

LLMs for Relational Reasoning: How Far are We?

Zhiming Li

zhiming001@e.ntu.edu.sg
Continental-NTU Corporate Lab,
Nanyang Technological University
Singapore

Yushi Cao*

yushi002@e.ntu.edu.sg
Continental-NTU Corporate Lab,
Nanyang Technological University
Singapore

Xiufeng Xu

xiufeng001@e.ntu.edu.sg
Nanyang Technological University
Singapore

Junzhe Jiang

junzhe.jiang@connect.polyu.hk
Hong Kong Polytechnic University
Hong Kong

Xu Liu

liuxu@comp.nus.edu.sg
National University of Singapore
Singapore

Yon Shin Teo

yon.shin.teo@continental-
corporation.com
Continental Automotive Singapore
Pte. Ltd.
Singapore

Shang-wei Lin

shang-wei.lin@ntu.edu.sg
Continental-NTU Corporate Lab,
Nanyang Technological University
Singapore

Yang Liu

yangliu@ntu.edu.sg
Continental-NTU Corporate Lab,
Nanyang Technological University
Singapore

ABSTRACT

Large language models (LLMs) have revolutionized many areas (e.g. natural language processing, software engineering, etc.) by achieving state-of-the-art performance on extensive downstream tasks. Aiming to achieve robust and general artificial intelligence, there has been a surge of interest in investigating the reasoning ability of the LLMs. Whereas the textual and numerical reasoning benchmarks adopted by previous works are rather shallow and simple, it is hard to conclude that the LLMs possess strong reasoning ability by merely achieving positive results on these benchmarks. Recent efforts have demonstrated that the LLMs are poor at solving sequential decision-making problems that require common-sense planning by evaluating their performance on the reinforcement learning benchmarks. In this work, we conduct an in-depth assessment of several state-of-the-art LLMs' reasoning ability based on the inductive logic programming (ILP) benchmark, which is broadly recognized as a representative and challenging measurement for

evaluating logic program induction/synthesis systems as it requires inducing strict *cause-effect* logic to achieve robust deduction on independent and identically distributed (IID) and out-of-distribution (OOD) test samples. Our evaluations illustrate that compared with the neural program induction systems which are much smaller in model size, the state-of-the-art LLMs are much poorer in terms of reasoning ability by achieving much lower performance and generalization using either natural language prompting or truth-value matrix prompting¹.

KEYWORDS

Large Language Models, Relational Reasoning, Program Induction

ACM Reference Format:

Zhiming Li, Yushi Cao, Xiufeng Xu, Junzhe Jiang, Xu Liu, Yon Shin Teo, Shang-wei Lin, and Yang Liu. 2024. LLMs for Relational Reasoning: How Far are We?. *ACM/IMS J. Data Sci.* 37, 4, Article 111 (August 2024), 8 pages. <https://doi.org/10.1145/3643795.3648387>

1 INTRODUCTION

Large language models (LLMs) have achieved great breakthroughs in various domains such as natural language processing [2, 28], software engineering [7, 16], finance [32, 33], etc. There has been a recent increase in interest in exploring the reasoning ability of LLMs, which is regarded as a crucial ability of Artificial General Intelligence (AGI) [17, 20, 31]. Prevalent reasoning ability evaluation benchmarks adopted by previous literature include arithmetic [5, 23], symbolic reasoning [25, 31], commonsense [24, 26] etc. And with the recent advancement of the in-context few-shot (zero-shot) *prompting* techniques [2, 19, 31], LLMs manage to achieve state-of-the-art few-shot (zero-shot) learning performance on these benchmarks without training. Concisely, in-context few-shot (zero-shot) *prompting* refers to the techniques that provide input to the

*Corresponding author

Authors' addresses: Zhiming Li, zhiming001@e.ntu.edu.sg, Continental-NTU Corporate Lab, Nanyang Technological University, Singapore; Yushi Cao, yushi002@e.ntu.edu.sg, Continental-NTU Corporate Lab, Nanyang Technological University, Singapore; Xiufeng Xu, xiufeng001@e.ntu.edu.sg, Nanyang Technological University, Singapore; Junzhe Jiang, junzhe.jiang@connect.polyu.hk, Hong Kong Polytechnic University, Hong Kong; Xu Liu, liuxu@comp.nus.edu.sg, National University of Singapore, Singapore; Yon Shin Teo, yon.shin.teo@continental-corporation.com, Continental Automotive Singapore Pte. Ltd., Singapore; Shang-wei Lin, shang-wei.lin@ntu.edu.sg, Continental-NTU Corporate Lab, Nanyang Technological University, Singapore; Yang Liu, yangliu@ntu.edu.sg, Continental-NTU Corporate Lab, Nanyang Technological University, Singapore.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2024 Copyright held by the owner/author(s).

