## A BALANCED NEURO-SYMBOLIC APPROACH FOR COM-MONSENSE ABDUCTIVE LOGIC

**Anonymous authors**Paper under double-blind review

### **ABSTRACT**

Although Large Language Models (LLMs) have demonstrated impressive formal reasoning abilities, they often break down when problems require complex proof planning. One promising approach for improving LLM reasoning abilities involves translating problems into formal logic and using a logic solver. Although off-the-shelf logic solvers are in principle substantially more efficient than LLMs at logical reasoning, they assume that all relevant facts are provided in a question and are unable to deal with missing commonsense relations. In this work, we propose a novel method that uses feedback from the logic solver to augment a logic problem with commonsense relations provided by the LLM, in an iterative manner. This involves a search procedure through potential commonsense assumptions to maximize the chance of finding useful facts while keeping cost tractable. On a collection of pure-logical reasoning datasets, from which some commonsense information has been removed, our method consistently achieves considerable improvements over existing techniques, demonstrating the value in balancing neural and symbolic elements when working in human contexts.

#### 1 Introduction

Large Language Models (LLMs) have demonstrated impressive abilities to reason formally, often via chain-of-thought reasoning (Wei et al., 2022). However, despite substantial efforts, there are increasing signs that such neural methods will never scale to problems that require complex proof planning (Saparov and He, 2023; Dziri et al., 2023). Such problems are exactly the type on which symbolic logical solvers excel: such solvers have a long history and were for a long time considered a key component of any path to artificial intelligence (Nilsson, 1991). Nevertheless, they are greatly restricted by their need for problems to be stated in symbolic language and for every relevant fact to be provided as input. These constraints have ultimately limited them to highly specialized applications, and they have never had the broad impact that was hoped for (Crevier, 1993).

These complimentary strengths of neural and symbolic methods have motivated a revival of interest in neuro-symbolic methods, where an LLM incorporates a logic solver to improve its reasoning abilities (Ye et al., 2023; Lee and Hwang, 2024; Lyu et al., 2023; Olausson et al., 2023). In these approaches, the LLM translates problems formulated in natural language into symbolic language, addressing one of the key deficiencies of a purely symbolic approach. Nonetheless, these hybrid systems remain impractical because they are ultimately purely deductive: that is, every relevant fact must be provided as input. This means that the symbolic solvers are often unable to reach a conclusion simply because obvious, commonsense assumptions are left unstated, and it is often difficult to predict which should be included until one is presented with a failed reasoning chain.

For example, consider the problem in Figure 1. A logic solver would return "unknown" for the target query as, formally speaking, neither its truth nor its falsehood is implied by the premises. A human, however, would easily solve this problem by supplying the additional commonsense fact that white surfaces reflect light  $(turns\_white(fox,winter) \rightarrow reflects(fox,sun))$ . This ability to supply missing information is usually themed abductive reasoning, and is a key mark of human intelligence.

The limitation of current neuro-symbolic LLM systems to deductive reasoning means that they have mostly been so far of theoretical interest, since they tend to break down when confronted with more complex problems where enumerating every possible background fact is not realistic. However, besides their translation skills, LLMs possess also another striking ability: their training on prodigious

Context: Some animals tough winter out. They do not leave. They do not hide. They must survive. Sometimes nature helps them out. Some animals grow thicker coats in the winter. Other animals change color. The arctic fox is brown in the summer. His coat turns white in the winter.

Question: The arctic fox's coat turns white in the winter because white absorbs the sun and is warmer.

```
\exists x : tough\_out(x, winter)
\exists x : \neg hide(x) \land \neg leave(x) \land survive(x)
\exists x : helps(nature, x)
\exists x : grow\_coat(x, winter)
\exists x : change\_color(x, winter)
brown(fox, summer)
turns\_white(fox, winter)
absorbs(white, sun)
warmer(coat)
```

turns\_white(fox, winter)

Figure 1: An example from a children's comprehension exercise booklet <sup>1</sup>. Left: the problem phrased in human language. Right: the same problem translated to first-order-logic.

amounts of internet data has made them very adept at recognizing commonsense statements, to the point where they have been regarded as potential universal databases (Petroni et al., 2019). In a way, LLMs seem to have internalized most commonsense knowledge.

This realization has led some works to use an LLM itself to supply missing but relevant clauses when reasoning. Notably, Toroghi et al. (2024) proposed a method that operates an exhaustive search over a heavily restrained set of rules in the symbolic space, whereas Liu et al. (2024) proposed a method that uses LLM prompting to produce new rules which might be deducible from the given logical context. Whiel these methods lie on opposite ends of the symbolic-linguistic reasoning spectrum (Figure 2), they both limit themselves to searching over such a restricted space of possible commonsense that they cannot solve practical problems.



Figure 2: Symbolic-Linguistic Spectrum depicting the positioning of LLM-Tres (Toroghi et al., 2024), Logic-of-Thought (Liu et al., 2024), and Chain-of-Thought (COT) (Wei et al., 2022) relative to our approach.

In this work, we seek to improve AI reasoning abilities by using an LLM to provide relevant unstated commonsense clauses to a logic solver, but unlike previous works, without imposing significant constraints on the shape or content of such clauses. Furthermore, and most importantly, our method ARGOS (Abductive Reasoning with Generalization Over Symbolics) can abduce clauses which contain variables not previously mentioned in the input problem. To compensate for the far more general search space, we guide the search using feedback from the logic solver in the form of the SAT problem backbone, another novel contribution. The resulting system strikes a balance between linguistic and symbolic approaches, allowing us to use both their strengths while minimizing their weaknesses to achieve true abductive reasoning.

The contributions of this paper are as follows.

- We propose a novel framing of the commonsense logical reasoning problem founded upon classical logical principles and an aim towards more realistic use-cases.
- We introduce a novel algorithm that (i) searches over larger spaces of commonsense facts; and (2) uses logic solver feedback in the form of the backbone graph to increase practicality and efficiency.
- We demonstrate empirically on multiple benchmarks and large language models that our method improves substantially over existing symbolic and neural methods on abductive reasoning problems where background information is missing.

<sup>&</sup>lt;sup>1</sup>Taken from https://www.ereadingworksheets.com/worksheets/reading/nonfiction-passages/wintertime. We selected a choice from multiple choice question 3 and re-phrased it as a True/False question, according to the logic-problem framing.

## 2 RELATED WORK

Previous LLM-related logical reasoning methods combine symbolic and neural approaches, but usually rely much more on one or the other. Appendix E provides an extended review.

Neural Methods Wei et al. (2022) were the first to present a framework for LLM-based reasoning, showing that providing examples of rationales for answers to questions can induce the LLM to do the same, leading to improved accuracy. Kojima et al. (2022) showed that this can be induced without any few-shot examples by prepending the sentence "Let's think step by step" before generating an answer. This is known as "Chain of Thought" (COT). Following this, Wang et al. (2023) proposed self-consistency (SC), using COT multiple times and taking the mode as the prediction. However, Saparov and He (2023) observed that COT and SC suffer from challenges in proof planning — rationale steps tend to be factual but of low value. This motivated guidance of the LLM at a step-level. Yao et al. (2023) proposed Tree of Thoughts (TOT), which explores hand-crafted trees using an LLM to solve reasoning tasks. TOT is poorly suited to logical reasoning settings as logic problems have highly variable tree-structures. Kazemi et al. (2023) and Lee and Hwang (2024) proposed more logic-focused methods, with reverse reasoning, starting at the answer and ending at the problem. These back-chaining methods, however, underperform symbolic approaches.

Symbolic Methods Acknowledging that LLMs are poor proof-planners, a series of methods, including F-COT (Lyu et al., 2023) and SAT-LM (Ye et al., 2023), proposed to offload the reasoning burden from the LLM to more specialized tools. In these works, the LLM converts the text to symbolic logic, and a solver is then employed. Logic-LM (Pan et al., 2023) extended this to include a self-refinement step. While these methods perform well on simple datasets, they fail to account for ambiguity and the exclusion of common knowledge. Addressing this, Liu et al. (2024) and Wang et al. (2022) proposed algorithms that produce new clauses via logical deduction and then add the logic back to the text for an LLM to solve. While this might help the LLM, it does not add information to the problem, because any added relations are already deducible. Instead of producing clauses via deduction, Toroghi et al. (2024) proposed a method that exhaustively searches for new single-variable modus-ponens clauses. However, the search is conducted only over variables from the question, and repeated until the problem is solvable by classical logic, diminishing robustness. This search space is highly restricted and leaves out nearly all necessary information for some logic problems.

#### 3 Background

**Propositional logic** is a logical system that involves propositions about variables. A variable is a statement of fact, such as A="It is sunny" or B="I need an umbrella", which can be true or false. A literal is either a variable A or the negation  $\neg A$  of a variable. A proposition, such as  $A \to \neg B$ , is some statement about literals tied together by logical connectives (such as  $\land$ ,  $\lor$  or  $\rightarrow$ ). Like variables, literals and propositions can be true or false.

