FLARE: FAITHFUL LOGIC-AIDED REASONING AND EXPLORATION

ABSTRACT

Modern Question Answering (QA) and Reasoning approaches based on Large Language Models (LLMs) commonly use prompting techniques, such as Chain-of-Thought (CoT), assuming the resulting generation will have a more granular exploration and reasoning over the question space and scope. However, such methods struggle with generating outputs that are faithful to the intermediate chain of reasoning produced by the model. On the other end of the spectrum, neuro-symbolic methods such as Faithful CoT (F-CoT) and Logic-LM propose to combine LLMs with external symbolic solvers. While such approaches boast a high degree of faithfulness, they usually require a model trained for code generation and struggle with tasks that are ambiguous or hard to formalise strictly. We introduce **F**aithful **L**ogic-**A**ided **R**easoning and **E**xploration (**FLARE**), a novel interpretable approach for traversing the problem space using task decompositions. We use the LLM to plan a solution, formalise the query into facts and predicates, which form the problem space, using a logic programming code and simulate that code execution using an exhaustive multi-hop search over the defined space. Our method allows us to compute the faithfulness of the reasoning process w.r.t. the generated code and explicitly trace the steps of the multi-hop search without relying on external solvers. Our methods achieve SOTA results on **7** out of **9** diverse reasoning benchmarks. We also show that model faithfulness positively correlates with overall performance and further demonstrate that **FLARE** allows pinpointing the decisive factors sufficient for and leading to the correct answer with optimal reasoning during the multi-hop search. Our findings reveal that successful traces exhibit, on average, a 18.1% increase in unique emergent facts, a 8.6% higher overlap between code-defined and execution-trace relations, and a 3.6% reduction in unused code relations.

1 INTRODUCTION

Complex Reasoning in natural Question Answering (QA) tasks assumes the capability to explore the problem space of the designated query with a formalised set of facts, relations, commonsense knowledge and logical implications. In line with this, LLMs have been enhanced with CoT (Wei et al., 2022) prompting, which supplements the QA process by generating intermediate reasoning chains given a set of in-context examples (Brown et al., 2020a), as shown in Figure 1. This allowed for advancement in commonsense (Madaan et al., 2022), symbolic (Wang et al., 2022; Sprague et al., 2024) and mathematical (Jie et al., 2023) reasoning. Although CoT allows for a problem exploration in natural language steps, such an approach has been shown to cause performance degradation for reasoning tasks involving multi-step planning (Valmeekam et al., 2022; Suzgun et al., 2023), problem exploration (Yao et al., 2022), and arithmetic tasks (Hendrycks et al., 2021b; Madaan & Yazdanbakhsh, 2022b). These discrepancies arise as CoT suffers from a limited ability to decompose, search, verify and backtrack using intermediate rationale chains (Yao et al., 2022), cascading hallucinations and errors (Ling et al., 2023) and that natural language might not be an optimal representation for describing the reasoning process (Li et al., 2024). Simultaneously, LLM output has been shown to be unfaithful and inconsistent w.r.t. the intermediate CoT rationale (Jacovi et al., 2024; Lanham et al., 2023b; Turpin et al., 2023).

To mitigate the problem of CoT faithfulness and allow for more robust reasoning during QA, Lyu et al. (2023, Faithful CoT) and Logic-LM (Pan et al., 2023) suggested generating code which is

Figure 1: A depiction of the *plan*, *code* and simulated *search* in FLARE. Each module is generated separately and iteratively, allowing us to obtain the final answer. The green and yellow highlighted text shows the overlap between the facts and the relations between the code and the simulated search.

further executed using an external symbolic solver. Producing and executing code enables the generation of outputs guided by external solvers, leveraging search with backtracking to explore the problem space effectively. However, strict translations of natural language queries into code, such as *autoformalisation* (Szegedy, 2020; Wang et al., 2018), is a non-trivial task involving direct inference of implicit commonsense and domain-specific knowledge and the ability to align abstract and informal concepts directly to constrained formal definitions for further execution (Wu et al., 2022). An example query, *"Do all parts of the aloe vera plant taste good?"*, is challenging to formalize or address with a strict algorithmic solution, as it requires interpretative, deductive and context-dependent reasoning, referred to as soft or fuzzy reasoning. Using external solvers makes such fuzzy reasoning impossible and requires consistently generating syntactically correct executable code. While some LLMs have coding capabilities stemming from their pretraining (Jiang et al., 2024; Aryabumi et al., 2024), relative code consistency is more probable with models explicitly trained for coding (Chen et al., 2021).

To overcome these problems, we propose Faithful Logic-Aided Reasoning and Exploration (FLARE), an interpretable method that allows for planning, fuzzy reasoning, and traversing the problem space with backtracking, exact task decomposition, and measuring faithfulness. In FLARE, given a natural language query, we prompt an LLM to sequentially generate a *plan* that includes an analysis and the logical steps necessary for formalising and answering the question, a logic programming (Wielemaker et al., 2012) *code* that allows formalising the query into a set of facts, relations and their composition forming the space for exploring that query and the *search*, which is an LLM-generated code execution simulation. An illustration of FLARE can be seen in Figure 1. In our framework, the generated code must not be consistently executable by an external solver, allowing for the soft-formalisation of natural language. Although we see that even generalist LLMs are able to produce executable code in $\geq 50\%$ of cases. FLARE allows us to measure the faithfulness of the outcome w.r.t. the simulated code execution by directly comparing the search paths produced by the external solver to that LLM generation. This comparison also allows for pinpointing model hallucinations and inconsistencies. We systematically study the effectiveness of our method using 4 general-purpose LLMs of varying scales across 9 diverse QA and 3 logical inference benchmarks, covering Math World Problems, Multi-hop QA, Relation inference, deductive and analytical reasoning and show that our method achieves state-of-the-art results in 7 out of 9 QA datasets and 2 out of 3 logic datasets in comparison to CoT, F-CoT and Logic-LM. We also show that the method is competitive for models tuned for coding, with an average overall increase of 16% over F-Cot and 9% over CoT. Our findings show that model accuracy strongly correlates with the faithfulness of the reasoning process towards search traces from the simulated code execution. We also provide ablations showing that the model can interpretably pinpoint hallucinations, underutilized knowledge, and the limitations of the search over the problem space. Our key contributions are the following:

- We introduce FLARE a novel paradigm for logic-aided and interpretable formalisation and search over the problem space in QA and logic reasoning tasks.

- We perform a systematic evaluation across 9 QA and 3 logical inference benchmarks and 4 models of varying scales, showing the advantages of using FLARE for QA in a few-shot setup over prior approaches.

- The modularity of FLARE allows defining a simple ingrained method for measuring model faithfulness, which is further shown to be strongly correlated with performance.

- We further show that using FLARE allows us to interpretably and rigorously detect hallucinations along with sub-optimal and inconsistent reasoning patterns.

2 RELATED WORK

Reasoning in Natural Language Few-shot prompting (Brown et al., 2020b) has been shown to be an effective approach for increasing the reasoning capabilities of LLMs in natural language generation (Gehrmann et al., 2021; Reif et al., 2022; Sanh et al., 2022). LLM reasoning can be further enhanced with prompting techniques such as CoT (Wei et al., 2022), which attempts to segment reasoning into explicitly written intermediate steps. Concurrent work has also proposed that models *"think step by step"* (Kojima et al., 2022), or divide the problem into subtasks before the solution (Zhou et al., 2023, Least-to-Most). These approaches have been shown to suffer from arithmetic inaccuracies (Lewkowycz et al., 2022; Hendrycks et al., 2021a) and reasoning inconsistencies (Madaan & Yazdanbakhsh, 2022a). Further attempts have been made to add a planning stage before reasoning by dividing the process into recursive plan formulation and execution steps (Yao et al., 2023b; Wang et al., 2023a). The *plan* generation step in FLARE is a hybrid technique inspired by these methods but focused on generating a natural language strategy for formalising the query into code.