2831-3194/2024/8-ART111

<https://doi.org/10.1145/3643795.3648387>

¹The implementation is available at: <https://sites.google.com/view/llm-rr>

language model to boost the performance of specific tasks. The input can be a few examples (few-shot) [3, 31] or merely instructions that describe the task (zero-shot) [30]. Despite the stunning achievement of the LLMs on these reasoning benchmarks. There are recent debate that these benchmarks are relatively simple in terms of task-solving logic and only require shallow reasoning to accomplish [29]. Thus it is insufficient to support the claims about LLMs' reasoning ability. Specifically, Valmeekam *et al.* [29] conducted evaluations of the LLMs' inherent emergent planning on a reinforcement learning benchmark: *blocksworld*. The *blocksworld* is a popular benchmark for evaluating reinforcement learning baselines in terms of sequential decision-making ability. Concretely, given a set of blocks, the goal of this task is to arrange the blocks in a particular order. While this task is simple enough for humans to solve, it is found that even some of the current state-of-the-art LLMs present poor performance on them.

Going beyond the evaluation of LLMs' sequential decision-making abilities on reinforcement learning benchmarks, we emphasize that the relational reasoning ability is another crucial ability to focus on. It is considered a crucial component of intelligence that directly correlates with the capacity to think logically and solve problems in novel situations [6, 14]. Specifically, relational reasoning ability is the ability to reason about relationships between objects.

The inductive logic programming (ILP) benchmarks [9, 10, 34] are broadly used for the evaluation of program induction/synthesis systems' relational reasoning ability. In specific, ILP is a task that aims to automatically induce a logic program given some positive examples and negative examples as specifications [10]. In this work, to achieve comprehensive evaluations of the LLMs' relational reasoning ability, we develop a universal evaluation pipeline that enables detailed evaluation of both the state-of-the-art LLMs and the neural program induction baseline methods which are dedicatedly designed for relational reasoning. To the best of our knowledge, we are the first to conduct a detailed analysis of the LLMs' relational reasoning ability as well as their comparison with the neural program induction models. We have the following key findings:

- We implement a universal evaluation pipeline for relational reasoning ability assessment, which is general for evaluating the state-of-the-art LLMs and the neural program induction models.
- We conduct comprehensive evaluations of LLMs' relational reasoning ability, which is the first of its kind.
- We unveil that the current state-of-the-art LLMs' relational reasoning ability is still far from perfect and is inferior compared to neural program induction models which are much smaller.

2 PRELIMINARIES

2.1 Inductive Logic Programming

Program induction & program synthesis [1, 8, 9] are tasks that aim to induce a program that satisfies a given specification. Program induction requires the model to induce the program implicitly and conduct inference directly by the model itself without executing an explicit program, while program synthesis requires the model to generate an explicit program and execute it for inference. Inductive logic programming (ILP) [11, 18, 22] is a sub-field in which specification comes in the form of input-output examples. It requires

the model to express the program with first-order logic rules. ILP has long been considered a task that is demanding in logical (relational) reasoning ability to accomplish. Specifically, given a set of background predicates \mathbf{B} (also called *premise*), a set of positive examples ρ , and a set of negative samples η , the goal of a program induction/synthesis model is to derive a logic program C that satisfies the following two conditions: (1) based on the premise \mathbf{B} , C entails all the positive examples, denoted by $\forall \rho : \mathbf{B}, C \models \rho$, and (2) based on the premises, C does not entail any of the negative examples, denoted by $\forall \eta : \mathbf{B}, C \not\models \eta$.

Differentiable logic machines (DLM) [34] model is a state-of-the-art neural program induction model for the ILP tasks that realizes first-order logic (FOL) rules in a neural manner. Specifically, the DLM model uses neural networks as computation units that implement "soft" logic operators. Then based on the computation units, the DLM model can approximate the forward-chaining mechanism of FOL by stacking multiple layers of computation units. It takes the truth value matrices of the background predicates as input and outputs the truth value matrix of the target predicate. With the strong inductive bias of the model design, DLM has been proven to achieve superior performance and generalization on benchmarks that demand strong reasoning ability such as relational reasoning and decision-making tasks.

2.2 Large Language Models

Large language models (LLMs) denote the pre-trained Transformer architecture-based AI models. By leveraging large amounts of multi-modal data and the pre-training & fine-tuning learning techniques, the LLMs are reported to achieve state-of-the-art performance on many downstream tasks (*e.g.* machine translation, numerical reasoning, code clone detection, *etc.*). Specifically, the GPT model family (*e.g.* GPT-3.5, GPT-4, *etc.*) are pre-trained with an autoregressive language modeling objective. The formal definition of the autoregressive loss \mathcal{L}_{aut} is as follows:

$$\mathcal{L}_{aut} = - \sum_{t=2}^n \log p(y_t | y_{t-1}, \dots, y_1) \quad (1)$$

Given a sentence of n tokens, the model is trained to maximize the likelihood of the ground-truth token y_t of the current time step t based on its previous sequence y_{t-1}, \dots, y_1 . In-context few-shot (zero-shot) *prompting* is a recently proposed technique that conditions the LLM with some initial input to improve performance. The input can be a few examples (few-shot) [3, 31] or instructions regarding the task (zero-shot) [30].