A **deductive propositional logic problem** is composed of a set of propositions, called the premises  $\mathcal{P} = \{P_1, \dots, P_K\}$ , which are given to be true (i.e.  $\vdash \mathcal{P}$ ), and a proposition Q, called the query. The goal of the problem is to determine whether the premises imply the query,  $\mathcal{P} \vdash Q$ , or its negation,  $\mathcal{P} \vdash \neg Q$ . Such problems are usually solved by translating them into two Boolean Satisfiability (SAT) problems, one for Q and one for  $\neg Q$ . The **backbone** of a propositional logic problem is the collection of all literals that are entailed by the premises, backbone( $\mathcal{P}$ ) = {literal  $L \mid \mathcal{P} \vdash L$ }. In effect, they are values for the variables which can be inferred from the propositions.

In an **abductive commonsense propositional logic problem** the premises  $\mathcal P$  entail neither the query Q nor its negation  $\neg Q$ : the problem is underdetermined. Instead, one must augment the premises with additional commonsense propositions  $\mathcal C$ , which represent background facts or knowledge left unstated in the problem, until either  $(\mathcal P \wedge \mathcal C) \vdash Q$  or  $(\mathcal P \wedge \mathcal C) \vdash \neg Q$ . Thus, the goal of an abductive problem is to not only find the truth-value of Q, but also a corresponding set of commonsense propositions  $\mathcal C$  to complete the problem. We assume that  $\neg(\mathcal P \wedge \mathcal C \to \bot)$ , that is that the premises  $\mathcal P$  are not contradictory with commonsense. One can show that, under this assumption, the answer to the problem will not depend on the choice of commonsense set  $\mathcal C$ : details are provided in Appendix A.

In practice, the problems we encounter in real life are often stated in terms of first-order logic. **First-order logic** is a logical system that involves statements about predicates. A predicate is a function,

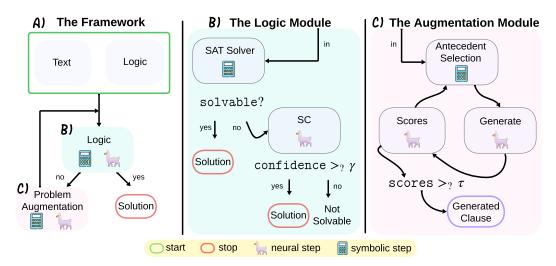


Figure 3: ARGOS at a glance. See Section 5.1 and Appendix D for details. (A) Given a propositional logic problem, we iteratively augment the problem with new propositions until it is solvable. (B) We attempt to solve the problem both with a logic solver, and with self-consistency (Wang et al., 2023). (C) If we fail, we attempt to add additional commonsense propositions by combining literals from the backbone as antecedents, and generating a right-hand-side using an LLM. We test the proposition for commonsense and relevance using this same LLM, and add it to the pool if it passes the tests.

such as MotherOf(x,y), which becomes true or false when entities, such as Alice and Bob, are substituted in its variables x and y. A statement, such as  $\forall (x,y)MotherOf(x,y) \rightarrow \neg Male(x)$ , is composed of predicates tied together using logical connectives (such as  $\land$ ,  $\lor$ ,  $\neg$  or  $\rightarrow$ ) and quantifiers ( $\exists$  and  $\forall$ ). Just like predicates, statements can also be true or false.

First-order logic problems over a finite set of entities can always be unrolled into propositional logic, by declaring a variable F(A) for every predicate F(x) and entity A, and expanding  $\forall x \, F(x)$  into the proposition  $(F(A) \land F(B) \land \ldots)$  and  $\exists x \, F(x)$  into  $(F(A) \lor F(B) \lor \ldots)$ . Given two propositional literals stemming from such an unrolling, we will declare them **related in first-order logic** if they have an entity in common. For example, MotherOf(Alice, Bob) and  $\neg Male(Alice)$  are related because both involve the entity Alice.

#### 4 PROBLEM STATEMENT

We are given an abductive propositional logic problem in both textual and logical form, as defined in Section 3, and we are also provided with a large language model and a SAT solver. As described, the task is to determine whether the target query is true or false given the premises and some additional commonsense propositions which must be found. Four annotated examples are provided, intended for few-shot prompting. In particular, the task is inference-only and no training phase is involved. We evaluate performance based on the number of correctly answered questions on a test dataset.

## 5 METHODOLOGY

We now describe our novel algorithm to tackle the problem described in Section 4. This algorithm is described by the diagram in Figure 3, and formally as Algorithm 1 in Appendix C.

#### 5.1 ALGORITHM

We start the algorithm by initializing our set of commonsense propositions as empty,  $C = \{\}$ . As shown in module B of Figure 3, we first try to solve the problem using the SAT Solver (sat\_solve) to test whether  $(P \land C) = P \vdash Q$  or  $\neg Q$ . If it reaches one of these conclusions, our job is finished; if not, we at least obtain from our call the backbone  $\mathcal{B} = \{\text{literal } L \mid P \vdash L\}$ .

Next, still in module B of Figure 3, we attempt to solve the problem using the LLM ( $11m\_solve$ ) by k-shot self-consistency (we use k=5 in our experiments). We ask the LLM whether the query is true or false, providing it the premises and the commonsense found so far. Details can be found in Appendix D.1. If the fraction of votes pass a certain threshold  $\gamma$ , we also conclude ( $\mathcal{P} \wedge \mathcal{C}$ ) =  $\mathcal{P} \vdash Q$  or  $\neg Q$  respectively, and the algorithm is finished. This parameter  $\gamma$  (initialized at  $\gamma=1$  in our experiments) is reduced by a fixed amount  $\gamma \leftarrow \gamma - \alpha$  at every iteration ( $\alpha=0.1$  in our experiments).

If neither solving method succeeds in establishing Q or  $\neg Q$ , we try to add a new commonsense proposition to our pool  $\mathcal C$ , as illustrated in module C of Figure 3. In practice, we define a proposition to be commonsense if it seems true to a large language model without any context. To guarantee that the added proposition will grow the problem's backbone, we search for commonsense propositions of the form  $L_1 \wedge L_2 \to L_{\text{right}}$ , where  $L_1$  and  $L_2$  are literals in the backbone  $\mathcal B$ , and  $L_{\text{right}}$  is a new literal suggested by the LLM. This search routine (find\_new\_commonsense) is described in Algorithm 2 in the Appendix. In detail, we start by iterating over pairs of literals in the backbone. We iterate by prioritizing the literals that share the most entities with others in the backbone,  $\operatorname{rank}_{\mathcal B}(L) = \#\{L' \in \mathcal B \mid L' \text{ has an entity in common with } L\}$ , which gives a measure of relevance of the literal to the problem. Next, for a given pair of literals  $L_1, L_2$ , we prompt the LLM (llm\_generate) to generate a right-hand-side literal  $L_{\text{right}}$  for  $L_1 \wedge L_2 \to L_{\text{right}}$ . In doing so, the LLM might introduce new variables not previously involved in the problem. Details can be found in Appendix D.2.

Finally, for each generated  $L_{\text{right}}$ , we use the LLM (llm\_score) twice to evaluate it. First, we use the LLM (llm\_commonsense\_score) to score whether  $L_1 \wedge L_2 \to L_{\text{right}}$  is likely to be commonsense. Second, we use the LLM again (llm\_relevance\_score) to score whether  $L_1 \wedge L_2 \to L_{\text{right}}$  is likely to be relevant to our current context. Each procedure returns a probability between 0 and 1. Details can be found in Appendices D.3 and D.4, respectively.

We stop the search at the first new proposition  $L_1 \wedge L_2 \to L_{\text{right}}$  whose commonsense and relevance scores are both above a given threshold  $\tau$  (we use  $\tau = 0.3$  in our experiments). When this happens, we update the commonsense set  $\mathcal C$  with this new proposition, and restart the process. If not, running new iterations will not change anything and we fall back on our best guess, namely the self-consistency estimate. In addition, if after multiple iterations the self-consistency threshold reaches zero, we also exit with the self-consistency estimate.

## 5.2 Example

Consider again the winter fox problem from the introduction section. Let us describe in Figure 4 a hypothetical run of our ARGOS algorithm to illustrate how it could solve the problem. To simplify the illustration, let us use only the SAT solver, and not self-consistency. We start with the premises (in black) and the query (in purple) on the top left-hand-side.

We first run the logic solver, which fails to reach any conclusion, but returns an initial backbone. The algorithm chooses the antecedents  $L_1 = L_2 = turns\_white(fox, winter)$  from this backbone, generating a new proposition  $turns\_white(fox, winter) \rightarrow reflects(fox, sun)$ . It is commonsensical and relevant to the question, so we add it to the question. We call the SAT solver again, which adds reflects(fox, sun) to the backbone. Next, the algorithm selects the antecedent  $L_1 = L_2 = reflects(fox, sun)$  from the new backbone and generates  $reflects(fox, sun) \rightarrow \neg absorbs(fox, sun)$ , which is similarly commonsensical and relevant. The SAT solver is called again and adds  $\neg absorbs(fox, sun)$  to the backbone. Finally, in the third iteration ARGOS picks  $L_1 = \neg absorbs(fox, sun)$  and  $L_2 = turns\_white(fox, winter)$  from the backbone and generates  $\neg absorbs(fox, sun) \wedge turns\_white(fox, winter) \rightarrow \neg absorbs(white, sun)$ , which is a logical conclusion it deems consistent with commonsense and relevant to the question. At this point, we call the SAT solver again, which concludes that  $\neg absorbs(white, sun)$  is true and therefore that the query must be false, returning  $\mathcal{P} \cup \mathcal{C} \vdash \neg Q$  as conclusion.