Reasoning with Search Several lines of work propose using techniques to expand the reasoning paths over the problem space. Self-consistency decoding (Wang et al., 2023b) is an approach used to sample many natural language reasoning paths and take a majority vote for an answer. Another popular approach is Tree-of-Thoughts (ToT; Yao et al., 2023a), which proposes to explore the query with reasoning similar to a tree traversal, where each state is created and evaluated using an LLM. Similar techniques try to adapt symbolic search approaches akin to DFS, BFS (Besta et al., 2024), A^* (Lehnert et al., 2024) or other combinations (Gandhi et al., 2024) with direct tuning (Lehnert et al., 2024), imitation training (Yang et al., 2022) or few-shot prompting (Zhang et al., 2024). It must be noted that all of these techniques have only been tested in constrained mathematical puzzle-solving and algorithmic domains like the 24 Game (Yang et al., 2022), Countdown (Wikipedia, 2024), Sorting (Besta et al., 2024), maze solving (Yang et al., 2022), Sokoban (Lehnert et al., 2024), and others. Although the *search* component of FLARE has some similarities to these techniques, we argue that our method allows for generalistic reasoning with interpretable multi-hop search through simulated code execution.

Reasoning with Formalisation Another line of research has tried formalising natural language queries into code (Gao et al., 2023; Li et al., 2024) or pseudo-code (Chae et al., 2024; Gandhi et al., 2024). This allows the translation of the query into a strict structure and offloads the reasoning and search components to deterministic solvers like Python Chen et al. (2023), PDDL Lyu et al. (2023); Liu et al. (2023), DataLog Lyu et al. (2023) and others. While models are capable of synthesising programs (Austin et al., 2021; Nijkamp et al., 2023) and benefit from the use of code in numerical and algorithmic reasoning settings (Chen et al., 2023; Gao et al., 2023), the usage of code for general QA has not been rigorously explored. The reasons are that formalisation from natural language into a strict and executable code is challenging (Wu et al., 2022), following the exact syntactic constraints of the programming language not abundantly used during pre-training is onerous (Liu et al., 2024) and can require models explicitly tuned for coding (Chen et al., 2021). Using an external solver for reasoning also limits the capability for soft reasoning in commonsense knowledge and implications. Although we formalise the natural language query into a logic programming Prolog program during the *code* generation part of FLARE, we do not explicitly require the code to be executable and do not use external solvers during inference. This allows for the further use of the LLM for soft-reasoning to simulate code execution in a logic-based problem space traversal similar to Prolog while circumventing the need for code tuning a generalist model.

Method	Math Word Problems					Multi-hop QA			Relation
	GSM8K	SVAMP	MultiArith	ASDiv	AQuA	StrategyQA	Date	Sport	CLUTRR
Llama-3.1-8B$_{\text{FLARE}}$	**72.7**	**86.0**	**96.3**	**83.1**	**62.9**	**70.2**	**59.3**	76.6	_36.8_
Llama-3.1-8B$_{\text{F-CoT}}$	_0_	_0_	_0_	_0_	_12.2_	53.2	_0_	_0_	_32_
Llama-3.1-8B$_{\text{CoT}}$	59.2	58.6	60.1	61.9	35	_2.9_	20.9	**95.8**	**42.2**
CmDR$_{\text{FLARE}}$	**52.4**	**74.0**	**84.5**	**72.2**	**43.7**	**67.0**	**52.3**	**78.9**	_29.1_
CmDR$_{\text{F-CoT}}$	_0_	_0_	_0_	_0_	_0_	59.7	_0_	_0_	_8.6_
CmDR$_{\text{CoT}}$	46.5	57.3	83.1	37.2	28.3	_21.3_	47.4	55.2	**29.5**
CmDR+$_{\text{FLARE}}$	**71.4**	**83.5**	**90.4**	**81.3**	**55.9**	**70.8**	61.8	**77.7**	_41.0_
CmDR+$_{\text{F-CoT}}$	_0_	_0_	_0_	_0_	_15.4_	57.6	_0_	_0_	35.3
CmDR+$_{\text{CoT}}$	48.7	81.1	86.6	44.6	44.1	_48.4_	**79.1**	62.6	**42.5**
GPT-3.5$_{\text{FLARE}}$	_68.1_	82.7	**98.3**	**85.4**	55.1	**65.5**	**82.4**	85.6	**49.8**
GPT-3.5$_{\text{F-CoT}}$	75.8	**83.0**	95.3	81.7	_53.5_	_51.5_	73.5	_52.3_	12.1
GPT-3.5$_{\text{CoT}}$	**79.8**	_82.4_	_98.2_	_75.8_	**59.4**	51.7	_69.9_	**95.8**	_4.3_

Table 1: The following table shows the performance of each of the tested models given a technique for reasoning. Each **bold**, underlined, and _italicised_ element highlights the best, second best and worst technique per specific model. The overall best method per dataset is highlighted in green .

Reasoning Faithfulness An explanation is considered _faithful_ if it explicitly and accurately describes the reasoning process of the model during inference (Gilpin et al., 2018; Jacovi & Goldberg, 2020). In the context of prompting techniques such as CoT, we are interested in the faithfulness of the intermediate reasoning chains towards the final output. Faithful intermediate reasoning chains should not just look _plausible_ (Herman, 2017) but have exact reflections of the problem exploration and reasoning used to arrive at the final answer. Natural language reasoning chains prevalent in CoT and similar methods are shown to be unfaithful, either masking the reasoning biases (Turpin et al., 2023) of the model or outright ignoring the intermediate reasoning (Lanham et al., 2023a). In FLARE, we introduce a method to seamlessly measure the faithfulness of the final outcome w.r.t. completed search.

3 METHODOLOGY

3.1 LLM SIMULATED SEARCH

FLARE comprises three modules for generating a _plan_, _code_ and simulated _search_ for answering a natural language query $\mathcal{Q} = \{T_1^{\mathcal{Q}} \ldots T_{|\mathcal{Q}|}^{\mathcal{Q}}\}$, where each $T_i^{\mathcal{Q}}$ is a token in the query $\mathcal{Q}$.

Generating A Plan For each query $\mathcal{Q}$, given an LLM $\mathcal{M}$, we initially use instructions $\mathcal{I}^{\mathcal{P}}$ to prompt it to generate a _plan_ $\mathcal{P}$, which should be comprised of task explanation, analysis and a plan for further formalising the query. An example of this can be seen in the _plan_ section in Figure 1. We use in-context few shot examples $\mathcal{E}_{\mathcal{P}}$ of such _plan_ generations along with **greedy** decoding for obtaining the final plan.

$$\mathcal{P}_i \sim \arg\max p_{\mathcal{M}}(T_i^{\mathcal{P}} \mid T_{:i-1}^{\mathcal{P}}, \mathcal{E}_{\mathcal{P}}, \mathcal{Q}, \mathcal{I}^{\mathcal{P}}) \tag{1}$$

Where $\mathcal{P}_i$ and $T_i^{\mathcal{P}}$ is the i-th token in the generated _plan_ $\mathcal{P}$ and $p_{\mathcal{M}}$ is the probability of the next token over the vocabulary obtained from model $\mathcal{M}$.

Generating Code After generating the _plan_, we use instructions $\mathcal{I}^{\mathcal{C}}$ to prompt the LLM $\mathcal{M}$ to generate a Prolog code $\mathcal{C}$, an example of which can be seen in Figure 1. We append executable code generation samples $\mathcal{C}_{\text{sample}}$ to the previous in-context examples $\mathcal{E}_{\mathcal{P}}$ and obtain few-shot code generation demonstrations $\mathcal{E}_{\mathcal{C}} = [\mathcal{E}_{\mathcal{P}}; \mathcal{C}_{\text{sample}}]$

$$\mathcal{C}_i \sim \arg\max p_{\mathcal{M}}(T_i^{\mathcal{C}} \mid T_{:i-1}^{\mathcal{C}} \mathcal{E}_{\mathcal{C}}, \mathcal{Q}, \mathcal{I}^{\mathcal{P}}, \mathcal{P}, \mathcal{I}^{\mathcal{C}}) \tag{2}$$

$$\mathcal{F}_{code}, \mathcal{R}_{code}, \mathcal{G}_{code} = \text{EXTRACT}(\mathcal{C}_i)$$

Where $\mathcal{C}_i$ and $T_i^{\mathcal{C}}$ is the i-th token in the generated _code_ $\mathcal{C}$.

Figure 2: The trend of mean model accuracy w.r.t mean faithfulness for all the models.