3 RELATIONAL REASONING ABILITY EVALUATION PIPELINE

In this section, we introduce the universal relational reasoning ability evaluation pipeline for the LLMs and neural program induction models (NPIMs). The overview of the pipeline is shown in Figure 1, which contains four major components: (1) sample generator, (2) (truth value (tv), natural language (nl)) modality compiler, and (3) the evaluation module. The content of this section is arranged as follows: we first illustrate the details of the relational reasoning benchmark which serves as the backend of the sample generator (Section 3.1). Secondly, we introduce the sample generator which

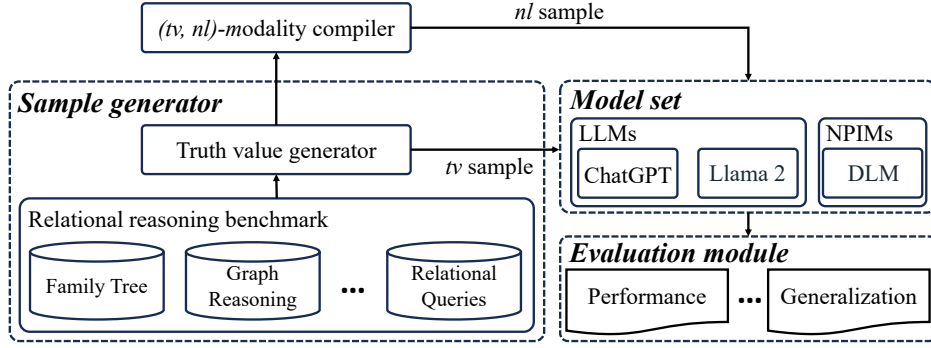


Figure 1: Overview of the evaluation pipeline.

generates random data points represented in the form of truth-value matrix prompting (Section 3.2). Then the details of the (tv, nl) -modality compiler are introduced. It transforms truth value prompting into corresponding natural language prompting to evaluate the models' relational reasoning ability using different data modalities (Section 3.3). Finally, we illustrate the evaluation module which allows measuring models' results from multiple aspects (Section 3.4).

3.1 Relational Reasoning Benchmark

There are many available relational reasoning benchmarks adopted by previous literature, *e.g.* family tree reasoning [9, 10], general graph reasoning [9, 13, 34], relational queries [27], *etc.* In this work, we conduct experiments on two broadly used relational reasoning benchmarks that are adopted by previous state-of-the-art neural program induction models [9, 10, 21, 34]: family tree reasoning and general graph reasoning.

Family tree reasoning. The family tree reasoning benchmark consists of tasks that require the model to induce programs that deduce more complex relations based on some basic properties of family members or relations between them. Specifically, a family tree is represented with four basic predicates: $IsMother(x, y)$, $IsSon(x, y)$, $IsDaughter(x, y)$. *E.g.* $IsMother(x, y)$ is True if y is x 's mother, the semantics of the other basic predicates are similar. This benchmark contains 5 target predicates to induce. The following shows the details of the $HasSister$ and $IsGrandparent$ tasks:

- $HasSister(x)$: the semantics of this predicate is to determine whether x has a sister. The ground-truth program to induce is:

$$HasSister(x) \leftarrow \exists y, z, IsDaughter(z, y) \wedge IsMother(x, z) \quad (2)$$

- $IsGrandparent(x, y)$: the semantics of this predicate is to determine whether y is the grandparent of x . The ground-truth program to induce is:

$$IsGrandparent(x, y) \leftarrow \exists z, ((IsSon(y, z) \wedge IsFather(x, z)) \vee (IsDaughter(y, z) \wedge IsMother(x, z))) \quad (3)$$

General graph reasoning. The general graph reasoning benchmark consists of tasks that require the models to infer the logic of high-level target predicates that describe properties/relations of a graph based on a basic predicate: $HasEdge(x, y)$ (*i.e.* whether

there is an undirected edge between node x and y in the graph). This benchmark contains 4 target predicates to infer. The following shows the details of the 4-Connectivity and 1-Outdegree tasks:

- 4-Connectivity(x, y): the semantics of 4-Connectivity(x, y) is to determine whether there exists a path between node x and node y within 4 edges. The ground-truth program to induce is:

$$4-Connectivity(x, y) \leftarrow \exists z, (HasEdge(x, y) \vee Invented(x, y) \vee (Invented(x, z) \wedge HasEdge(z, y)) \vee (Invented(x, z) \wedge Invented(z, y))) \quad (4)$$

$$Invented(x, y) \leftarrow \exists z, (HasEdge(x, z) \wedge HasEdge(z, y))$$

- 1-Outdegree(x): the semantics of this predicate is to determine whether the outdegree of node x in a graph is exactly 1. The ground-truth program to induce is:

$$1-Outdegree(x) \leftarrow \exists y, \forall z, (HasEdge(x, y) \wedge \neg HasEdge(x, z)) \quad (5)$$

3.2 Sample Generator

Given the specifications of each relational reasoning benchmark, the sample generator implements a truth value generator tool. Specifically, when it is called, the truth value generator tool generates the truth value matrices that represent the input & output relations of the desired program to be induced. Figure 2 shows a concrete truth value table representation of a $HasSister$ task sample. It represents a family with four members $\{P0, P1, P2, P3\}$. $P0$ is the mother, $P1$ is the father, $P2$ and $P3$ are the two daughters of them. The input prompting therefore contains the truth value matrices of the input basic predicate $IsMother(x, y)$, $IsDaughter(x, y)$ ². *E.g.* for the $IsMother(x, y)$, since $P1$ is the mother of both $P2$ and $P3$, the items within the matrix that represent $IsMother(P2, P1)$ and $IsMother(P3, P1)$ equal to 1. Then, given the input representation, the models are required to deduce the truth value matrix of the desired target predicate $HasSister$.