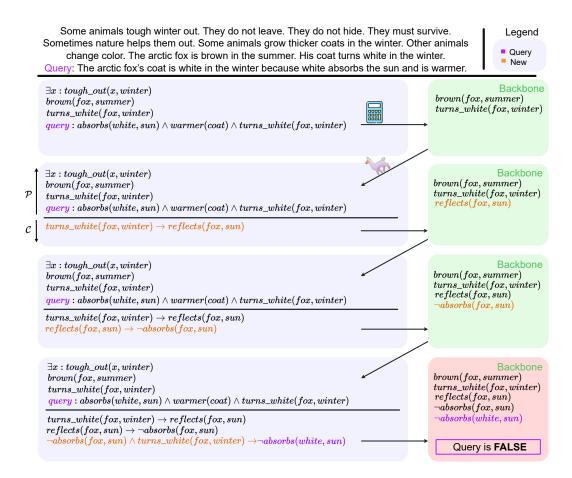


Figure 4: Overview of ARGOS with the winter fox example. We iteratively add to the logic problem and query a logic solver to look for conflicts within the backbone compared to the query. Eventually, we find that absorbs(white, sun) is False, contradicting the query.

## 6 EXPERIMENTS

**Models** We employ Llama3-8B (L8B), Llama3-70B (L70B), and Mistral 7B (M7B) as LLMs. Our method is dependent on access to logit-level outputs, so closed-source models are excluded. <sup>2</sup>

**Benchmarks** Unfortunately, there are few natural language reasoning datasets that are strongly logically-structured and commonsense-abductive. However, given a dataset of classical commonsense-based logic problems, data transformations to introduce the need for abductions are typically achievable. For a list of common datasets which have proven unsuitable for our setting, and corresponding explanation, see Appendix G. For our experiments, we use abductive versions of ProntoQA (Saparov and He, 2023), CLUTRR (Sinha et al., 2019), and FOLIO (Han et al., 2024). CLUTRR is not originally True/False, but it is multiple-choice. We modify it to be True/False output by making the question randomly either ask if the correct or an incorrect choice is True. To test our method's generalizability, we also include some datasets that are not strictly logical. CosmosQA (Huang et al., 2019) and QUAIL (Rogers et al., 2020) are reading comprehension MCQA datasets. ESNLI (Camburu et al., 2018) is a short-form natural-language-inference dataset. Each of these datasets requires some form of reasoning, but the structure of both the text and the necessary reasoning is generally fuzzy, requiring subjective interpretation. For the MCQA datasets, we process them into True/False questions similarly to how it was done for CLUTRR. We note that ProntoQA, CosmosQA and ESNLI performances are already saturated by self-consistency. Despite this, the results are valuable as they demonstrate that on these apparently simple tasks ARGOS is able to compare with

<sup>&</sup>lt;sup>2</sup>Experiments are each conducted on 1 or 2 NVIDIA Tesla V100 GPUs, depending on the LLM's GPU memory requirement.

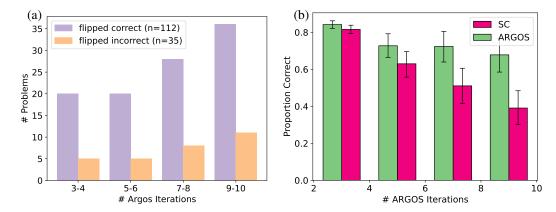


Figure 5: (a) The number of CLUTRR problems for which ARGOS flips SC predictions correctly and incorrectly. (b) SC and ARGOS accuracies on CLUTRR subsets, partitioned by the number of ARGOS iterations each datapoint receives.

purely neural methods, avoiding the performance collapse that more symbolic methods encounter. For more details, please see Appendix F.

**Evaluation** We compare against *COT* (Wei et al., 2022), *Self-Consistency* Wang et al. (2023), *SAT-LM* (Ye et al., 2023), *Logic-of-Thoughts* (Liu et al., 2024) and *LLM-Tres* (Toroghi et al., 2024). For a fair comparison with 20-shot self-consistency and LOT, we set ARGOS's hyperparameters such that it makes no more than 20 COT calls per problem on average. Details are provided in Appendix B. We report accuracy on the abduction-modified evaluation sets and report results in Table 1.

#### 6.1 RESULTS AND DISCUSSION

As can be seen in Table 1, ARGOS provides significant performance improvements over existing methods (up to +13%). Of the datasets, FOLIO is the most representative of human-generated logical reasoning problems. ARGOS outperforms the baselines for FOLIO, improving performance by 3-10%. For more structured problems (CLUTRR), the symbolic components of ARGOS become more reliable, and we see more consistent gains of 6-8%. On QUAIL, a highly ambiguous dataset that is also formatted in strange ways due to it being constructed by scraping forums and wikis, ARGOS improves compared to self-consistency by up to 13%, demonstrating its ability to adapt to even non-logical contexts. On ProntoQA, ESNLI and CosmosQA, despite the very competitive neural baseline performances, ARGOS performs comparably. Symbolic baselines (SAT-LM, LoT-20, LLM-Tres) see large performance gaps, at times being reduced to guessing.

**RQ1:** How useful are the scoring and backbone-tracking elements? In Table 2, we test the importance of two elements of ARGOS: (i) score thresholding and (ii) backbone computation. The ablation of each element in isolation results in a decrease in performance. In addition, the ablation of both results in a larger performance drop than even the sum of the two single ablations' decreases. The fully-ablated method, however, still shows strong performance relative to the next strongest baseline (SC-20), highlighting the strength of the general concept behind the method.

**RQ2:** Are ARGOS' added clauses corruptive? An important criterion when adding clauses is that they do not corrupt the logic of the problem, undesirably changing the outcome of the logic. It can be shown (see Appendix A) that so long as clauses are commonsensical, their addition will not corrupt the problem. However, it is possible that our method adds non-commonsensical clauses, since the commonsense scoring is not perfectly reliable. Given CLUTRR's strict structure, since we know the full knowledge base from which it was constructed, we can re-construct the full problems and test if ARGOS' added clauses corrupt the logical arithmetic such that a different answer is found for the logic problem. We find that on CLUTRR, ARGOS never corrupts a problem. It is then not surprising that ARGOS sees significant performance gains: added information should in principle never negatively affect a wholly rational reasoner's solution and so performance should only improve.

Table 1: Binary classification accuracy (True/False) of all methods on the datasets, using the chosen language models. Bolded text indicates that the method has the best performance, and that its performance is better than the next-best-performing method in a statistically significant way (*p*-value < 0.005 according to a Wilcoxon pair-wise rank test). Small-font numbers to the right indicate the bounds of the 95% confidence interval, derived via a bootstrap approach.