Benefits of Prolog Prolog is a symbolic logic-programming engine (Bowen, 1979) designed for heuristic search over Horn Clauses (Chandra & Harel, 1985). As a declarative programming paradigm (Lloyd, 1994), the code is expressed as the logic of computation, expressed as a set of facts $\mathcal{F}$ and relations $\mathcal{R}$ defining the problem space, with the goal $\mathcal{G}$ being a first-order logic combination of them. Prolog employs depth-first search (DFS) (Bowen, 1979) for sub-goal decomposition and traversal of the problem space, satisfying $\mathcal{G}$ through a sequence of steps known as the *trace*. Each step either confirms/invalidates a sub-goal, expands the search tree, or retries failed sub-goals with new combinations. An example of such a search is shown in Figure 1. Prolog supports exhaustive search by exploring all paths that satisfy or fail the goal. This explicit segmentation of facts, relations, and search strategies simplifies query formalization. As a declarative language, Prolog enables segmentation using a simple regexp heuristic, referred to as EXTRACT in Equation (2) and Equation (3). Including exhaustive traces in-context allows an LLM to simulate sub-goal decomposition, backtracking, and intermediate goal invalidation, discussed further in the next paragraph.

Simulating Search After generating the logic-programming *code*, we want to simulate program execution by generating a problem space traversal trace with our LLM $\mathcal{M}$. We use instructions $\mathcal{I}^{\mathcal{S}}$ and update our in-context samples by appending search traces $\mathcal{S}_{\text{sample}}$ constructed from Prolog execution of sample codes $\mathcal{C}_{\text{sample}}$, i.e. $\mathcal{E}_{\mathcal{S}} = [\mathcal{E}_{\mathcal{C}}; \mathcal{S}_{\text{sample}}]$.

$$\mathcal{S}_i \sim \arg\max p_{\mathcal{M}}(T_i^{\mathcal{S}} \mid T_{:i-1}^{\mathcal{S}}\mathcal{E}_{\mathcal{C}}, \mathcal{Q}, \mathcal{I}^{\mathcal{P}}, \mathcal{P}, \mathcal{I}^{\mathcal{C}}, \mathcal{C}, \mathcal{I}^{\mathcal{S}}) \qquad (3)$$
$$\mathcal{A}_{search}, \mathcal{F}_{search}, \mathcal{R}_{search} = \text{EXTRACT}(\mathcal{S}_i)$$

Where $T_i^{\mathcal{S}}$ is the i-th token in the generated *search* trace $\mathcal{S}$. During iterative problem space traversal, we can segment the facts $\mathcal{F}_{search}$, relations $\mathcal{R}_{search}$, completed and backtracked paths with their answers $\mathcal{A}_{search}$ used during the search simulation. To get the final answer we update in-context samples with their correct final answers $\mathcal{A}_{\text{sample}}$ from the executed search $\mathcal{S}_{\text{sample}}$, $\mathcal{E}_{\mathcal{A}} = [\mathcal{E}_{\mathcal{S}}; \mathcal{A}_{\text{sample}}]$ and use instructions $\mathcal{I}^{\mathcal{A}}$ to obtain the final answer from the model.

$$\mathcal{A}_{Final} \sim \arg\max p_{\mathcal{M}}(T_i^{\mathcal{A}} \mid T_{:i-1}^{\mathcal{A}}\mathcal{E}_{\mathcal{C}}, \mathcal{Q}, \mathcal{I}^{\mathcal{P}}, \mathcal{P}, \mathcal{I}^{\mathcal{C}}, \mathcal{C}, \mathcal{I}^{\mathcal{S}}, \mathcal{S}, \mathcal{I}^{\mathcal{A}}) \qquad (4)$$

The prompts used for generating each part in FLARE can be seen in Appendix A along with a complete example in Table 8.

3.2 DETECTING REASONING INCONSISTENCIES

For each query $\mathcal{Q}$ given the *code* $\mathcal{C}$ and the simulated *search* $\mathcal{S}$ along with the extracted facts $\mathcal{F}_{code}, \mathcal{F}_{search}$ and relations $\mathcal{R}_{code}, \mathcal{R}_{search}$ from each designated module, we aim to detect the inconsistencies during the reasoning process of the LLM. We use exact string matching between all

these facts and relations in *code* and simulated *search*.

$$\forall i, \exists j \quad \text{such that} \quad \mathcal{F}_{\text{code}}^i = \mathcal{F}_{\text{search}}^j \quad \text{and} \quad \forall v, \exists q \quad \mathcal{R}_{\text{code}}^v = \mathcal{R}_{\text{search}}^q \tag{5}$$

$$\forall j, \exists i \quad \text{such that} \quad \mathcal{F}_{\text{code}}^i = \mathcal{F}_{\text{search}}^j \quad \text{and} \quad \forall q, \exists v \quad \mathcal{R}_{\text{code}}^v = \mathcal{R}_{\text{search}}^q \tag{6}$$

With this framework in mind, we define two reasoning failure modes. In the *first* failure mode, given that some fact or relation was used in the simulated *search* but did not exist in the generated *code*, i.e. $\exists j$ such that $\mathcal{F}_{\text{search}}^j \notin \mathcal{F}_{\text{code}}$, we claim that the LLM has *hallucinated*. We postulate that the model either produced incomplete knowledge during formalisation to *code* or created a piece of non-existing information during the *search*. We do not consider facts that emerged during a direct inference step within the simulated search during our calculation. For example, if

Dataset	ChatGPT (gpt-3.5-turbo)			
	Standard	CoT	Logic-LM	FLARE
PrOntoQA	47.40	67.80	61.00	**73.40**
LogicalDeduction	40.00	42.33	**65.67**	58.60
AR-LSAT	20.34	17.31	26.41	**27.39**

Table 2: Comparison of Direct Prompting, CoT, Logic-LM and FLARE.

we are dealing with a mathematical query $4 \cdot (5 + 6) = ?$, the search would involve separately evaluating the expression $5 + 6 = 11$. In this case, 11 will not be treated as a hallucinated fact within the search but rather as an emergent fact obtained from direct inference. The *second* failure mode is the reciprocal case, where a fact or relation present in the *code* is not used during the *search*. We refer to this phenomenon as *sub-optimal reasoning* as it shows that the LLM could not explore the problem space completely or injected unsuitable knowledge during formalisation into *code*.

3.3 MEASURING FAITHFULNESS

We propose a method to measure the faithfulness of the LLM reasoning process when using FLARE. As mentioned in Section 3.1, for each query in a dataset $\mathcal{D} = [\mathcal{Q}_1, \ldots, \mathcal{Q}_{|\mathcal{D}|}]$, we generate a set of codes $\Phi = [\mathcal{C}_1, \ldots, \mathcal{C}_{|\Phi|}]$ and simulated problem space searches $\Psi = [\mathcal{S}_1, \ldots, \mathcal{S}_{|\Psi|}]$. We use the Prolog engine to execute all of the codes Φ and obtain a set of correctly written programs Φ' and exact search paths Ψ'. As we do not require explicit programmatic correctness during inference in FLARE for any code $\mathcal{C}_i$, some Prolog executions resulting in an error are filtered out in Ψ'. To assess model reasoning faithfulness towards code formalisations, we compare the search paths Φ' obtained from Prolog execution with their designated counterparts Φ'_{gen} generated by the LLM from the same code. We use ROUGE (Lin, 2004) to compute the matching score for each executed and simulated search path. In particular, we use ROUGE-Lsum, which uses the longest common subsequence (LCS) over each line to obtain the final score. This method fits our cause as a line in a Prolog search execution represents a single logic step within the traversal. This allows us to measure the similarity of the reasoning contents and structure in exact and simulated searches.