3.3 (tv, nl) -Modality Compiler

As the LLMs are natural language models whose primary training data source is natural language, we aim to evaluate their performance when the samples are represented in the form of natural language. Therefore, we implement a (tv, nl) -modality compiler to convert the sample from truth value prompting to natural language

²The matrices of $IsFather(x, y)$, $IsSon(x, y)$ are not presented for better illustration

$x y$	$P0$	$P1$	$P2$	$P3$
$P0$	0	0	0	0
$P1$	0	0	0	0
$P2$	0	1	0	0
$P3$	0	1	0	0

$x y$	$P0$	$P1$	$P2$	$P3$
$P0$	0	0	1	1
$P1$	0	0	1	1
$P2$	0	0	0	0
$P3$	0	0	0	0

$P0$	0
$P1$	0
$P2$	1
$P3$	1

$\text{IsMother}(x, y)$ $\text{IsDaughter}(x, y)$ $\text{HasSister}(x)$

Figure 2: Truth value prompting of a HasSister task sample. Given the truth values of the input predicates IsMother and IsDaughter, the model is required to deduce the results of the target predicate HasSister.

prompting. Specifically, the semantics of the predicates are directly conveyed via natural language prompting as background and the LLMs are then asked to answer by listing all the pairs that satisfy the target predicate. E.g. the example in Figure 2 is converted into the natural language prompting shown in Figure 3.

3.4 Evaluation Module

We implement the evaluation module for evaluations of results from both the truth value source and natural language source. This module allows evaluating test performance under the independent and identically distribution (IID) setting and generalization under the out-of-distribution (OOD) setting. Specifically, for the family tree benchmark, IID is defined as the test sample whose number of family members is the same as that of the training sample. OOD is defined as the test sample with a larger family size than that of the training sample. Similarly, for the general graph reasoning benchmark, IID refers to the test sample with the same number of graph nodes as the training data while OOD refers to those with a larger number of graph nodes.

4 EXPERIMENTAL SETTINGS

4.1 Data setup.

We illustrate the details of the data setup of the two benchmarks used in our study. For all the tasks of the family tree reasoning benchmark, the number of family members for the few-shot learning (training) data is set to 10. We use data samples with the same number of family members ($n = 10$) for IID performance evaluation; and samples with a larger number of family members ($n = 20$) for OOD generalization evaluation. Similarly, for the general graph reasoning benchmarks, the number of nodes of each graph for the few-shot learning is set to 10. The number of nodes of each IID test sample is the same as the training samples ($n = 10$), and each OOD test graph contains 20 ($n = 20$) nodes. For standard prompting, the number of training samples is set to 10, i.e., 10 families/graphs. The numbers of IID and OOD test samples are set to 10. For chain-of-thought prompting, due to the limitation of tokens, the number of training samples is set to one. The numbers of IID and OOD test samples are set to 10. As an example, for the IID test, 10 queries are conducted, each containing all the training samples and one unique testing sample. For the truth value prompting, we use Python list to represent the matrices. We conduct all experiments on a server with 48 cores of Intel(R) Xeon(R) Silver 4214 CPU @ 2.20GHz, 4 Quadro RTX 8000 GPUs, and 256G RAM.

$\text{HasSister}(x)$ <i>nl</i> prompt	
You are an agent who determines the relations in a family. I will first show you some examples, and then I will give you a new question for you to answer. The examples are shown as follows.	
<i>Question:</i>	For a family tree containing 6 family members, which is depicted with 4 kinds of relations: father, mother, son, daughter. The relations are: P0 is P2's father; P0 is P3's father; P1 is P2's mother; P1 is P3's mother; P2 is P0's daughter; P2 is P1's daughter; P3 is P0's daughter; P3 is P1's daughter. For a new relation sister, given the above-given facts, who has the relation sister? You must give the final answer in this format: Who has sister.
<i>Answer:</i>	P2 has sister; P3 has sister.
Now this is the question for you to answer.	
<i>Question:</i>	For a family tree containing 5 family members, which is depicted with 4 kinds of relations: father, mother, son, daughter. The relations are: P1 is P0's father; P2 is P4's father; P4 is P0's mother; P3 is P4's mother; P0 is P1's son; P0 is P4's son; P4 is P2's daughter; P4 is P3's daughter. For a new relation sister, from the above-given facts, who has relation sister? You must give the final answer in this format: Who has sister.

Figure 3: Natural language prompting of a HasSister task.

4.2 Model setup.

To comprehensively evaluate the relational reasoning abilities of the LLMs and NPI models, we choose five representative LLMs baseline models and one state-of-the-art NPI model. Specifically, for LLMs, we evaluate the GPT-3.5 Turbo, GPT-4, and GPT-4 Turbo, and the open-sourced Llama 2 (7B) and Llama 2 (13B) models. For the NPI, we evaluate the differentiable logic machines (DLM) model. The details of the models are as follows:

GPT-3.5 Turbo: The GPT-3.5 Turbo LLM is a Transformer-based decoder network with 175 billion parameters, which is pre-trained on 45 terabytes of text data from various sources, including books, articles, and websites. It is additionally fine-tuned on specific tasks such as language translation, summarization, and question-answering and allows a context window (i.e. maximum input sequence length) of up to 16,385 tokens.