	FOLIO			CLUTRR			PQA		
	M7B	L8B	L70B	M7B	L8B	L70B	M7B	L8B	L70B
SC20	66% <sup>66.4</sup> <sub>65.5</sub>	$71\%_{70.1}^{71.7}$	77% <sup>77.7</sup> <sub>75.9</sub>	59% <sup>59.3</sup> <sub>58.8</sub>	69% <sup>69.5</sup> <sub>68.8</sub>	69% <sup>69.4</sup> <sub>68.8</sub>	97% <sup>97.2</sup> <sub>95.6</sub>	95% <sup>95.6</sup> <sub>94.4</sub>	93% 94.1 92.4
COT	66% <sup>66.4</sup> <sub>65.5</sub>	$68\%  {}^{69.1}_{67.2}$	$72\%  {}^{72.5}_{71.8}$	59% <sup>59.3</sup> <sub>58.8</sub>	$68\%  {}^{68.4}_{67.8}$	$66\% \frac{66.3}{65.6}$	82% <sup>82.9</sup> <sub>81.7</sub>	$90\% \frac{91.2}{89.6}$	$93\% \frac{94.1}{92.4}$
SAT-LM	$43\%  {}^{43.2}_{42.8}$	$43\%  {}^{43.2}_{42.8}$	$43\%  {}^{43.2}_{42.8}$	50% 50.4 49.9	$50\%  {}^{50.4}_{49.9}$	$50\%  {}^{50.4}_{49.9}$	50% 50.3 49.8	$50\%_{\ 49.8}^{\ 50.3}$	$50\%  {}^{50.3}_{49.8}$
LoT-20	57% <sup>57.3</sup> <sub>56.6</sub>	$69\% \frac{69.5}{68.7}$	$70\% \frac{70.4}{69.5}$	$71\%_{70.7}^{71.6}$	$70\%_{\ 69.7}^{\ 70.2}$	$69\% \frac{69.3}{68.7}$	88% <sup>88.4</sup> <sub>87.5</sub>	$97\% \frac{98.2}{96.3}$	$95\%  {}^{95.7}_{94.3}$
LLM-Tres	66% <sup>66.6</sup> <sub>65.9</sub>	$63\% \frac{63.2}{62.4}$	$63\% \frac{63.2}{62.4}$	51% <sup>51.5</sup> <sub>50.8</sub>	$51\% \frac{51.6}{50.8}$	$53\% \frac{53.2}{52.8}$	80% 81.4	$83\% \frac{83.8}{82.3}$	$76\%  {}^{76.6}_{75.2}$
ARGOS	$70\% \frac{70.6}{69.8}$	$81\%{}^{81.8}_{80.0}$	$80\% \frac{80.5}{78.8}$	78% 78.4	$76\%  {}^{76.3}_{75.8}$	$78\%  {}^{78.2}_{77.7}$	98% <sup>98.7</sup> <sub>97.9</sub>	$97\% \frac{98.2}{96.3}$	$97\% \ {}^{98.1}_{96.2}$
				1					
		CosmosQA	A		ESNLI			QUAIL	
	M7B	CosmosQA 88 7	L70B	M7B	L8B	L70B	M7B	QUAIL 88 7	L70B
SC20	<u> </u>			97% 97.7 97.1			8 W 70% <sup>71.0</sup> 67.9		80 02 75% 75.6 72.4
SC20 COT	M7B	L8B	L70B		L8B	99% 99.5		L8B	
	84% 84.3 82.9	81% 81.3 81% 79.7	90% 91.1 90% 89.9	97% <sup>97.7</sup> <sub>97.1</sub>	96% 97.0 96% 96.3	99% 99.5	70% 71.0 67.9	8% 68% 69.1 65.6	75% 75.6 72.4
COT	84% 84.3 84% 82.9 81% 81.4 80.1	81% 81.3 79.7 76% 77.5 74.2	90% 91.1 90% 89.9 88% 88.7 86.9	97% 97.7 97.1 96% 95.8 96.4	96% 97.0 96% 96.3 88% 86.9 88.4	99% 99.5 99.2 99% 98.9 99.4	$70\% \frac{71.0}{67.9}$ $71\% \frac{71.9}{68.8}$	68% <sup>69.1</sup> 65% <sup>66.6</sup> 65% <sup>66.2</sup> 63.6	75% 75.6 72.4 75% 75.6 72.4
COT SAT-LM	84% 84.3 82.9 81% 81.4 80.1 35% 37.5 34.8	81% 81.3 79.7 76% 77.5 74.2 35% 37.5 34.8	90% 91.1 90% 89.9 88% 88.7 86.9 35% 37.5 34.8	97% 97.7 97.1 96% 95.8 96.4 49% 50.0 47.7	96% 97.0 96% 96.3 88% 86.9 88.4 49% 50.0 47.7	99% 99.5 99% 98.9 99% 50.0 49% 50.0 75% 76.1 75% 74.4	70% <sup>71.0</sup> 67.9 71% <sup>71.9</sup> 68.8 53% <sup>55.0</sup> 51.6	68% 69.1 65% 65.6 65% 63.6 53% 55.0 51.6	75% 75.6 72.4 75% 75.6 72.4 53% 55.0 51.6

**RQ3:** Does ARGOS attribute more compute to harder problems? How does this affect the solution of harder problems? In Figure 5 (b), we examine the proportion of CLUTRR problems that are solved correctly by SC and ARGOS, over subsets of the dataset grouped by the number of ARGO iterations before termination. The error bars are 5/95% confidence intervals. As the number of ARGOS iterations increases, the problems become harder for SC to solve (indicated by a lower proportion of correct solutions by SC). This tells us that ARGOS' method of evaluating solvability is working as intended; harder problems are being assigned more computation. Another interpretation of this result is that problems which have more missing information, or for which the missing information is more difficult to infer, are attributed more ARGOS iterations (in order for ARGOS to find the necessary information). This is supported by the fact that the decrease in proportion seen in SC is not present for ARGOS: if SC's performance is dropping due to missing information, then ARGOS is successfully recovering the necessary missing information.

This ability to address the obstacles which cause SC performance to drop contribute to a large number of answers being flipped from incorrect (when solved by self-consistency) to correct (when solved by ARGOS). Changes to the answers caused by new information are more often than not in the right direction. On CLUTRR L70B, we find 112 correct and 35 incorrect flips. Figure 5 (a) shows the number of correct and incorrect flips ARGOS achieves. As the number of ARGOS iterations increases, both the correct and incorrect flip counts increase, but the correct flip counts increase much faster. For a closer look at confidence-score vs. iteration behavior, see Appendix I.

Figure 6 provides an example of a question from CLUTRR that is misclassified by self-consistency but flipped to correct by ARGOS. The COT seems confused, displaying its characteristic inability to plan out a proof: in steps 1-3 it provides disjoint pieces of information that neither follow from each

Table 2: Ablations. We ablate elements of ARGOS: (i) the score thresholding, taking the first clause sampled at each iteration (ARGOS - No T), (ii) the backbone-tracking, generating prompts by randomly selecting two variables (ARGOS - No BB).

	FOLIO L8B
SC-20	$71\% \frac{71.7}{70.1}$
ARGOS - No T	79% <sup>79.8</sup> <sub>78.5</sub>
ARGOS - No BB	$79\% \frac{79.4}{78.4}$
ARGOS - No Both	76% <sup>75.2%</sup> <sub>77.2%</sub>
ARGOS	$81\% \ {}^{81.8}_{80.0}$

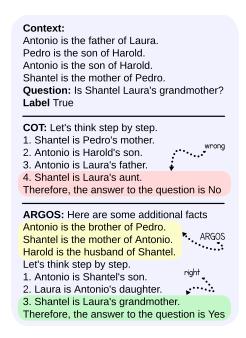


Figure 6: COT vs ARGOS on a CLUTRR problem.

other nor move towards the target conclusion. This confusion eventually leads to an incorrect step: "Shantel is Laura's aunt", resulting to an incorrect conclusion. ARGOS, after 3 iterations, provides several pieces of key information which would require at least one additional reasoning step to find, halving the necessary chain-length. For some examples in which ARGOS fails, see Appendix H.

## 6.2 IMPACT OF IMPERFECT LOGICAL TRANSLATION

In our experiments, we assumed that we started from a propositional logic formulation. Some datasets came with an official formulation, while for the others we translated from text using Claude Opus 4, filtering to remove failed translations. This was done to fairly evaluate the methods on abductive reasoning, regardless of the quality of translation. In general, logical translation is an orthogonal problem and intrinsically simpler for LLMs, since it is linguistic rather than cognitive. LLMs have already demonstrated strong abilities at logic translation (Yang et al., 2024), and are expected to continue improving faster than at reasoning. To validate this claim in our context, we re-tested ARGOS with Llama 8B on FOLIO using a translator, but including failed translations. Performance only decreased marginally, from 80% to 78%, still outperforming the next best method (SC at 71%).

## 7 Conclusion

We have presented a method for addressing realistic natural-language logic problems, where "realistic" entails a need for abduction and commonsense. Whether neural or symbolic, we demonstrate empirically that existing methods struggle in this setting. The method we present addresses this weakness by (a) balancing neural and symbolic elements and allowing them to speak to each-other; and (b) avoiding the commonplace design choice of heavily restricting the abduction-clause search space. On both general and highly structured logic problems, our method demonstrates the power of a balanced neuro-symbolic approach, outperforming all existing work meaningfully.

**Limitations** A limitation of our work is that it is currently restricted to problems which are strictly True or False, eliminating cases where logic might be used to select an option from a list of choices, or cases where the correct answer is "Maybe". In our experimental work, we addressed the consequences this had on dataset selection by converting datasets to be True/False. The method could however be extended to multiple-choice questions by asking each question as an individual True/False question, combined with a decision heuristic for when no/multiple choices are determined True.

## REPRODUCIBILITY STATEMENT

In the supplementary material, we provide our full code which was used to implement and benchmark our method as well as the baselines. The code also includes data processing steps. We took great care to include in the Appendix, as well, detailed descriptions of our algorithm and our prompts. While our human modification of FOLIO text is not provided, the process for generating it is described carefully in the Appendix.