4 EXPERIMENTAL SETUP

4.1 DATASETS

To evaluate FLARE, we use a benchmark of 9 tasks covering Math Word Problems (MWP), multi-hop QA and relation inference, and 3 common logical reasoning datasets. For testing numeric and mathematical reasoning, we follow CoT (Wei et al., 2022) by including GSM8K (Cobbe et al., 2021), SVAMP (Patel et al., 2021), MultiArith (Roy & Roth, 2015), ASDiv (Miao et al., 2020) and AQuA (Ling et al., 2017). Among these, GSM8K, SVAMP, MultiArith and ASDiv cover elementary and middle school arithmetic word problems with a set of integers or decimals as the answer. AQuA is a multiple-choice numerical, symbolic reasoning dataset where each answer is a mathematical expression containing notations, values and expressions not defined in the query. We also test FLARE using three multi-hop QA datasets. We use StrategyQA (Geva et al., 2021), which is a boolean QA task that requires sub-goal decomposition and a multi-hop reasoning strategy to answer. The example *"Do all parts of the aloe vera plant taste good?"* used in Figure 1, is taken from StrategyQA. The multi-hop QA testing also includes Date and Sports Understanding, subsets of BIG-Bench (bench authors, 2023). The tasks involve inferring an exact date given some calculations in the relative time period and understanding if an artificially created sports statement is feasible. Furthermore, we assess FLARE on Relational Inference using CLUTRR (Sinha et al., 2019), which involves inferring

Method	CmDR$_{FLARE}$	CmDR$_{plan-only}$	CmDR+$_{FLARE}$	CmDR+$_{plan-only}$	GPT-3.5$_{FLARE}$	GPT-3.5$_{plan-only}$
GSM8K	**52.4**	24.7	**71.4**	40.7	**68.1**	36.1
AQuA	**43.7**	35.0	**55.9**	55.1	**55.1**	54.3
StrategyQA	**67.0**	65.5	70.8	**75.7**	**65.5**	62.3

Table 3: The table shows the accuracy of an LLM with FLARE compared to prompting for a final answer directly after generating (plan-only) a plan $\mathcal{P}$.

Figure 3: The figure shows the percentage of executable code per model (right) and the accuracy of the executable code when answering the queries (left).

the familial relation between two entities mentioned in a natural language description of the partial family graph. We evaluate FLARE on challenging logic inference datasets: ProntoQA (Saparov & He, 2023), AR-LSAT (Zhong et al., 2021), and LogicalDeductions from BigBench (et al., 2023), focusing on harder subsets proposed by Pan et al. (2023). These datasets, covering deductive, analytical, and logical reasoning, allow us to assess FLARE's performance. Details, including descriptions and examples, are in Table 9 of Appendix A.

4.2 BENCHMARKS

We compare FLARE with CoT (Wei et al., 2022) as a prompting method that reasons using natural language chains and with F-CoT (Lyu et al., 2023) and Logic-LM (Pan et al., 2023) that formalise the query into a code and offload the reasoning to an external symbolic solver. We use Llama3.1 (8B) (Dubey et al., 2024), CmDR (30B) (Cohere, 2024), CmDR+ (100B) (Cohere, 2024) and GPT3.5 (Brown et al., 2020b) ($\geq 100B$ (Ye et al., 2023)). As the coding model OpenAI Codex (code-DaVinci-002) (Chen et al., 2021) used in F-CoT has been deprecated, we replace it with the new GPT3.5 as suggested by OpenAI and recalculate the results accordingly.

5 RESULTS

5.1 FEW-SHOT PROMPTING

To evaluate FLARE, we use a set of models of varying sizes on diverse benchmarks, as defined in Section 4. We compare the performance of each model while using FLARE, CoT and F-CoT prompting. The results for F-CoT and CoT on all the models are computed using the codebase of the original study (Lyu et al., 2023). We additionally compare Logic-LM and FLARE using the logic reasoning benchmarks proposed in Pan et al. (2023).

Model	Avg. Number of Paths	Avg. #Hops per path	Avg. #Fails per path	Avg. Total Hops	Avg. Total Fails
		Incorrect Answers			
Llama-3.1-8B$_{FLARE}$	1.55	11.12	1.52	15.09	2.26
CmDR$_{FLARE}$	1.51	6.55	0.68	10.56	1.39
CmDR+$_{FLARE}$	0.92	7.52	1.13	8.57	1.32
GPT-3.5	0.68	5.22	0.71	5.32	0.74
		Correct Answers			
Llama-3.1-8B$_{FLARE}$	1.43	9.12	0.62	12.36	0.96
CmDR$_{FLARE}$	1.19	7.10	0.42	11.29	0.66
CmDR+$_{FLARE}$	0.97	7.19	0.42	8.22	0.61
GPT-3.5$_{FLARE}$	0.82	5.65	0.26	5.69	0.27

Table 4: The table depicts the difference in the average explored paths, hops, and fails during the reasoning process, which leads to incorrect or correct answers. The purple colour illustrates that incorrect reasoning paths have fewer explorations that led to Failed search paths.

LLMs for general reasoning   Our results, presented in Table 1, show that using FLARE allows the LLMs to achieve state-of-the-art results on 7 out of 9 datasets, with an average 28% increase over CoT. We can see a clear trend that FLARE increases the performance compared to CoT and F-CoT for all the models of varying scales. We also see that LLMs that are not explicitly tuned for coding suffer massive degeneracies when using F-CoT. We postulate that they are unable to consistently produce executable programs that satisfy a predefined scheme in F-CoT, thus resulting in an error during execution. This further highlights the value of simulating program execution using an LLM instead of using external solvers. The results show that using FLARE yields more benefit on datasets that require longer chains of multi-hop and symbolic reasoning, like AQuA and StrategyQA. Our findings in Table 2 show that FLAREachieves state-of-the-art results on 2 out of 3 logic inference benchmarks with 10% increase over CoT and 7% increase over Logic-LM.

LLMs for code generation   To understand the effect of FLARE on models tuned for coding, we use GPT3.5 (Brown et al., 2020a) as it was the OpenAI suggested succession model for Codex (Chen et al., 2021) which is used in F-CoT and possesses strong coding capabilities (Ye et al., 2023). The results in Table 1 show that using FLARE is beneficial for models that are tuned for coding and boost the accuracy with a 16% increase over F-CoT and 9% over CoT. The reason is that many natural language queries with non-trivial formalisations are more suited to be tackled with more commonsense soft reasoning than direct code execution. This is evident in Table 1 where FLARE and CoT are consistently better than F-CoT in StrategyQA, Sports and CLUTRR. The opposite case of numeric and algorithmic heavy reasoning tasks is also covered by FLARE as it maintains strong performance similar to F-CoT on MWP problems Table 1. Consequently, FLARE allows combining algorithmic formalisation with simulated soft-reasoning, circumventing the pitfalls of using a deterministic external solver while still producing a query formalisation and problem space traversal.

5.2 IS SIMULATING SEARCH USEFUL?

To understand if simulating a search over the problem space is useful, we compare the performance of FLARE where we only generate the *plan* without the subsequent *code* or *search* components. We refer to this framework setup as *plan-only*, which can be seen in Figure 1 if we were to use only the *plan* for answer generation. We completed this ablation using CmDR, CmDR+, and GPT-3.5, and we used GSM8K, AQuA, and StrategyQA as our baselines. The results in Table 3 confirm that all of the models suffer massive performance degradation from $61.1 \rightarrow 49.9$ when omitting the *code* and the *search* components of FLARE. We hypothesise that this is caused by insufficient problem space exploration when using the *plan-only* setting. Furthermore, we have already seen in Table 1 that in methods, like F-CoT, that do not use simulated problem space exploration for soft-reasoning and only rely on *plan* and *code*, the performance also deteriorates even resulting in a complete breakdown of reasoning over the designated datasets. This can be viewed as a constrained version of FLARE with *code-only* execution. Consequently, our results show that simulating problem space traversal is highly beneficial as it avoids the pitfalls posed by *plan-only* and *code-only* modes by exploring the problem space more rigorously and soft-reasoning during that traversal instead of using external solvers.