GPT-4: GPT-4 LLM is a successor of GPT-3.5 with an estimated 100 trillion parameters. It is a multimodal model capable of analyzing text, images, and voice data. It is capable of longer sequences compared to the GPT-3.5 model, which is up to 8,192 tokens.

GPT-4 Turbo: GPT Turbo is an enhanced version of GPT-4, featuring a significantly expanded context window (32k tokens). This enables it to efficiently process a substantial volume of data in a single run. Additionally, GPT Turbo possesses up-to-date knowledge of global events until April 2023.

Llama 2: Llama 2 is an open-sourced pre-trained foundation model which adopts the standard Transformer architecture. It introduces using the Grouped Query Attention (GQA) and is capable of processing sequences with longer context length (4096 tokens) compared with the Llama 1 model (2048 tokens). It comes in four model sizes: 7B, 13B, 34B, and 70B. In this work, we evaluate the 7B and 13B models.

DLM: DLM is a neural program induction model (NPIM) architecture that approximates the inductive definition of logic formulas. DLM is the current state-of-the-art NPIM which achieves great performance and generalization of tasks ranging from relational

Table 1: Performance and generalization results of the LLMs using standard natural language prompting and DLM using few-shot training. Red number denotes the best performance/generalization result among all the LLMs for a task, number in the grey box denotes the best result among all the evaluated baselines.

Task	GPT-4 Turbo		GPT-4		GPT3.5 Turbo		llama-7b		llama-13b		DLM	
	n=10	n=20	n=10	n=20	n=10	n=20	n=10	n=20	n=10	n=20	n=10	n=20
HasFather	100.00%	99.17%	100.00%	99.60%	77.30%	69.52%	61.34%	46.87%	47.66%	59.29%	100.00%	100.00%
HasSister	74.17%	81.41%	70.61%	81.18%	46.89%	54.96%	46.77%	40.49%	48.17%	24.19%	100.00%	100.00%
IsGrandparent	57.00%	35.42%	45.83%	13.98%	32.37%	9.54%	8.60%	2.28%	5.22%	2.78%	100.00%	100.00%
IsUncle	26.67%	17.08%	49.05%	10.15%	2.86%	2.15%	2.00%	0.00%	0.00%	3.15%	85.00%	29.32%
IsMGUncle	10.00%	10.00%	48.33%	16.67%	0.00%	3.33%	0.00%	0.00%	0.00%	0.00%	55.24%	11.33%
4-Connectivity	62.43%	11.77%	73.28%	12.08%	44.25%	6.29%	28.40%	9.39%	61.48%	26.80%	80.82%	56.58%
6-Connectivity	63.01%	1.43%	73.98%	37.85%	25.67%	6.08%	37.86%	8.06%	61.03%	23.61%	83.26%	62.95%
1-Outdegree	100.00%	100.00%	57.77%	64.47%	3.33%	4.00%	24.99%	8.89%	47.06%	13.17%	100.00%	100.00%
2-Outdegree	88.57%	86.59%	86.67%	75.74%	6.67%	0.00%	31.00%	0.00%	34.94%	15.34%	100.00%	70.73%

IsUncle(x, y) case study 1

Question: For a family tree containing 10 family members, which is depicted with 4 kinds of relations: father, mother, son, daughter. The relations are: ...; P1 is P0's mother; **P8 is P3's mother; P3 is P8's daughter**. For a new relation uncle, from the above-given facts, who is whose uncle? Please answer by listing all the pairs that satisfied. You must answer in this format: Who is Whose uncle.

GPT-4 Answer: To determine the uncle relationships, we must find males who have siblings with children. An uncle is a brother to one of the parents. First, we analyze the relations to establish sibling and parent-child relationships:

...
 - **P8 is the father of P3 and also the mother of P3, which raise a contradiction (P8 holds two different and incompatible genders according to the relations).**
 - P5 is the mother of P4, P7, and P8; P5 doesn't have any siblings mentioned.
 ...
 Hence, the uncle relationship is:
 P0 is P4's uncle
 P0 is P7's uncle

Figure 4: An input-conflicting reasoning case of GPT-4.

reasoning to decision making. We strictly follow all the model setups in the original DLM paper for the evaluation. For all the tasks, the DLM model is trained on 10 training samples (the same set of samples used for LLMs' in-context few-shot prompting) for at most 500 epochs, and we early stop the training if the training loss is lower than 10^{-8} following the original DLM paper [34]. We use the F1-score as the evaluation metric for all experiments since the number of positive and negative samples is imbalanced.

5 EVALUATION RESULTS

In this section, we study the following research questions:

- **RQ1:** How good is LLMs' relational reasoning ability with standard natural language prompting?