## REFERENCES

- Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. e-SNLI: Natural Language Inference with Natural Language Explanations. In *Proc. Conf. Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), Vol. 31. Curran Associates, Inc.
- Peter Clark, Oyvind Tafjord, and Kyle Richardson. 2021. Transformers as soft reasoners over language. In *Proc. Int. Joint Conf. on Artificial Intelligence*. 3882–3890.
- Daniel Crevier. 1993. *AI: The Tumultuous Search for Artificial Intelligence*. Basic Books, Chapter 6: The Tree of Knowledge.
- Nouha Dziri, Ximing Lu, Melanie Sclar, Xiang (Lorraine) Li, Liwei Jiang, Bill Yuchen Lin, Sean Welleck, Peter West, Chandra Bhagavatula, Ronan Le Bras, Jena Hwang, Soumya Sanyal, Xiang Ren, Allyson Ettinger, Zaid Harchaoui, and Yejin Choi. 2023. Faith and Fate: Limits of Transformers on Compositionality. In *Proc. Conf. Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Eds.), Vol. 36. Curran Associates, Inc., 70293–70332.
- Simeng Han, Hailey Schoelkopf, Yilun Zhao, Zhenting 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, Ansong Ni, Linyong Nan, Jungo Kasai, Tao Yu, Rui Zhang, Alexander Fabbri, Wojciech Maciej Kryscinski, Semih Yavuz, Ye Liu, Xi Victoria Lin, Shafiq Joty, Yingbo Zhou, Caiming Xiong, Rex Ying, Arman Cohan, and Dragomir Radev. 2024. FOLIO: Natural Language Reasoning with First-Order Logic. In *Proc Conf. on Empirical Methods in Natural Language Processing*. 22017–22031.
- Lifu Huang, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. Cosmos QA: Machine Reading Comprehension with Contextual Commonsense Reasoning. In *Proc. Conf. Empirical Methods in Natural Language Processing and Int. Joint Conf. on Natural Language Processing (EMNLP-IJCNLP)*. Hong Kong, China, 2391–2401.
- Mehran Kazemi, Najoung Kim, Deepti Bhatia, Xin Xu, and Deepak Ramachandran. 2023. LAM-BADA: Backward Chaining for Automated Reasoning in Natural Language. In *Proc. Conf. Association for Computational Linguistics*. 6547–6568.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. In *Proc. Conf. Neural Informations Processing Systems*. 22199–22213.
- Jinu Lee and Wonseok Hwang. 2024. SymBa: Symbolic Backward Chaining for Structured Natural Language ReasoningSymBa: Symbolic Backward Chaining for Structured Natural Language Reasoning. arXiv:2402.12806 [cs.CL]
- Jian Liu, Leyang Cui, Hanmeng Liu, Dandan Huang, Yile Wang, and Yue Zhang. 2021. LogiQA: a challenge dataset for machine reading comprehension with logical reasoning. In *Proc. Int. Joint Conf. on Artificial Intelligence*. 3622–3628.
- Tongxuan Liu, Wenjiang Xu, Weizhe Huang, Yuting Zeng, Jiaxing Wang, Xingyu Wang, Hailong Yang, and Jing Li. 2024. Logic-of-thought: Injecting logic into contexts for full reasoning in large language models. In *arXiv preprint arXiv:2409.17539*.

Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. 2023. Faithful Chain-of-Thought Reasoning. In *Int. Joint Conf. Natural Language Processing and Conf. Asia-Pacific Chapter of the Association for Computational Linguistics*.

- Nils J. Nilsson. 1991. Logic and artificial intelligence. Artificial Intelligence 47, 1 (1991), 31–56.
- Theo Olausson, Alex Gu, Ben Lipkin, Cedegao Zhang, Armando Solar-Lezama, Joshua Tenenbaum, and Roger Levy. 2023. LINC: A Neurosymbolic Approach for Logical Reasoning by Combining Language Models with First-Order Logic Provers. In *Proc. Conf. on Empirical Methods in Natural Language Processing*. 5153–5176.
- Liangming Pan, Alon Albalak, Xinyi Wang, and William Wang. 2023. Logic-LM: Empowering Large Language Models with Symbolic Solvers for Faithful Logical Reasoning. In *Findings of the Association for Computational Linguistics*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). 3806–3824.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language Models as Knowledge Bases?. 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)*, Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (Eds.). Association for Computational Linguistics, Hong Kong, China, 2463–2473.
- Aske Plaat, Annie Wong, Suzan Verberne, Joost Broekens, Niki van Stein, and Thomas Back. 2024. Reasoning with large language models, a survey. In *arXiv preprint arXiv:2407.11511*.
- Anna Rogers, Olga Kovaleva, Matthew Downey, and Anna Rumshisky. 2020. Getting closer to AI complete question answering: A set of prerequisite real tasks. In *Proc. AAAI Conf. Artificial Intelligence*, Vol. 34. 8722–8731.
- Abulhair Saparov and He He. 2023. Language Models Are Greedy Reasoners: A Systematic Formal Analysis of Chain-of-Thought. In *Proc. Int. Conf. Learning Representations*.
- Koustuv Sinha, Shagun Sodhani, Jin Dong, Joelle Pineau, and William L. Hamilton. 2019. CLUTRR: A Diagnostic Benchmark for Inductive Reasoning from Text. In *Proc. Conf. Empirical Methods of Natural Language Processing*.
- Oyvind Tafjord, Bhavana Dalvi, and Peter Clark. 2021. ProofWriter: Generating Implications, Proofs, and Abductive Statements over Natural Language. In *Proc. Conf. Association for Computational Linguistics: ACL-IJCNLP*. 3621–3634.
- Jidong Tian, Yitian Li, Wenqing Chen, Liqiang Xiao, Hao He, and Yaohui Jin. 2021. Diagnosing the first-order logical reasoning ability through LogicNLI. In *Proc. Conf. Empirical Methods in Natural Language Processing*. 3738–3747.
- Armin Toroghi, Willis Guo, Ali Pesaranghader, and Scott Sanner. 2024. Verifiable, Debuggable, and Repairable Commonsense Logical Reasoning via LLM-based Theory Resolution. In *Proc. Conf. on Empirical Methods in Natural Language Processing*. 6634–6652.
- Siyuan Wang, Wanjun Zhong, Duyu Tang, Zhongyu Wei, Zhihao Fan, Daxin Jiang, Ming Zhou, and Nan Duan. 2022. Logic-Driven Context Extension and Data Augmentation for Logical Reasoning of Text. In *Proc. Conf. Association for Computational Linguistics*, Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (Eds.). 1619–1629.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. Self-Consistency Improves Chain of Thought Reasoning in Language Models. In *Proc. Int. Conf. on Learning Representations*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Proc. Conf. Neural Information Processing Systems*. 24824–24837.

- Jundong Xu, Hao Fei, Liangming Pan, Qian Liu, Mong-Li Lee, and Wynne Hsu. 2024. Faithful Logical Reasoning via Symbolic Chain-of-Thought. In Proc. Conf. Association for Computational Linguistics. 13326-13365. Shangzi Xue, Zhenya Huang, Jiayu Liu, Xin Lin, Yuting Ning, Binbin Jin, Xin Li, and Qi Liu. 2024. Decompose, analyze and rethink: Solving intricate problems with human-like reasoning cycle. In Proc. Conf. Neural Information Processing Systems. 357–385. Yuan Yang, Siheng Xiong, Ali Payani, Ehsan Shareghi, and Faramarz Fekri. 2024. Harnessing the Power of Large Language Models for Natural Language to First-Order Logic Translation. In Annual Meeting of the Association for Computational Linguistics 2024. Association for Computational Linguistics (ACL), 6942–6959.
  - Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. In *Proc. Conf. Neural Information Processing Systems*. 11809–11822.
  - Xi Ye, Qiaochu Chen, Isil Dillig, and Greg Durrett. 2023. SatLM: Satisfiability-Aided Language Models Using Declarative Prompting. In *Proc. Conf. Neural Information Processing Systems*.
  - Weihao Yu, Zihang Jiang, Yanfei Dong, and Jiashi Feng. 2020. ReClor: A Reading Comprehension Dataset Requiring Logical Reasoning. In *Proc. Int. Conf. on Learning Representations*.

#### A ABDUCTIVE LOGIC PROBLEMS ARE WELL-DEFINED

In this section we prove that the solution of an abductive propositional logic problem, given in Section 3, does not depend on the choice of commonsense set C.

**Proposition 1.** Let  $\mathcal{P}$  be a set of premises, and  $\mathcal{C}_1$  and  $\mathcal{C}_2$  sets of commonsense propositions, that is  $\mathcal{C}_1, \mathcal{C}_2 \subseteq \mathcal{C}$  such that  $\vdash \mathcal{P}, \mathcal{C}$ . Then the logical expressions  $\mathcal{P} \land \mathcal{C}_1$  and  $\mathcal{P} \land \mathcal{C}_2$  cannot be contradictory, i.e.  $\forall_{\mathcal{C}_1, \mathcal{C}_2 \subseteq \mathcal{C}} : \neg \vdash \neg (\mathcal{P} \land \mathcal{C}_1) \land (\mathcal{P} \land \mathcal{C}_2)$ .