5.3 FAITHFUL REASONING IMPROVES PERFORMANCE

As described in Section 3, using FLARE allows us to measure the faithfulness of the LLM reasoning process by comparing the simulated problem space traversals Φ'_{gen} with actual traces Φ' produced from a symbolic Prolog solver. To do this, we initially compute the percentage of syntactically correct executable code each LLM produces. We can see from the right part of Figure 3 that all of the models are capable of producing correct executable Prolog code in 67% of cases on average and $\geq 50\%$ of cases at the very least. This shows that the simulated searches Φ'_{gen} can be considered a representative sample that will be further used to accurately measure the faithfulness of the simulated search w.r.t. the generated code. After measuring the reasoning faithfulness for each model, we want to understand what impact it has on the performance of the LLM. In Figure 2, we segment the models w.r.t. their ROUGE-Lsum scores. The results show that model perfor-

Model	UEF (%) in Search	OR (%)	UR (%) in code
Correct Answers			
Llama-3.1-8B$_{FLARE}$	74.14	43.65	5.73
CmDR$_{FLARE}$	59.06	35.96	4.02
CmDR+$_{FLARE}$	64.30	34.47	4.54
GPT-3.5$_{FLARE}$	64.46	37.55	1.90
Avg. (Correct)	65.49	37.91	4.05
Incorrect Answers			
Llama-3.1-8B$_{FLARE}$	54.69	35.04	9.28
CmDR$_{FLARE}$	54.50	32.76	6.23
CmDR+$_{FLARE}$	44.12	24.98	8.22
GPT-3.5$_{FLARE}$	36.02	24.44	6.94
Avg. (Incorrect)	47.33	29.31	7.67
Δ	18.16	8.60	-3.62

Table 5: The table shows how the percentage of unique emergent facts (UEF) in search, overlapping relations (OR) between code and search, and unused relations (UR) in code impact answer correctness.

mance is strongly positively correlated with reasoning faithfulness. However, we also observe in the left part of Figure 3 that executing semantically precise code results in an accurate answer only in 47% of cases on average. Indeed, having a simulated search trace with a ROUGE-Lsum faithfulness score of 1, would be equivalent to simply executing the program as proposed in F-CoT. Yet we have priorly shown that F-CoT struggles with reasoning tasks that are hard to formalise and require multi-hop commonsense and soft reasoning. These two discoveries show that optimal LLM reasoning, conditioned on a search in the problem space, should be increasingly faithful toward the facts, relations and the search strategy defined within the code while simultaneously maintaining the capability for soft-reasoning along more abstractly defined concepts. Our results show that FLARE allows LLMs to maintain a similar reasoning capacity.

5.4 WHAT IS IMPORTANT DURING THE SEARCH?

We expand the analysis of the simulated search traces to detect the reasons which can lead to optimal reasoning within an LLM. For this purpose, we calculate several statistics, like the average number of explored paths, average and total hops and failures per path, for each model during the simulated traversal. The failure in a path is an invalidation of a solution for a sub-goal explored during the search, which is used for backtracking, as explained in Section 3. Calculating these statistics is simple as the *search* component of FLARE, seen in Figure 1, is a structured simulation of a Prolog trace, where each line contains a hop of reasoning inference. We split these statistics for the

Model	Hops per Paths	Halluc. (%)	Unutil. (%)
Llama-3.1-8B	9.4	63.3	62.9
CmDR	6.7	54.7	56.9
CmDR+	7.2	54.3	56.3
GPT-3.5	5.5	49.3	52.1

Table 6: The table shows the changes in simulated search averaged statistics when using FLARE w.r.t model scale from 8B to 100B+. Hallucinations (Halluc.) refer to facts and predicates only used in trace, while unutilised knowledge (Unutil.) relates to the facts and relations only seen in the code.

reasoning paths that lead to correct or incorrect outcomes. Our results in Table 4 show that LLM performance and reasoning optimally are not directly connected to the amount of explored paths or multi-hop inferences per path. We also see that traces that lead to incorrect answers have a higher number of failures per path and in total. We explain this phenomenon with the hypothesis that LLMs with traces that were optimal for reasoning and led to correct answers could skip exploring degenerate solutions due to strong commonsense reasoning capabilities. Further analyses focus on identifying inconsistencies and failure modes (Section 3.2). By comparing relations in code with those in search traces, we measure emergent hallucinations and unused relations, highlighting areas

Figure 4: The effect of the model parameter scale from 8B to 100B+ on model accuracy (left) and faithfulness (right).

of sub-optimal reasoning. Additionally, we assess the uniqueness of emergent facts per inference hop, which indicates the extent of problem-space exploration (Table 5). The results in Table 5 show consistently over each model that, on average, traces that lead to correct answers had a higher percentage of unique emergent facts (UEF) and overlap in the relations (OR) used between the code and search, while the portion of underutilized relations (UR) was lower. This means that optimal reasoning with an LLM requires a great degree of problem-space exploration with fewer relation hallucinations during the search and more relation utilization from the defined code. This aligns with our prior discoveries, which show a strong correlation between simulated search faithfulness towards the formalised code and model performance. Our framework FLARE has these reasoning patterns ingrained within its inference pipeline.

5.5 THE EFFECT OF SCALE

We want to assess the impact of the number of parameters in the model on the overall performance and faithfulness. The results in Figure 4 show no precise relation between model scale, performance and faithfulness. However, scaled models from the same family, i.e. CmDR (30B) and CmDR+ (100B), show improvements in reasoning faithfulness and model performance. We can also see in Table 6 that as the model size increases, the average number of hops and the portion of hallucinations and unutilised knowledge decreases. This further confirms our prior assumptions that models with strong commonsense soft-reasoning capabilities can skip steps during the search while maintaining the knowledge and structure of the traversal strategy outlined in the code.

6 CONCLUSION

This work introduces FLARE, a novel approach for logic-aided interpretable formalisation and reasoning with simulated search over the problem space. We show that models of varying scales obtain state-of-the-art results compared to prompting paradigms like CoT and F-CoT. We further pinpoint that using FLARE allows us to perform soft-reasoning with simulated search, making it flexible for diverse reasoning benchmarks. We introduce a method to measure model reasoning faithfulness w.r.t. the problem formalization ingrained within FLARE. Our results show that model performance is positively correlated with the faithfulness of the reasoning process. The systematic studies of the method show the benefits of using simulated search compared to natural language reasoning and external symbolic solvers. We further show that using FLARE allows us to interpretably and rigorously detect hallucinations and sub-optimal and inconsistent reasoning patterns.

REPRODUCIBILITY REPORT

To reproduce the results of our study, we provide the complete codebase, processing pipelines and prompts for each dataset. The only model hyper-parameter we explicitly fix is the temperature for greedy decoding. We also make the inference of all of the models using FLARE, F-CoT and CoT across all of the datasets publicly available for further experimentation and exploration.

REFERENCES

Viraat Aryabumi, Yixuan Su, Raymond Ma, Adrien Morisot, Ivan Zhang, Acyr Locatelli, Marzieh Fadaee, Ahmet Üstün, and Sara Hooker. To code, or not to code? exploring impact of code in pre-training. *CoRR*, abs/2408.10914, 2024. doi: 10.48550/ARXIV.2408.10914. URL https://doi.org/10.48550/arXiv.2408.10914.

Jacob Austin, Augustus Odena, Maxwell I. Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie J. Cai, Michael Terry, Quoc V. Le, and Charles Sutton. Program synthesis with large language models. *CoRR*, abs/2108.07732, 2021. URL https://arxiv.org/abs/2108.07732.

BIG bench authors. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL https://openreview.net/forum?id=uyTL5Bvosj.

Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michal Podstawski, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczyk, and Torsten Hoefler. Graph of thoughts: Solving elaborate problems with large language models. In Michael J. Wooldridge, Jennifer G. Dy, and Sriraam Natarajan (eds.), *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada*, pp. 17682–17690. AAAI Press, 2024. doi: 10.1609/AAAI.V38I16.29720. URL https://doi.org/10.1609/aaai.v38i16.29720.

Kenneth A. Bowen. Prolog. In *Proceedings of the 1979 Annual Conference, Detroit, Michigan, USA, October 29-31, 1979*, pp. 14–23. ACM, 1979. doi: 10.1145/800177.810020. URL https://doi.org/10.1145/800177.810020.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 1877–1901. Curran Associates, Inc., 2020a. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020b. URL https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfcb4967418bfb8ac142f64a-Abstract.html.

Hyungjoo Chae, Yeonghyeon Kim, Seungone Kim, Kai Tzu-iunn Ong, Beong-woo Kwak, Moohyeon Kim, Seonghwan Kim, Taeyoon Kwon, Jiwan Chung, Youngjae Yu, and Jinyoung Yeo. Language models as compilers: Simulating pseudocode execution improves algorithmic reasoning in language models. *CoRR*, abs/2404.02575, 2024. doi: 10.48550/ARXIV.2404.02575. URL https://doi.org/10.48550/arXiv.2404.02575.

Ashok K. Chandra and David Harel. Horn clauses queries and generalizations. *J. Log. Program.*, 2 (1):1–15, 1985. doi: 10.1016/0743-1066(85)90002-0. URL https://doi.org/10.1016/0743-1066(85)90002-0.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Pondé de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Joshua Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code. *CoRR*, abs/2107.03374, 2021. URL https://arxiv.org/abs/2107.03374.

Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W. Cohen. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks. *Trans. Mach. Learn. Res.*, 2023, 2023. URL https://openreview.net/forum?id=YfZ4ZPt8zd.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *CoRR*, abs/2110.14168, 2021. URL https://arxiv.org/abs/2110.14168.

Cohere. Command r: Retrieval-augmented generation at production scale. https://txt.cohere.com/command-r, March 2024.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurélien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Rozière, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Grégoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel M. Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, and et al. The llama 3 herd of models. *CoRR*, abs/2407.21783, 2024. doi: 10.48550/ARXIV.2407.21783. URL https://doi.org/10.48550/arXiv.2407.21783.

Aarohi Srivastava et al. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *Trans. Mach. Learn. Res.*, 2023, 2023. URL https://openreview.net/forum?id=uyTL5Bvosj.

Kanishk Gandhi, Denise Lee, Gabriel Grand, Muxin Liu, Winson Cheng, Archit Sharma, and Noah D. Goodman. Stream of search (sos): Learning to search in language. *CoRR*,

abs/2404.03683, 2024. doi: 10.48550/ARXIV.2404.03683. URL https://doi.org/10.48550/arXiv.2404.03683.

Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. PAL: program-aided language models. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pp. 10764–10799. PMLR, 2023. URL https://proceedings.mlr.press/v202/gao23f.html.

Sebastian Gehrmann, Tosin P. Adewumi, Karmanya Aggarwal, Pawan Sasanka Ammanamanchi, Aremu Anuoluwapo, Antoine Bosselut, Khyathi Raghavi Chandu, Miruna-Adriana Clinciu, Dipanjan Das, Kaustubh D. Dhole, Wanyu Du, Esin Durmus, Ondrej Dusek, Chris Emezue, Varun Gangal, Cristina Garbacea, Tatsunori Hashimoto, Yufang Hou, Yacine Jernite, Harsh Jhamtani, Yangfeng Ji, Shailza Jolly, Dhruv Kumar, Faisal Ladhak, Aman Madaan, Mounica Maddela, Khyati Mahajan, Saad Mahamood, Bodhisattwa Prasad Majumder, Pedro Henrique Martins, Angelina McMillan-Major, Simon Mille, Emiel van Miltenburg, Moin Nadeem, Shashi Narayan, Vitaly Nikolaev, Rubungo Andre Niyongabo, Salomey Osei, Ankur P. Parikh, Laura Perez-Beltrachini, Niranjan Ramesh Rao, Vikas Raunak, Juan Diego Rodriguez, Sashank Santhanam, João Sedoc, Thibault Sellam, Samira Shaikh, Anastasia Shimorina, Marco Antonio Sobrevilla Cabezudo, Hendrik Strobelt, Nishant Subramani, Wei Xu, Diyi Yang, Akhila Yerukola, and Jiawei Zhou. The GEM benchmark: Natural language generation, its evaluation and metrics. *CoRR*, abs/2102.01672, 2021. URL https://arxiv.org/abs/2102.01672.

Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. Did aristotle use a laptop? A question answering benchmark with implicit reasoning strategies. *Trans. Assoc. Comput. Linguistics*, 9:346–361, 2021. doi: 10.1162/TACL_A_00370. URL https://doi.org/10.1162/tacl_a_00370.

Leilani H. Gilpin, David Bau, Ben Z. Yuan, Ayesha Bajwa, Michael A. Specter, and Lalana Kagal. Explaining explanations: An overview of interpretability of machine learning. In Francesco Bonchi, Foster J. Provost, Tina Eliassi-Rad, Wei Wang, Ciro Cattuto, and Rayid Ghani (eds.), *5th IEEE International Conference on Data Science and Advanced Analytics, DSAA 2018, Turin, Italy, October 1-3, 2018*, pp. 80–89. IEEE, 2018. doi: 10.1109/DSAA.2018.00018. URL https://doi.org/10.1109/DSAA.2018.00018.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In Joaquin Vanschoren and Sai-Kit Yeung (eds.), *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*, 2021a. URL https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/be83ab3ecd0db773eb2dc1b0a17836a1-Abstract-round2.html.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*, 2021b.

Bernease Herman. The promise and peril of human evaluation for model interpretability. *arXiv preprint arXiv:1711.07414*, 2017.

Alon Jacovi and Yoav Goldberg. Towards faithfully interpretable NLP systems: How should we define and evaluate faithfulness? In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel R. Tetreault (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pp. 4198–4205. Association for Computational Linguistics, 2020. doi: 10.18653/V1/2020.ACL-MAIN.386. URL https://doi.org/10.18653/v1/2020.acl-main.386.

Alon Jacovi, Yonatan Bitton, Bernd Bohnet, Jonathan Herzig, Or Honovich, Michael Tseng, Michael Collins, Roee Aharoni, and Mor Geva. A chain-of-thought is as strong as its weakest link: A benchmark for verifiers of reasoning chains. In Lun-Wei Ku, Andre Martins, and

Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pp. 4615–4634. Association for Computational Linguistics, 2024. URL https://aclanthology.org/2024.acl-long.254.

Juyong Jiang, Fan Wang, Jiasi Shen, Sungju Kim, and Sunghun Kim. A survey on large language models for code generation. *CoRR*, abs/2406.00515, 2024. doi: 10.48550/ARXIV.2406.00515. URL https://doi.org/10.48550/arXiv.2406.00515.

Zhanming Jie, Trung Quoc Luong, Xinbo Zhang, Xiaoran Jin, and Hang Li. Design of chain-of-thought in math problem solving. *arXiv preprint arXiv:2309.11054*, 2023.

Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/8bb0d291acd4acf06ef112099c16f326-Abstract-Conference.html.

Tamera Lanham, Anna Chen, Ansh Radhakrishnan, Benoit Steiner, Carson Denison, Danny Hernandez, Dustin Li, Esin Durmus, Evan Hubinger, Jackson Kernion, Kamile Lukosiute, Karina Nguyen, Newton Cheng, Nicholas Joseph, Nicholas Schiefer, Oliver Rausch, Robin Larson, Sam McCandlish, Sandipan Kundu, Saurav Kadavath, Shannon Yang, Thomas Henighan, Timothy Maxwell, Timothy Telleen-Lawton, Tristan Hume, Zac Hatfield-Dodds, Jared Kaplan, Jan Brauner, Samuel R. Bowman, and Ethan Perez. Measuring faithfulness in chain-of-thought reasoning. *CoRR*, abs/2307.13702, 2023a. doi: 10.48550/ARXIV.2307.13702. URL https://doi.org/10.48550/arXiv.2307.13702.

Tamera Lanham, Anna Chen, Ansh Radhakrishnan, Benoit Steiner, Carson Denison, Danny Hernandez, Dustin Li, Esin Durmus, Evan Hubinger, Jackson Kernion, et al. Measuring faithfulness in chain-of-thought reasoning. *arXiv preprint arXiv:2307.13702*, 2023b.

Lucas Lehnert, Sainbayar Sukhbaatar, Paul McVay, Michael Rabbat, and Yuandong Tian. Beyond a*: Better planning with transformers via search dynamics bootstrapping. *CoRR*, abs/2402.14083, 2024. doi: 10.48550/ARXIV.2402.14083. URL https://doi.org/10.48550/arXiv.2402.14083.

Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay V. Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. Solving quantitative reasoning problems with language models. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/18abbeef8cfe9203fdf9053c9c4fe191-Abstract-Conference.html.

Chengshu Li, Jacky Liang, Andy Zeng, Xinyun Chen, Karol Hausman, Dorsa Sadigh, Sergey Levine, Li Fei-Fei, Fei Xia, and Brian Ichter. Chain of code: Reasoning with a language model-augmented code emulator. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024. URL https://openreview.net/forum?id=vKtomqlSxm.

Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pp. 74–81, Barcelona, Spain, July 2004. Association for Computational Linguistics. URL https://aclanthology.org/W04-1013.

Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. Program induction by rationale generation: Learning to solve and explain algebraic word problems. In Regina Barzilay and Min-Yen Kan (eds.), *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pp. 158–167. Association for Computational Linguistics, 2017. doi: 10.18653/V1/P17-1015. URL https://doi.org/10.18653/v1/P17-1015.