IsUncle(x, y) case study 2

Question: For a family tree containing 20 family members, which is depicted with 4 kinds of relations: father, mother, son, daughter. The relations are: P13 is P0's father; P8 is P1's father; P13 is P2's father; P8 is P3's father; ... For a new relation uncle, from the above-given facts, who is whose uncle? Please answer by listing all the pairs that satisfied. You must answer in this format: Who is Whose uncle"

GPT-4 Turbo Answer: To determine the uncle-nephew/niece relationships, we must first find siblings within the family tree and then identify their nephews and nieces. Here is the analysis:

...
 Using this information, we can list the uncles and nieces/nephews:
 - **P0 is the uncle of P5 and P16 since he is a brother to their parents (sharing P13 and P19 as common parents).**
 - **P5 is the uncle of P0 and P2 since he is a brother to their parents (sharing P13 and P19 as common parents).**
 Now, let's list the uncle-nephew/niece relationships:
 ...
 - P4 is P1's uncle
 - P4 is P3's uncle

Figure 5: A common sense-conflicting reasoning case of the GPT-4 Turbo model.

- **RQ2:** How good is LLMs' relational reasoning ability with truth value prompting?
- **RQ3:** Is the state-of-the-art *prompting* technique effective in boosting LLMs' relational reasoning ability?

5.1 RQ1: How good is LLMs' relational reasoning ability with standard natural language prompting?

Table 1 shows the in-context few-shot prompting results of the LLMs using standard natural language prompting and the few-shot training results of the DLM. Overall, we have the following findings.

Among all the evaluated LLMs models, GPT-4 Turbo and GPT-4 perform the best in terms of both test performance and OOD

generalization. Their superiority can be attributed to the fact that the two models are much larger in size and have access to a much larger amount of pretraining data than other LLMs baselines. Furthermore, we observe that although the LLMs can achieve great performance/generalization on relatively easier tasks, their results compromise significantly on tasks that require more complex task-solving logic (*i.e.* tasks that contain more predicates in the ground-truth program). For example, while the GPT-4 model achieves 100% F1-score in terms of performance on the HasFather task, it only achieves 49.05% on the IsUncle. To better understand why the LLMs perform poorly on harder tasks, we conduct two case studies of the GPT-4 and GPT-4 Turbo models on the IsUncle test sample, shown in Figure 4 and Figure 5. For Figure 4, though it is stated in the input prompt that “P8 is P3’s mother; P3 is P8’s daughter” (bold blue text), GPT-4 generates an input-conflicting reasoning process stating that “P8 is the father of P3 and also the mother of P3” (bold red text). For Figure 5, as shown in GPT-4 Turbo’s answer, the model generates a common sense-conflicting reasoning process that claims that P0 and P5 are each other’s uncle (bold red text). These hallucination phenomena would therefore undermine the model’s relational reasoning ability, leading to erroneous conclusions.

Finally, though the DLM model is much smaller in size compared to the LLMs (DLM has only 60k parameters, while the smallest evaluated LLM has 7 billion parameters (llama-7B)), it manages to achieve the best results with only few-shot training. Note that the hardware requirement required for training the DLM model is much lower than even loading the LLMs (*e.g.* it only requires 429 MiB of GPU memory for training DLM while 36,957 MiB is required for loading llama-7B). The reason is that the logical reasoning knowledge injected in the DLM’s model design allows it to conduct robust program induction which makes it suitable for the relational reasoning tasks.

Answer to RQ1: GPT-4 and GPT-4 Turbo models present the strongest relational reasoning ability among all the evaluated LLMs. However, the LLMs’ performance/generalization drops significantly on tasks that require complex task-solving logic. Besides, the relational reasoning ability of the LLMs’ in-context few-shot prompting is generally worse than the few-shot trained DLM model which is much less hardware-demanding.

5.2 RQ2: How good is LLMs’ relational reasoning ability with truth value prompting?

We further assess the relational reasoning ability of the LLMs when using the same data modality as the DLM model (*i.e.* representing input/output predicate relations with truth value matrices). Table 2 shows the in-context few-shot learning results of the LLMs using truth value prompting. The result is marked as N/A if the model cannot generate the complete output array. We omit the results of the GPT-4, llama-7b, and llama-13b models as all their results are N/A. The reason for the complete N/A results for these three models is because they are relatively small in terms of context window whereas the GPT-4 Turbo and GPT-3.5 Turbo are capable

Table 2: Performance and generalization results of DLM and LLMs using truth table prompting.

Task	GPT-4-Turbo		GPT-3.5 Turbo		DLM	
	n=10	n=20	n=10	n=20	n=10	n=20
HasFather	79.17%	76.99%	47.57%	N/A	100.00%	100.00%
HasSister	47.06%	50.49%	38.41%	N/A	100.00%	100.00%
IsGrandparent	9.52%	N/A	6.62%	N/A	100.00%	100.00%
IsUncle	6.67%	N/A	0.00%	N/A	85.00%	29.32%
IsMGUncle	0.00%	N/A	0.00%	N/A	55.24%	11.33%
4-Connectivity	60.76%	N/A	53.35%	N/A	80.82%	56.58%
6-Connectivity	65.77%	N/A	51.67%	N/A	83.26%	62.95%
1-Outdegree	98.90%	99.19%	50.05%	N/A	100.00%	100.00%
2-Outdegree	100.00%	95.31%	42.65%	N/A	100.00%	70.73%

of processing long context window and generating long output sequence³.