*Proof.* Let us assume that  $\mathcal{P} \wedge S_1$  and  $\mathcal{P} \wedge S_2$  have contradictory outcomes, say

$$\vdash (\mathcal{P} \land S_1 \to q),$$
 (1)

$$\vdash (\mathcal{P} \land S_2 \to \neg q),$$
 (2)

where q is some logical expression. Since  $C_1, C_2 \subseteq C$ , and  $\vdash C$ , we have  $\vdash C_1$  and  $\vdash C_2$ , so

$$\vdash (\mathcal{C}_1 \leftrightarrow \mathcal{C}_2).$$
 (3)

Substituting  $C_1$  for  $C_2$  in Eq. 1 we have

$$\vdash (\mathcal{P} \land \mathcal{C}_2 \to q).$$
 (4)

Since  $\vdash \to A$  allows us to infer  $\vdash A$ , and since we have  $\vdash (\mathcal{P} \land \mathcal{C}_2)$ , Eq. 2, 4 allow us to infer both  $\vdash q$  and  $\vdash \neg q$ , which is impossible so long as  $q \neq \emptyset$  and thus shows by contradiction the impossibility of the initial conditions Eq. 1, 2. Therefore, we can say that if  $\mathcal{C}_1, \mathcal{C}_2 \subseteq \mathcal{C}, \mathcal{P} \land \mathcal{C}_2$  can not be contradictory

$$\forall_{\mathcal{C}_1,\mathcal{C}_2 \subset \mathcal{C}} : \neg \vdash \neg(\mathcal{P} \land \mathcal{C}_1) \land (\mathcal{P} \land \mathcal{C}_2)$$
 (5)

#### B Cost Discussion

While in theory COT generation is meant to be done until an answer is found, in practice it is necessary that an upper-limit on number of tokens generated is enforced. This is in case (a) the LLM continues generating past its answer, or (b) the LLM goes off-track and never answers the question. In any case, this means that each COT generation will be, at worst-case-assumption, equal in cost. In addition, the various method-specific LLM generations that are employed require small token-limits relative to COT, and so the number of COT calls made dominates the total number of tokens generated by any method. Given this, budgeting method costs in terms of COT calls is well-justified. For SC and COT, the cost evaluation is trivial: COT always makes 1 COT call and SC makes a fixed number of COT calls, specified as a hyper-parameter. Similarly, LOT makes some small generative calls followed by SC, so its number of COT calls is fixable. LLM-Tres makes no COT calls, and neither does SAT-LM. ARGOS' cost varies according to the entry, but its hyper-parameters (number of COT calls per-iteration and threshold/annealing constants) can be set such that its average number (or worst-case) number of calls is less than a budget. A summary of method cost in terms of COT calls is provided in Table 3.

Table 3: Average number of COT calls required by each method.

	Cost (Avg # COT)
COT	1
SC	20
LOT	20
SAT-LM	0
LLM-Tres	0
ARGOS	18.4

## C ALGORITHM DESCRIPTION

In this section we provide a detailed description of our ARGOS algorithm. The main procedure is summarized as Algorithm 1, which uses the find\_new\_commonsense subroutine in Algorithm 2.

## Algorithm 1 ARGOS

702

703 704

705

706

708

735 736

737

738 739

740

741

742

743

744

745

746

747

748

749

750

751

752753754

```
709
           Require: premises \mathcal{P}, query Q, SC sample-count k, scoring threshold \tau \in (0,1],
710
                self-consistency threshold \gamma \in (0,1] and decay \alpha \in (0,1]
711
            1: commonsense set \mathcal{C} \leftarrow \{\}
712
            2: while \gamma > 0 do
                     // Attempt solving with SAT solver
713
            3:
            4:
                     sat_conclusion, backbone \mathcal{B} \leftarrow \text{sat\_solve}(\mathcal{P} \land \mathcal{C} \vdash Q, \neg Q)
714
                     if sat_conclusion is (\mathcal{P} \wedge \mathcal{C}) \vdash Q or (\mathcal{P} \wedge \mathcal{C}) \vdash \neg Q then
            5:
715
            6:
                          return C, sat_conclusion
716
            7:
                     end if
717
                     // Else attempt solving with the LLM (k-shot self-consistency)
            8:
718
            9:
                     llm\_conclusion, llm\_confidence \leftarrow llm\_solve(\mathcal{P} \land \mathcal{C} \vdash Q, \neg Q)
719
           10:
                     if llm_confidence > \gamma then
720
                          return C, llm_conclusion
           11:
721
           12:
                     end if
722
           13:
                     // Else find a new commonsense proposition to add to the pool
723
           14:
                     C \leftarrow \text{find\_new\_commonsense}(\mathcal{P}, \mathcal{C}, \mathcal{B}, \tau)
                     if C is not None then
724
           15:
                          // New commonsense has been found, we try again with an enlarged C and smaller \gamma
725
           16:
           17:
                          \mathcal{C} \leftarrow \mathcal{C} \land \{C\}
726
           18:
                          \gamma \leftarrow \gamma - \alpha
727
           19:
                     else
728
           20:
                          // We failed, return best guess
729
                          return C, LLM_conclusion
           21:
730
                     end if
           22:
731
           23: end while
732
           24: // We ran out of time, return best guess
733
           25: return C, LLM conclusion
734
```

## Algorithm 2 find\_new\_commonsense

```
Require: premises \mathcal{P}, commonsense \mathcal{C}, backbone \mathcal{B} = \text{backbone}(\mathcal{P} \wedge \mathcal{C}), scoring threshold \tau \in
 1: for L_1 \in \mathcal{B} from highest to lowest rank<sub>\mathcal{B}</sub>(L_1) do
 2:
          for L_2 \in \mathcal{B} from highest to lowest rank_{\mathcal{B}}(L_2) do
 3:
               for L_{\text{right}} in llm_generate_{\mathcal{P} \wedge \mathcal{C}}(L_1 \wedge L_2 \rightarrow ?) do
 4:
                     commonsense_score \leftarrow llm_commonsense_score(L_1 \land L_2 \rightarrow L_{\text{right}})
 5:
                     relevance_score \leftarrow llm_relevance_score_{\mathcal{P} \wedge \mathcal{C}}(L_1 \wedge L_2 \rightarrow L_{\text{right}})
 6:
                     if commonsense_score > \tau and relevance_score > \tau then
                          // We found a new relevant commonsense clause
 7:
 8:
                          return L_1 \wedge L_2 \rightarrow L_{\text{right}}
 9:
                     end if
10:
                end for
          end for
11:
12: end for
13: // We failed to find a new relevant commonsense clause
14: return None
```

## D ALGORITHM LLM DETAILS

In this section we provide details about the parts of the algorithm that involve interactions with the large language model.

## D.1 SOLVING WITH 5-ROUND SELF-CONSISTENCY (LLM\_SOLVE)

This subroutine aims to establish whether the premises and commonsense entails the query, i.e.  $(\mathcal{P} \wedge \mathcal{C}) \vdash Q$  or  $(\mathcal{P} \wedge \mathcal{C}) \vdash \neg Q$ . This routine always returns an answer, but can make mistakes. We few-shot prompt the LLM 5 times with the prompt in Table 4.

Table 4: COT prompt

```
Here are some facts and rules: [premises \mathcal{P}] Here is some additional info we found: [commonsense \mathcal{C}] True or false: [query Q]? Answer:
```

Each call returns an answer  $a_1, \ldots, a_5 \in \{\text{True}, \text{False}\}\$ , with a certain confidence  $c_1, \ldots, c_5 \in (0, 1]$ . The algorithm returns the most common answer (True or False), and a total confidence score given by the sum of confidences of the most common answer  $a^*$ , divided by 5:

$$c^* = \frac{1}{5} \sum_{i=1}^{5} c_i \mathbb{1} [a_i = a^*].$$

## D.2 GENERATING NEW COMMONSENSE LITERALS $L_{\text{RIGHT}}$ (LLM\_GENERATE)

This subroutine aims to find a plausible right-hand-side literal  $L_{\text{right}}$  for a proposition  $L_1 \wedge L_2 \to L_{\text{right}}$ . The new literal might potentially involve new variables not previously seen in the problem. We use a slightly different prompt for CLUTRR and for the others, because of CLUTRR's more distinct structure.

Table 5: Clause generation prompt for all datasets but CLUTRR. e is an entity appearing in  $L_1$  or  $L_2$ .

```
Fill in the blank with a known predicate: [L_1 \wedge L_2] implies ___([e]). Known predicates are: [all predicates appearing in the premises \mathcal P and the commonsense \mathcal C] Answer:
```

Table 6: Clause generation prompt for CLUTRR.  $e_1$  and  $e_2$  are entities appearing in  $L_1$  or  $L_2$ .

```
If [L_1 \wedge L_2] then ___([e_1], [e_2]). Fill in the blank. Answer:
```

Thus, for example, in FOLIO if  $L_1 = drinksCoffee(Rina)$  and  $L_2 = Loves(Mary, Sam)$ , then we would make three calls, one for  $drinksCoffee(Rina) \land Loves(Mary, Sam) \rightarrow F(Rina)$ ,  $drinksCoffee(Rina) \land Loves(Mary, Sam) \rightarrow F(Mary)$  and  $drinksCoffee(Rina) \land Loves(Mary, Sam) \rightarrow F(Sam)$  respectively. With such calls, the method might return the set  $\{productive(Rina), hasFeelings(Mary), isLoved(Sam)\}$ , for example.

## D.3 SCORING PROPOSITIONS $L_1 \wedge L_2 \rightarrow L_{\text{RIGHT}}$ FOR COMMONSENSE (LLM\_COMMONSENSE\_SCORE)

This procedure uses the LLM to score how much our new proposition  $L_1 \wedge L_2 \rightarrow L_{\text{right}}$  is likely to be commonsense. In detail, we ask the LLM whether, without any context, the

rule seems contradictory (FOLIO/ProntoQA) or true (CLUTRR). We record the logits of the "Yes" and "No" tokens following the prompt, and we return as commonsense score  $P[\text{No}] = \exp(\log i t_{\text{No}})/(\exp(\log i t_{\text{Yes}}) + \exp(\log i t_{\text{No}}))$ , except for CLUTRR where we return  $P[\text{Yes}] = \exp(\log i t_{\text{Yes}})/(\exp(\log i t_{\text{Yes}}) + \exp(\log i t_{\text{No}}))$  since the question is inverted.