Zhan Ling, Yunhao Fang, Xuanlin Li, Zhiao Huang, Mingu Lee, Roland Memisevic, and Hao Su. Deductive verification of chain-of-thought reasoning. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/72393bd47a35f5b3bee4c609e7bba733-Abstract-Conference.html.

Bo Liu, Yuqian Jiang, Xiaohan Zhang, Qiang Liu, Shiqi Zhang, Joydeep Biswas, and Peter Stone. LLM+P: empowering large language models with optimal planning proficiency. *CoRR*, abs/2304.11477, 2023. doi: 10.48550/ARXIV.2304.11477. URL https://doi.org/10.48550/arXiv.2304.11477.

Fang Liu, Yang Liu, Lin Shi, Houkun Huang, Ruifeng Wang, Zhen Yang, and Li Zhang. Exploring and evaluating hallucinations in llm-powered code generation. *CoRR*, abs/2404.00971, 2024. doi: 10.48550/ARXIV.2404.00971. URL https://doi.org/10.48550/arXiv.2404.00971.

John W. Lloyd. Practical advtanages of declarative programming. In María Alpuente, Roberto Barbuti, and Isidro Ramos (eds.), *1994 Joint Conference on Declarative Programming, GULP-PRODE'94 Peñiscola, Spain, September 19-22, 1994, Volume 1*, pp. 18–30, 1994.

Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. Faithful chain-of-thought reasoning. In Jong C. Park, Yuki Arase, Baotian Hu, Wei Lu, Derry Wijaya, Ayu Purwarianti, and Adila Alfa Krisnadhi (eds.), *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics, IJCNLP 2023 -Volume 1: Long Papers, Nusa Dua, Bali, November 1 - 4, 2023*, pp. 305–329. Association for Computational Linguistics, 2023. doi: 10.18653/V1/2023.IJCNLP-MAIN.20. URL https://doi.org/10.18653/v1/2023.ijcnlp-main.20.

Aman Madaan and Amir Yazdanbakhsh. Text and patterns: For effective chain of thought, it takes two to tango. *CoRR*, abs/2209.07686, 2022a. doi: 10.48550/ARXIV.2209.07686. URL https://doi.org/10.48550/arXiv.2209.07686.

Aman Madaan and Amir Yazdanbakhsh. Text and patterns: For effective chain of thought, it takes two to tango. *arXiv preprint arXiv:2209.07686*, 2022b.

Aman Madaan, Shuyan Zhou, Uri Alon, Yiming Yang, and Graham Neubig. Language models of code are few-shot commonsense learners. *arXiv preprint arXiv:2210.07128*, 2022.

Shen-Yun Miao, Chao-Chun Liang, and Keh-Yih Su. A diverse corpus for evaluating and developing english math word problem solvers. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel R. Tetreault (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pp. 975–984. Association for Computational Linguistics, 2020. doi: 10.18653/V1/2020.ACL-MAIN.92. URL https://doi.org/10.18653/v1/2020.acl-main.92.

Erik Nijkamp, Bo Pang, Hiroaki Hayashi, Lifu Tu, Huan Wang, Yingbo Zhou, Silvio Savarese, and Caiming Xiong. Codegen: An open large language model for code with multi-turn program synthesis. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL https://openreview.net/forum?id=iaYcJKpY2B_.

Liangming Pan, Alon Albalak, Xinyi Wang, and William Yang Wang. Logic-lm: Empowering large language models with symbolic solvers for faithful logical reasoning. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pp. 3806–3824. Association for Computational Linguistics, 2023. doi: 10.18653/V1/2023.FINDINGS-EMNLP.248. URL https://doi.org/10.18653/v1/2023.findings-emnlp.248.

Arkil Patel, Satwik Bhattamishra, and Navin Goyal. Are NLP models really able to solve simple math word problems? In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tür, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou (eds.), *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pp. 2080–2094. Association for Computational Linguistics, 2021. doi: 10.18653/V1/2021.NAACL-MAIN.168. URL https://doi.org/10.18653/v1/2021.naacl-main.168.

Emily Reif, Daphne Ippolito, Ann Yuan, Andy Coenen, Chris Callison-Burch, and Jason Wei. A recipe for arbitrary text style transfer with large language models. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (eds.), *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pp. 837–848. Association for Computational Linguistics, 2022. doi: 10.18653/V1/2022.ACL-SHORT.94. URL https://doi.org/10.18653/v1/2022.acl-short.94.

Subhro Roy and Dan Roth. Solving general arithmetic word problems. In Lluís Màrquez, Chris Callison-Burch, Jian Su, Daniele Pighin, and Yuval Marton (eds.), *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pp. 1743–1752. The Association for Computational Linguistics, 2015. doi: 10.18653/V1/D15-1202. URL https://doi.org/10.18653/v1/d15-1202.

Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal V. Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Févry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M. Rush. Multitask prompted training enables zero-shot task generalization. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. URL https://openreview.net/forum?id=9Vrb9D0WI4.

Abulhair Saparov and He He. Language models are greedy reasoners: A systematic formal analysis of chain-of-thought. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL https://openreview.net/forum?id=qFVVBzXxR2V.

Koustuv Sinha, Shagun Sodhani, Jin Dong, Joelle Pineau, and William L. Hamilton. CLUTRR: A diagnostic benchmark for inductive reasoning from text. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pp. 4505–4514. Association for Computational Linguistics, 2019. doi: 10.18653/V1/D19-1458. URL https://doi.org/10.18653/v1/D19-1458.

Zayne Sprague, Fangcong Yin, Juan Diego Rodriguez, Dongwei Jiang, Manya Wadhwa, Prasann Singhal, Xinyu Zhao, Xi Ye, Kyle Mahowald, and Greg Durrett. To cot or not to cot? chain-of-thought helps mainly on math and symbolic reasoning. *arXiv preprint arXiv:2409.12183*, 2024.

Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc Le, Ed Chi, Denny Zhou, and Jason Wei. Challenging BIG-bench tasks and whether chain-of-thought can solve them. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 13003–13051, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-acl.824. URL https://aclanthology.org/2023.findings-acl.824.

Christian Szegedy. A promising path towards autoformalization and general artificial intelligence. In Christoph Benzmüller and Bruce R. Miller (eds.), *Intelligent Computer Mathematics - 13th*

International Conference, CICM 2020, Bertinoro, Italy, July 26-31, 2020, Proceedings, volume 12236 of *Lecture Notes in Computer Science*, pp. 3–20. Springer, 2020. doi: 10.1007/978-3-030-53518-6_1. URL https://doi.org/10.1007/978-3-030-53518-6_1.

Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/ed3fea9033a80fea1376299fa7863f4a-Abstract-Conference.html.

Karthik Valmeekam, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. Large language models still can't plan (a benchmark for llms on planning and reasoning about change). In *NeurIPS 2022 Foundation Models for Decision Making Workshop*, 2022.

Lei Wang, Wanyu Xu, Yihuai Lan, Zhiqiang Hu, Yunshi Lan, Roy Ka-Wei Lee, and Ee-Peng Lim. Plan-and-solve prompting: Improving zero-shot chain-of-thought reasoning by large language models. In Anna Rogers, Jordan L. Boyd-Graber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pp. 2609–2634. Association for Computational Linguistics, 2023a. doi: 10.18653/V1/2023.ACL-LONG.147. URL https://doi.org/10.18653/v1/2023.acl-long.147.

Qingxiang Wang, Cezary Kaliszyk, and Josef Urban. First experiments with neural translation of informal to formal mathematics. In Florian Rabe, William M. Farmer, Grant O. Passmore, and Abdou Youssef (eds.), *Intelligent Computer Mathematics - 11th International Conference, CICM 2018, Hagenberg, Austria, August 13-17, 2018, Proceedings*, volume 11006 of *Lecture Notes in Computer Science*, pp. 255–270. Springer, 2018. doi: 10.1007/978-3-319-96812-4_22. URL https://doi.org/10.1007/978-3-319-96812-4_22.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2022.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023b. URL https://openreview.net/forum?id=1PL1NIMMrw.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.

Jan Wielemaker, Tom Schrijvers, Markus Triska, and Torbjörn Lager. Swi-prolog. *Theory and Practice of Logic Programming*, 12(1-2):67–96, 2012.

Wikipedia. Countdown (game show) — Wikipedia, the free encyclopedia. http://en.wikipedia.org/w/index.php?title=Countdown%20(game%20show)&oldid=1248084922, 2024. [Online; accessed 09-September-2024].