Interestingly, we observe that the results of GPT-4 Turbo and GPT-3.5 Turbo on the family tree reasoning tasks decrease as compared to the corresponding standard natural language prompting results. This is because the natural language pretraining data regarding family relations is common and sufficient which allows the LLMs to perform well with natural language prompting on the family tree reasoning tasks. While the truth value pretraining data is much fewer comparatively which makes it challenging for the LLMs to reason well with truth value prompting. Whereas interestingly, the IID performance is consistent or even improves on the general graph reasoning tasks as compared to the standard natural language prompting results, the GPT-4 Turbo even manages to achieve the best performance and generalization on the 2-outdegree task. We believe that LLMs have the ability to perform graph reasoning with truth value matrices, which makes them promising for truth value-related tasks (*e.g.* logic synthesis [15]). We leave the further investigation to future work.

Answer to RQ2: LLMs with small context window are incapable of conducting relational reasoning with truth value prompting, while the LLMs with large context window present consistent or even improved IID performance on the general graph reasoning tasks compared to the standard natural language prompting results.

5.3 RQ3: Is the state-of-the-art *prompting* technique effective in boosting LLMs’ relational reasoning ability?

We further evaluate whether the recently proposed *chain-of-thought* (CoT) [31] prompting can improve the LLMs’ relational reasoning ability. Concretely, each exemplar in few-shot prompting is augmented with the direct cause-effect logic that represents the deduction process of the answer. Figure 6 shows an example CoT prompting of a HasSister sample, the bold text represents the example

³Note that for all the evaluated tasks, the ground-truth output truth value matrices required to generate never exceed the maximum output token length of any LLMs.

Table 3: Performance and generalization results of DLM and LLMs using chain-of-thought prompting.

Task	GPT-4-Turbo		GPT-4		GPT-3.5 Turbo		llama-7b		llama-13b		DLM	
	n=10	n=20	n=10	n=20	n=10	n=20	n=10	n=20	n=10	n=20	n=10	n=20
HasFather	100.00%	99.57%	100.00%	100.00%	76.32%	85.04%	59.49%	63.77%	86.98%	64.30%	100.00%	100.00%
HasSister	86.39%	76.26%	79.88%	73.60%	16.39%	20.26%	57.94%	47.98%	52.91%	53.33%	100.00%	100.00%
IsGrandparent	63.57%	38.38%	70.32%	39.02%	11.96%	0.82%	2.99%	1.03%	7.08%	2.75%	100.00%	100.00%
IsUncle	44.44%	23.84%	96.67%	16.35%	23.33%	1.55%	0.00%	0.00%	0.00%	1.43%	85.00%	29.32%
IsMGUncle	10.00%	10.00%	40.00%	4.68%	0.00%	0.00%	2.00%	0.00%	0.00%	0.00%	55.24%	11.33%
4-Connectivity	26.00%	3.44%	79.78%	19.33%	1.18%	28.41%	40.82%	18.64%	39.33%	27.13%	80.82%	56.58%
6-Connectivity	54.82%	3.50%	86.31%	11.46%	2.61%	24.27%	50.88%	18.98%	33.55%	23.02%	83.26%	62.95%
1-Outdegree	86.39%	83.18%	38.19%	38.45%	5.00%	2.86%	24.81%	20.51%	14.00%	6.19%	100.00%	100.00%
2-Outdegree	100.00%	89.24%	76.67%	19.16%	5.00%	0.00%	28.71%	0.00%	20.67%	4.72%	100.00%	70.73%

$\text{HasSister}(x)_{nl}(\text{CoT})$ prompt
<p>You are an agent who determines the relations in a family. I will first show you some examples, and then I will give you a new question for you to answer. The examples are shown as follows. Now this is the question for you to answer.</p> <p>Question: For a family tree containing 6 family members, which is depicted with 4 kinds of relations: father, mother, son, daughter. The relations are: P3 is P0's father; P3 is P2's father; P5 is P0's mother; P5 is P2's mother; P0 is P3's daughter; P2 is P3's daughter; P0 is P5's daughter; P2 is P5's daughter. For a new relation sister, from the above-given facts, who has relation sister? You must give the reasoning process and you must give the final answer in this format: Who has sister.</p> <p>Answer: For P0: Since P5 is P0's mother and P2 is P5's daughter, P0 has 1 sister. Therefore, we can derive that P0 has sister. For P2: Since P5 is P2's mother and P0 is P5's daughter, P2 has 1 sister. Therefore, we can derive that P2 has sister.</p> <p>Now this is the question for you to answer.</p> <p>Question: For a family tree containing 6 family members, which is depicted with 4 kinds of relations: father, mother, son, daughter. The relations are: P0 is P1's father; P5 is P1's mother; P1 is P0's daughter; P1 is P5's daughter. For a new relation sister, from the above-given facts, who has relation sister? You must give the reasoning process and you must give the final answer in this format: Who has sister.</p>

Figure 6: Chain-of-thought prompting of a HasSister sample.

deduction process. The evaluation results of LLMs using natural language CoT prompting are shown in Table 3. Though it is reported in previous literature that CoT can generally improve the LLMs' performance on benchmarks such as arithmetic, commonsense, *etc.* We observe that for the evaluated relational reasoning benchmarks, CoT cannot consistently improve the performance/generalization of the LLMs. For some tasks, the results are even significantly decreased, *e.g.* the performance and generalization of GPT-4 Turbo on the 1-outdegree task are decreased by 13.61%, 16.82% respectively.

