUCTION PROMPT

Show that the provided argument is invalid by giving a countermodel—one where all premises are true and the conclusion is false.

Instructions

1. Provide assignments for all constants and predicates used in the argument.
2. Respect predicate arity:
 - Monadic predicates take one argument (e.g., Mx).
 - Binary predicates take two arguments (e.g., Pxy).
3. Use the fixed domain $[0,1,2]$.

Required Format

- Domain: $[0,1,2]$
- Constants: “a”: 0
- Monadic predicates: “F”: $[0,2]$
- Binary predicates: “R”: $[[0,1],[2,0]]$

Argument

$$\begin{aligned} & (\forall x (Mx \vee (Nx \vee Ox)) \wedge (\neg \exists x (Mx \wedge Nx) \wedge (\neg \exists x (Mx \wedge Ox) \wedge \neg \exists x (Nx \wedge Ox))))), \\ & (\exists x Mx \wedge (\exists x Nx \wedge \exists x Ox)), \\ & \neg \exists x \exists y ((Mx \wedge My) \wedge Pxy), \\ & (\neg \exists x (Ox \wedge Qxc) \rightarrow Mb), \\ & (\neg \exists x (Nx \wedge Rxc) \rightarrow \exists x (Nx \wedge \neg Rxb)), \\ & (\forall x (Ox \rightarrow Rxc) \rightarrow (Mb \rightarrow \forall x (Nx \rightarrow Pxe))), \\ & (Pcc \wedge \forall x (Ox \rightarrow Rxa)) \\ & \models Mc \end{aligned}$$

Figure 5: Illustrative instance of the *Countermodel Construction* task.

Sentence type	English example	Carroll example
atomic_monadic	Hazel drank.	Bungo gyred.
atomic_dyadic	Hazel kicked Lewis.	Bungo galumphed over Rafin.
negation	Lewis isn't happy.	Rafin isn't uffish.
quantified_universal_affirmative	Every monkey is asleep.	Every rath is mimsy.
quantified_particular_affirmative	A donkey will run.	A tove will whiffle.
quantified_universal_negative	No humans chased Alfred.	No borogoves snicker-snacked Zindle.
quantified_particular_negative	Not all donkeys chased Alfred.	Not all toves snicker-snacked Zindle.
quantified_only_restrictor	Only humans will run.	Only borogoves will whiffle.
quantified_name_all	Hazel chased every human.	Bungo snicker-snacked every borogove.
quantified_name_some	Alfred chased a donkey.	Zindle snicker-snacked a tove.
quantified_all_all	Every monkey kicked every donkey.	Every rath galumphed over every tove.
quantified_all_all_all	Every donkey saw every donkey that saw every monkey.	Every tove chortled at every tove that chortled at every rath.
quantified_all_all_back	Every donkey chased every monkey that chased it.	Every tove snicker-snacked every rath that snicker-snacked it.
quantified_all_some	Every monkey kicked a human.	Every rath galumphed over a borogove.
quantified_all_some_back	Every human saw a donkey that saw it.	Every borogove chortled at a tove that chortled at it.
quantified_some_all	A monkey chased every human.	A rath snicker-snacked every borogove.
quantified_some_all_back	A monkey saw every human that saw it.	A rath chortled at every borogove that chortled at it.
quantified_some_some	A human chased a donkey.	A borogove snicker-snacked a tove.
quantified_some_some_back	A donkey kicked a monkey that kicked it.	A tove galumphed over a rath that galumphed over it.
quantified_some_some_some	A human kicked a human that kicked a monkey.	A borogove galumphed over a borogove that galumphed over a rath.
quantified_no_all	No humans kicked every human.	No borogoves galumphed over every borogove.
quantified_no_some	No humans chased a human.	No borogoves snicker-snacked a borogove.
quantified_no_some_back	No donkeys kicked a monkey that kicked it.	No toves galumphed over a rath that galumphed over it.
quantified_rev_some_all	There is a donkey that every monkey kicked.	There is a tove that every rath galumphed over.
quantified_rev_no_all	There is not a monkey that every human chased.	There is not a rath that every borogove snicker-snacked.
quantified_some_self	A donkey kicked itself.	A tove galumphed over itself.
conjunction_simple	Hazel saw Lewis and not all monkeys kicked Lewis.	Bungo chortled at Rafin and not all raths galumphed over Rafin.
disjunction_simple	Alfred drank or a monkey kicked Alfred.	Zindle gyred or a rath galumphed over Zindle.
conditional_if_then	If Hazel is a human, then Alfred isn't a human.	If Bungo is a borogove, then Zindle isn't a borogove.
biconditional_just_in_case	Hazel is happy just in case every donkey chased Lewis.	Bungo is uffish just in case every tove snicker-snacked Rafin.

Table 2: Sentence types used in LogicSkills with representative English and Carroll examples.

Model	Task	Language	Total	Correct	Accuracy
meta_llama-3.1-8B-instruct	Validity	Carroll	300	138	0.46
	Validity	English	300	135	0.45
	Symbolization	Carroll	300	42	0.14
	Symbolization	English	300	41	0.14
	Countermodel	–	300	0	0.00
meta_llama-3.1-70B-instruct	Validity	Carroll	300	251	0.84
	Validity	English	300	257	0.86
	Symbolization	Carroll	300	198	0.66
	Symbolization	English	300	206	0.69
	Countermodel	–	300	36	0.12
qwen2.5-math-72B-instruct	Validity	Carroll	300	241	0.80
	Validity	English	300	239	0.80
	Symbolization	Carroll	300	241	0.80
	Symbolization	English	300	240	0.80
	Countermodel	–	300	44	0.15
qwen3-32B	Validity	Carroll	300	284	0.95
	Validity	English	300	292	0.97
	Symbolization	Carroll	300	247	0.82
	Symbolization	English	300	255	0.85
	Countermodel	–	300	261	0.89
anthropic_claude-3.7-sonnet	Validity	Carroll	300	287	0.96
	Validity	English	300	282	0.94
	Symbolization	Carroll	300	214	0.71
	Symbolization	English	300	228	0.76
	Countermodel	–	300	126	0.42
openai/gpt-4o	Validity	Carroll	300	262	0.87
	Validity	English	300	256	0.85
	Symbolization	Carroll	300	223	0.74
	Symbolization	English	300	217	0.72
	Countermodel	–	300	30	0.10
google_gemini-2.5-flash	Validity	Carroll	300	288	0.96
	Validity	English	300	283	0.94
	Symbolization	Carroll	300	229	0.76
	Symbolization	English	300	238	0.79
	Countermodel	–	300	40	0.13
deepseek_chat	Validity	Carroll	300	244	0.81
	Validity	English	300	229	0.76
	Symbolization	Carroll	300	240	0.80
	Symbolization	English	300	242	0.81
	Countermodel	–	300	104	0.35
microsoft_phi-4	Validity	Carroll	300	239	0.80
	Validity	English	300	220	0.73
	Symbolization	Carroll	300	152	0.51
	Symbolization	English	300	172	0.57
	Countermodel	–	300	43	0.14

Table 3: Accuracy across tasks, languages, and models.