Table 7: Commonsense scoring prompt for FOLIO/ProntoQA.

```
Does the following rule seem contradictory? Rule: [L_1 \wedge L_2 	o L_{\text{right}}] Answer:
```

Table 8: Commonsense scoring prompt for CLUTRR.

```
Does the following rule seem true? Rule: [L_1 \wedge L_2 	o L_{	ext{right}}] Answer:
```

# D.4 SCORING PROPOSITIONS $L_1 \wedge L_2 \rightarrow L_{\text{RIGHT}}$ FOR CONTEXT-RELEVANCE (LLM\_RELEVANCE\_SCORE)

This procedure uses the LLM to score how much our new proposition  $L_1 \wedge L_2 \to L_{\rm right}$  is likely to be relevant to the context. This helps eliminate propositions, like "The sky is blue", that are true and commonsense but unlikely to help prove our query. In this case, we use the same prompt for all datasets. We record the logits of the tokens "Yes" and "No" following the text, and return  $P[{\rm Yes}] = \exp(\log i t_{{\rm Yes}})/(\exp(\log i t_{{\rm Yes}}) + \exp(\log i t_{{\rm No}}))$  as relevance score.

Table 9: Context scoring prompt for all datasets.

```
Here are some facts and rules: [premises \mathcal P and commonsense \mathcal C] Does the following new rule seem contextually relevant to the facts and rules? [L_1 \wedge L_2 \to L_{\text{right}}] Answer:
```

### E IN-DEPTH LITERATURE REVIEW

Since 2022, when Wei et al. (2022) found that models had the capacity to solve (at the time) difficult reasoning problems by simply prompting the model to output a detailed rationale (a "chain of thought" (COT)) before making a decision, there has been significant interest in leveraging/improving LLMs' capacity for reasoning.

**Semantic Logic Parsers** More recently, a series of methods were proposed which effectively bypassed the reasoning process by prompting LLMs to translate the given input to symbolic language (parsing the logic) and then using external, programmatic solvers to solve the problem in the symbolic space. F-COT, proposed by Lyu et al. (2023) and SAT-LM, proposed by Ye et al. (2023) are two contemporaneous works which prompt the LLM to translate the given problem into its corresponding symbolic language, to be solved by the appropriate solver. Logic-LM, proposed by Pan et al. (2023), also published around the same time, includes a self-refinement step in order to catch mistranslations and to re-translate them, but this module provided only minor improvements.

The logic-parsing strategy proved extremely effective, converting reasoning tasks into the more linguistic translation task. Since the algorithmic tools never make mistakes, if the translation is accurate, then the solution will be too. These methods, however, rely upon the strong assumption which goes widely unacknowledged that all necessary information is provided at input. By motivation, this assumption implies a well-informed, precise, and careful end-user. This critical assumption greatly injures the applicability and generalizability of these symbolic methods. In order to address

this, ARGOS aims specifically to provide missing information by exploring the logical space and leveraging logical tools. Xu et al. (2024), argue that using external solvers is not relevant to the LLM's actual reasoning capacity and present a method which still parses the text into symbolic language, but uses the LLM as the symbolic solver by inputting the symbolic expressions directly to the LLM. In fact, this was an alternative presented in SAT-LM, and was shown to be weaker than directly using an external solver in cases where the problems are fully defined (which is the case in the setting chosen by both works).

**Deductive Logical Algorithms** Given the findings discussed above that LLMs were efficient logic-parsers, the door was opened to more agentic algorithms which operate in the logic-space (as opposed to the textual space). Algorithms were proposed by Kazemi et al. (2023) and Lee and Hwang (2024) which attempted to reason backwards through the problem in logic space, using the LLM after each deduction step to choose the next deduction. These methods proved to perform worse than the previous parsing methods as they relied on the same assumptions of completeness in the input, but also required the LLM to greedily trace a reasoning path, opening another avenue for errors. This step is unnecessary since logic tools are able to search exhaustively over the solution space at very low cost.

Clause-generative Algorithms (Abduction) In response to the literature's inability to address cases where the completeness assumption fails, some work was proposed which explicitly aims to generate logical clauses in the textual space. The very similar algorithms proposed by Liu et al. (2024) and Wang et al. (2022) first translate the text to logic and then produce some new clauses which are deducible via classical logic given the logical translation of the input. By then translating the newly deduced logic back to text, the textual representation of the logic is now richer from a linguistic point of view (although no new information was added to the underlying logic), and then COT is used to solve the augmented problem in text-space. In the logical sense, this method is not truly abductive as no new logical information is produced.

Instead of producing clauses via deduction on the input, Toroghi et al. (2024) proposed a method which leverages LLMs' knowledge of commonsense to supply missing clauses during the reasoning process. Given the logical translation of the input text, the space of all 2-variable clauses which are possible to construct using the variables given in the question is explored exhaustively, with each rule being given a probability of being true by the LLM. Reasoning paths are formed with the various rules.

This method suffers from a rigidity regarding the search space: due to the exhaustive search the space must be restricted to only 2-variable rules constructible from the input. Often the missing information from logic problems may include variables which are not named in the problem and so are unseen in the abduction input. In other cases, the necessary rules might be of different sizes. These cases are easily seen even in the commonly used datasets for logic-reasoning evaluation. For example, CLUTRR Sinha et al. (2019), which is a dataset of reasoning problems related to family relationships, requires 3-variable rules of the form  $mom(A,B) \wedge sister(B,C) \rightarrow mom(A,C)$ . This algorithm is by construction incapable of producing such information. The evaluation procedures presented by Toroghi et al. (2024) use simple and rigidly structured data such that only simple 2-variable rules constructed from existing variables are required.

Since the algorithm output is dependent on the logical proof, if the necessary clauses cannot be found then a full proof will never be found, and the algorithm is forced to provide either no output or a guess. In contrast to these works, ARGOS aims to introduce logical clauses which provide new information about potentially unseen variables and which can be of varying sizes. In order to accommodate the very large search space, we introduce several innovative methods to dynamically select sub-spaces and to efficiently search them. Our proposed method aims to build its solution from both the language-reasoning space and the logic space in order to leverage the exactness of logical tools while remaining robust to failures to find the necessary logical clauses.

**Generalist Reasoners** Despite much research, Chain-of-Thought (COT), proposed by Wei et al. (2022), remains one of the most generally applicable and robust reasoning approaches Plaat et al. (2024). Thus, much research has been done on how to augment the COT process (the previously discussed work by Liu et al. (2024) could be viewed as part of this category). While the focus of our paper is logical algorithms, there are some generalist methods which cannot be ignored. A simple but

highly effective approach known as self-consistency (SC), proposed by Wang et al. (2023), prompts via COT several times and takes the mode of the outputs as the prediction. This method benefits from adding further robustness by smoothing potential outlier errors which might be present in a few chains, and is easily scalable according to computational budget by choosing the appropriate number of samples to take. However, the marginal return decreases as the number of sampled chains increases, as is shown in the original paper Wang et al. (2023). COT and SC's poor performance in proof planning (Saparov and He, 2023) motivated Tree of Thoughts (TOT), proposed by Yao et al. (2023), in which a reasoning tree is hand-crafted for a problem and then the LLM is prompted to explore this tree for likely paths. This method demands a very specific tree topology and exploration designs for each task and so is more of a general framework than an explicit method applicable to logical reasoning. The tree structure of logical reasoning problems even in the same dataset are highly varied. Given this, TOT is poorly suited to logical reasoning settings. There are many works which go beyond TOT in terms of exploring potential chains. Most recently, Xue et al. (2024) proposed a recursive method which also conducts a tree-like search, but allows for dynamically-structured trees.

## F USED DATASETS

The following is a description of the datasets which we have used in our experiments.

**ProntoQA** (Saparov and He, 2023) ProntoQA is a dataset which comes in three types of groundings over the same symbolic structure, which is a string of single-variable modus ponens operations. One of these types is a hand-crafted grounding designed to be true according to commonsense. This dataset is, at this point, fairly easy for language models to handle. However, it remains a common dataset in logical reasoning, and comes with the convenience that a simple random removal of inference rules included in the problem will build abduction cases, as all rules are commonsense. There are 59 problems in the dataset.

CLUTRR (Sinha et al., 2019) CLUTRR is a dataset of family-relational reasoning problems in which some family relations (i.e. "Sam is the mother of John") are given in the form of simple stories (i.e. "John went with his mom Sam to the mall"). Traditionally, the task in CLUTRR is to deduce the relationship between two people, given the context. In order to structure the problem as a true/false classical logic problem, however, we restructure the task to determine if a given relationship between two people is true or not. Practically, we construct the labels by taking the ground-truth relationship as the query 50% of the time (so the new label is "True"), and taking a random other relationship between the given two people ("False") the other 50% of the time, in order to balance the dataset. While the task is naturally abductive in that the input contexts do not include abstract or even grounded relational inference rules (i.e. "if A is the mother of B then B is the child of A"), most symbolic methods rely upon a practitioner to hand-craft a knowledge base of relational rules which are appended to each problem in the dataset. Simply by forbidding this provision, the task becomes truly abductive for the reasoning model. There are 1000 problems in the dataset.