Yuhuai Wu, Albert Qiaochu Jiang, Wenda Li, Markus N. Rabe, Charles Staats, Mateja Jamnik, and Christian Szegedy. Autoformalization with large language models. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/d0c6bc641a56bebee9d985b937307367-Abstract-Conference.html.

Mengjiao Yang, Dale Schuurmans, Pieter Abbeel, and Ofir Nachum. Chain of thought imitation with procedure cloning. In Sanmi Koyejo, S. Mohamed, A. Agarwal,

Danielle Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/ebdb990471f653dffb425eff03c7c980-Abstract-Conference.html.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*, 2022.

Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023a. URL http://papers.nips.cc/paper_files/paper/2023/hash/271db9922b8d1f4dd7aaef84ed5ac703-Abstract-Conference.html.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R. Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023b. URL https://openreview.net/forum?id=WE_vluYUL-X.

Junjie Ye, Xuanting Chen, Nuo Xu, Can Zu, Zekai Shao, Shichun Liu, Yuhan Cui, Zeyang Zhou, Chao Gong, Yang Shen, Jie Zhou, Siming Chen, Tao Gui, Qi Zhang, and Xuanjing Huang. A comprehensive capability analysis of GPT-3 and GPT-3.5 series models. *CoRR*, abs/2303.10420, 2023. doi: 10.48550/ARXIV.2303.10420. URL https://doi.org/10.48550/arXiv.2303.10420.

Yadong Zhang, Shaoguang Mao, Tao Ge, Xun Wang, Adrian de Wynter, Yan Xia, Wenshan Wu, Ting Song, Man Lan, and Furu Wei. LLM as a mastermind: A survey of strategic reasoning with large language models. *CoRR*, abs/2404.01230, 2024. doi: 10.48550/ARXIV.2404.01230. URL https://doi.org/10.48550/arXiv.2404.01230.

Wanjun Zhong, Siyuan Wang, Duyu Tang, Zenan Xu, Daya Guo, Jiahai Wang, Jian Yin, Ming Zhou, and Nan Duan. AR-LSAT: investigating analytical reasoning of text. *CoRR*, abs/2104.06598, 2021. URL https://arxiv.org/abs/2104.06598.

Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V. Le, and Ed H. Chi. Least-to-most prompting enables complex reasoning in large language models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL https://openreview.net/forum?id=WZH7099tgfM.

A   APPENDIX

A.1   LLM PROMPTS

We define straight-forward prompts for generating *plan*, *code* and *search* simulation in FLARE, which can be observed in Appendix A.2.

A.2   DATASET STATISTICS

The datasets used in this study encompass a variety of domains, specifically targeting the performance of the models in interpreting Math Word Problems, multi-hop question answering, and relational inference. Table 9 provides a detailed breakdown of each dataset, including the number of few-shot in-context samples (shots), the number of test samples, and representative examples from each dataset. The datasets provide a comprehensive basis for evaluating the models' abilities to handle complex tasks across different domains, facilitating an in-depth analysis of model performance under few-shot conditions.

Task	Prompt	Description
Plan Generation	Generate an explanation and analysis, and plan to generate a prompt for writing a swi-prolog code for the last task. The 3 sections should be exactly outlined. Your plan should show enough intermediate reasoning steps towards the answer. Construct the plan as much as you can and describe the logic specifically. When constructing the plan for the code prompt, actively use swi prolog search capabilities.	Detailed instructions for generating an outline and plan, with an emphasis on reasoning steps and using Prolog's search capabilities.
Code Generation	Write a Prolog code to solve using the plan. If there are unknown or stochastic atoms or predicates, fill in the values for them as a logical assumption and add a comment in the same line Assumed atom/predicate". Do not use write and read commands within the code. The code should be very detailed and utilize swi prolog capabilities to the fullest. To run the program, at the end create a predicate named "query" that returns the correct numerical answer. The last line of the program should be the commented-out driver predicate "query". Write only the code.	Instructions for generating a Prolog code based on the plan with assumptions for unknown atoms. Emphasizes code details and a final "query" predicate.
Simulated Search	Ignoring the read commands, explicitly write out the search paths that are explored by the code: #### Here are the paths [Starting Search Simulation]: #### [Path 1]:	A task to simulate and display the search paths that the Prolog code would follow during execution.
Final Answer	Given the plan, the code and the explored search paths answer the question above. Answer with the correct numerical answer. ##### Here is the answer:	Final prompt asking for the correct numerical answer based on the previous steps.

Table 7: Table of Prompts for Plan, Code, Simulated Search, and Final Answer generation for GSM8K (Cobbe et al., 2021).

Section	Content
Question	"Can Reiki be stored in a bottle?"
Explanation	Reiki is a form of alternative therapy that involves the channeling of energy to promote healing and relaxation. It is often administered through the practitioner's hands, with the energy flowing directly to the recipient. The question raises the possibility of storing Reiki energy in a bottle, which may be a common query due to the abstract nature of energy healing.
Analysis	To address the question of whether Reiki can be stored in a bottle, it is important to consider the nature of Reiki energy, the principles of energy healing, and the practical implications of containing and preserving such energy.
Plan	- Understand Reiki Energy: ... - Explore Energy Storage: ... - Consider Symbolism and Intention: ... - Evaluate Practicality: ... - Reflect on Tradition: ...
Code	% Define the nature of Reiki energy reiki_energy(spiritual_energy).
reiki_energy(channeling_through_touch).
% Define the feasibility of storing energy in a bottle
store_energy_in_bottle(feasible) :-
reiki_energy(spiritual_energy),
reiki_energy(channeling_through_touch).
% Define the assumption that storing Reiki energy in a bottle is not a common practice not_common_practice(store_reiki_in_bottle).
% Query to determine if Reiki can be stored in a bottle
query :-
store_energy_in_bottle(feasible),
not_common_practice(store_reiki_in_bottle).
Search Paths	Path 1:
– 0: Start of execution: Beginning Search
– 1: Call: query
– 2: Call: store_energy_in_bottle(feasible)
– 3: Call: reiki_energy(spiritual_energy)
– 4: Call: reiki_energy(channeling_through_touch)
– 5: Call: not_common_practice(store_reiki_in_bottle)
– 6: Call: feasible=no
– 7: Fail: feasible=no — 'Result': 'Search Failed'
– 8: Fail: query — 'Result': 'Search Failed'
Answer	"No, Reiki energy cannot be stored in a bottle based on the logical evaluation of its abstract, non-physical nature and traditional practices of Reiki."

Table 8: Complete example of FLARE

Domain	Dataset	Shots	Test Samples	Example
Math Word Problems	GSM8K	8	1,319	Q: A robe takes 2 bolts of blue fiber and half that much white fiber. How many bolts in total does it take?
A: 3
	SVAMP	8	1,000	Q: Dan had $3 left with him after he bought a candy bar. If he had $4 at the start, how much did the candy bar cost?A: 1
	MultiArith	8	600	Q: A pet store had 13 siamese cats and 5 house cats. During a sale they sold 10 cats. How many cats do they have left?
A: 8
	ASDiv	8	2,096	Q: Adam has five more apples than Jackie. Jackie has nine apples. How many apples does Adam have?
A: 14
	AQuA	8	254	Q: A man walks at 5 kmph for 6 hrs and at 4 kmph for 12 hrs. His average speed is
Answer option: A)4 1/3 km/h, B)7 2/3 km/h, C)9 ½ km/h, D)8 km/h, E)81 km/h
A: A
Multi-hop QA	StrategyQA	6	2,290	Q: Did Aristotle use a laptop?
A: False
	Date Understanding	10	359	Q: Yesterday was April 30, 2021. What is the date tomorrow in MM/DD/YYYY?
A: "05/02/2021"
	Sports Understanding	10	977	Q: Is the following sentence plausible? Lionel Messi was called for icing?
A: False
Relational Inference	CLUTRR	8	1,042	Q: [Carlos] is [Clarence]'s brother. [Carlos] and his sister, [Annie], went shopping. asked her mom [Valerie] if she wanted anything, but [Valerie] said no. How is [Valerie] related to [Clarence]?
A: "mother"

Table 9: The statistics and examples of the datasets used in benchmarking. Shots refers to the number of few-shot in-context samples used during benchmarking.