FOLIO (Han et al., 2024) FOLIO is a dataset of logic problems which were hand-crafted by "expert annotators". The dataset is the most diverse of the three we use, both linguistically and structurally. While not perfectly commonsensical, it is generally based in commonsense simply because the annotators were humans who exist in and tend to operate within real-world contexts. Given this pseudo-commonsensicality, random removal of rules to introduce commonsense abduction is not an option, as non-commonsense may be removed from the context. Thus, we engaged human annotators to replace randomly selected phrases from problems with semantically equivalent expressions, so that each replacement would require a minimum of one new rule, indicating that the replaced phrase is implied by its replacement. For example, "NBA Player"  $\rightarrow$  "Pro Basketball Player". Annotators were instructed to reject problems where no replacement could easily be found, and some annotations failed to impact the problem due to either weak replacements or non-interaction with the solve-path of the problem. These problems were discarded, leaving us an abduction-variant of FOLIO with 108 True-or-False problems.

**ESNLI** (Camburu et al., 2018) ESNLI is a dataset of premise-conclusion pairs, stemming from the human-explanation NLI field. A machine is asked to explain how the premise yields or contradicts the conclusion. We adopted this dataset by inverting the task, so that given the premise and conclusion

the machine must determine whether the conclusion is entailed (True) or contradicted (False). The task naturally ensures that abduction is necessary, as we leave out the human explanation with which the pairs are annotated. The dataset, on inspection, is mostly common-sensical.

**CosmosQA** (Huang et al., 2019) CosmosQA is an MCQA dataset designed to test machine reading comprehension. We adapt it to our setting by constructing problems for which the answer is "False" by taking questions for which "None of the Above" is the correct answer and randomly selecting another of the answer choices to be the query (for problems for which the answer is "True", the process is obvious). The dataset is mostly commonsensical.

QUAIL (Rogers et al., 2020) QUAIL is an MCQA dataset also designed to test machine reading comprehension. It is derived primarily by scraping wiki and forum pages on online, and so it often contains artifacts such as timestamps or CSS formatting quirks. We adopt it from MCQA to True/False in an identical fashion to CosmosQA. Becasue of its provenance, questions are often extremely vague; the larger context of the webpage from which the problem comes is not included but is often key to answering the problem. On manual solution of QUAIL problems, the intuitive approach is often to start by inferring the original context of the scraping, making the dataset of high value for abductive settings and for robust evaluation of commonsense flexibility.

#### G UNUSED DATASETS

While there are many logical-reasoning-related datasets available, many are unsuitable because they are either not truly logically-structured or are not commonsensical. Here, we will list some of the most commonly used datasets for logic-adjacent applications and explain their weaknesses/unsuitability for our setting.

**LogiQA** While LogiQA (Liu et al., 2021) is generally commonsensical and logically themed, its questions do not in fact impose an immediate logical problem. In fact, many of the questions are in fact *meta-logical*, in that they ask questions about the underlying logic of the text. For example: "Which of the following makes the same logical mistake as above". These questions could indeed have a formal-logical re-framing, but this would require far more logical aptitude than is currently held by language models, and hand-translating LogiQA questions to logic problems is too time-consumptive.

**RECLOR** (Yu et al., 2020) suffers the same weakness as LogiQA. Questions are metalogical or ask for subjective qualifications regarding some described commonsense logic. Again, this type of question both fails to evaluate true logical reasoning in real-world contexts, and proves problematic for careful evaluation given the inconsistency of the task in that different questions ask different things of the reasoner.

**Soft Reasoner** The Soft Reasoner dataset (Clark et al., 2021) is strictly logical, but is plainly non-commonsensical by construction. The logical problems are constructed without consideration of commonsense or real-world contexts. The dataset is constructed by building clauses from variables/predicates which are randomly selected from a hand-selected bag of words.

**LogicNLI** LogicNLI (Tian et al., 2021) suffer from the Soft Reasoner dataset's weakness to an even greater degree - while the problems are also randomly generated and do not comply with commonsense, they also often do not comply with grammar. For example, phrases such as "Quinlan does not entire" appear frequently. While this may suit the authors' aims of producing arbitrary text as a stand-in for symbolic logic, it is not amenable to the evaluation of real-world logical reasoning in human contexts.

**ProofWriter** ProofWriter (Tafjord et al., 2021) again suffers from the same weakness which all semi-random auto-generated datasets suffer: non-commonsenseness. By picking clauses randomly by sampling bags of predicates, no guarantees can be made on the realism of the data. Examination of the dataset will show that the "facts" described in problem contexts range from unlikely to non-sensical - the very first problem includes a rule "All red things are rough". Within the real world this is obviously not true, as we can find examples of red things which are not rough. Of course, it was not the dataset creators' aim to build a commonsensical dataset and so this is of little surprise.

## H CONFUSING PROBLEMS FOR ARGOS

 Here' we give some examples of problems which confused ARGOS. Note that all of the problems provided are examples in which the given premises are at least somewhat contradictory with commonsense, breaking the setting.

Table 10: Confusing since Fido is typically specifically a dog's name (and not a cat's). If we ask the LLM if Fido is a dog, with no context, it will say yes.

All tigers are cats. No cats are dogs. All Bengal tigers are tigers. All huskies are dogs. Fido is either a Bengal tiger or a cat.

True or False: Fido is a husky animal.

#### Table 11: Confusing since Detroit City is probably not a horse according to commonsense.

Detroit City is a horse. Some horses are racehorses. If a horse falls in a race, it poses risks to its rider. Detroit City fell in a race. A horse is a racehorse if it is in a race. True or False: Detroit City has been in multiple races.

Table 12: Confusing because (1) edible can refer to beverages when the contextual distinction is between safe and unsafe for consumption, but not when the distinction is between eating and drinking; (2) Coke is not apple juice.

All drinks on the counter are edible. All juices on the counter are drinks. Orange juice is a type of juice. Everything on the counter is either orange juice or apple juice. All apple juices on the counter are sweet. The coke is on the counter and if the coke is apple juice, then the coke is a drink. If the coke is not apple juice, then the coke is not edible. True or False: The coke is edible and sweet.

## I SOLVABILITY PROGRESSION ILLUSTRATION

In this section we provide figures for CLUTRR, CosmosQA and QUAIL, illustrating how ARGOS' SC-solvability measure progresses per-question, over a 100-question clipping from each dataset. Positive % values indicate a positive classification, and negative the like.



Figure 7: CosmosQA: This dataset, being mostly solved by vanilla SC, sees little fluctuation and exits from ARGOS early. With that said, we still see a flip from low-confidence negative to high-confidence positive.

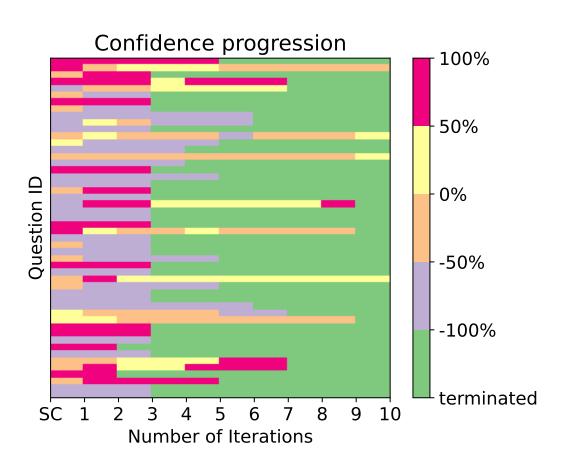


Figure 8: CLUTRR: We see a significant amount of confidence fluctuation and flipping, indicating that meaningful elements within the logic of the problem are being modified by ARGOS in order to affect the answer. This is not surprising, since our construction of generated propositions as taking purely backbone variables as antecedents ensures that added ARGOS propositions will be effectual.

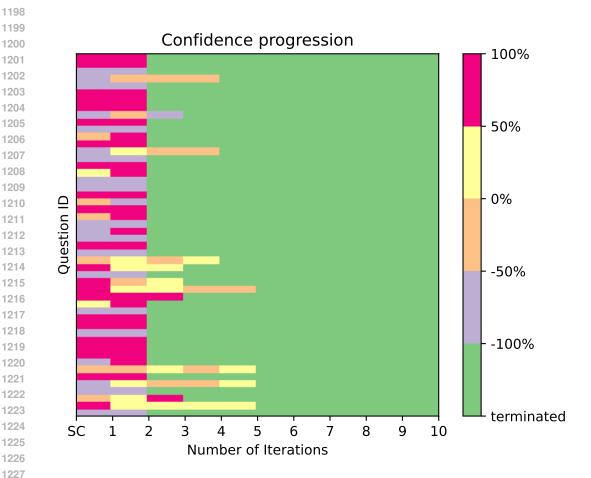


Figure 9: QUAIL: despite ARGOS's strong performance on QUAIL relative to SC, we find that in fact few ARGOS iterations are necessary. While QUAIL is made complicated for language models by its irregular form and often disjoint nature, its logical structure (while very ambiguous) is simple, meaning that we can solve QUAIL problems with only a small number of well-chosen propositions.