Answer to RQ3: The CoT prompting technique is not generally effective for the relational reasoning benchmarks. The

CoT prompting can even significantly decrease the LLMs' performance/generalization for some tasks.

6 RELATED WORKS

6.1 LLMs for logical reasoning

With the rapid advancement of LLMs, there has been a recent surge in research that leverages LLMs to solve logical reasoning tasks, such as common-sense reasoning [24, 26], arithmetic reasoning [5, 23], symbolic reasoning [25], *etc.* It is observed in previous works that the standard prompting (only questions are given) is not sufficient as the model performs poorly [31]. Aiming to improve the LLMs' performance, Brown *et al.* [3] propose prompting the model with a few input-output exemplars demonstrating the task, which is proven to be effective for a range of simple question-answering tasks. Wei *et al.* [31] propose providing the model with the concrete chain-of-thought prompting (*i.e.* I/O examples with the corresponding explicit reasoning steps) to the LLMs. It is demonstrated that the chain-of-thought prompting technique improves performance on arithmetic reasoning, commonsense reasoning, and symbolic reasoning tasks.

6.2 Neural Program Induction & Synthesis

Leveraging neural networks for program induction and program synthesis has become increasingly popular. Recent work has demonstrated the effectiveness of neural network-based program induction and synthesis methods in relational reasoning tasks with better noise tolerance and less manual effort [8, 10, 12]. Dong *et al.* [9] propose Neural Logic Machines (NLMs) which approximate logic predicates and logic rules with neural modules and achieves state-of-the-art performance on the program induction tasks. To improve the interpretability of the NLM model, Zimmer *et al.* [34] propose using soft logic operators as the computation unit which allows the model to generate interpretable logic programs. Cao *et al.* [4] propose a sketch-based hierarchical program synthesis framework for complex sequential decision-making problems based on ∂LLP , which achieves high performance and generalizability.

7 CONCLUSION

In this work, we propose a general and extensible relational reasoning ability assessment pipeline for the LLMs and the neural program induction models. Based on the pipeline, we conduct the first comprehensive evaluations of the LLMs' relational reasoning ability. Our results show that when using standard natural language prompting, the relational reasoning ability of LLMs' in-context few-shot prompting is generally far from satisfaction compared with the program induction model which is much smaller in size. While the LLMs with large context window present consistent or even improved IID performance on the general graph reasoning tasks. And we further show that the current state-of-the-art prompting technique is not generally effective for improving LLMs' relational reasoning ability.

8 ACKNOWLEDGEMENT

This research is supported by the National Research Foundation, Singapore, and the Cyber Security Agency under its National Cybersecurity R&D Programme (NCRP25-P04-TAICeN) and NRF Investigatorship NRF-NRFI06-2020-0001, and the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contributions from the industry partner(s). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore and Cyber Security Agency of Singapore.

REFERENCES

- [1] Matej Balog, Alexander L Gaunt, Marc Brockschmidt, Sebastian Nowozin, and Daniel Tarlow. 2016. DeepCoder: Learning to Write Programs. In *International Conference on Learning Representations*.
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [4] Yushi Cao, Zhiming Li, Tianpei Yang, Hao Zhang, Yan Zheng, Yi Li, Jianye Hao, and Yang Liu. 2022. GALOIS: boosting deep reinforcement learning via generalizable logic synthesis. *Advances in Neural Information Processing Systems* 35 (2022), 19930–19943.
- [5] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168* (2021).
- [6] Eveline A Crone, Carter Wendelken, Linda Van Leijenhorst, Ryan D Honovich, Kalina Christoff, and Silvia A Bunge. 2009. Neurocognitive development of relational reasoning. *Developmental science* 12, 1 (2009), 55–66.
- [7] Gelei Deng, Yi Liu, Victor Mayoral-Vilches, Peng Liu, Yuekang Li, Yuan Xu, Tianwei Zhang, Yang Liu, Martin Pinzger, and Stefan Rass. 2023. Pentestgpt: An llm-empowered automatic penetration testing tool. *arXiv preprint arXiv:2308.06782* (2023).
- [8] Jacob Devlin, Jonathan Uesato, Surya Bhupatiraju, Rishabh Singh, Abdel-rahman Mohamed, and Pushmeet Kohli. 2017. Robustfill: Neural program learning under noisy i/o. In *International conference on machine learning*. PMLR, 990–998.
- [9] Honghua Dong, Jiayuan Mao, Tian Lin, Chong Wang, Lihong Li, and Denny Zhou. [n. d.]. Neural Logic Machines. In *International Conference on Learning Representations*.
- [10] Richard Evans and Edward Grefenstette. 2018. Learning explanatory rules from noisy data. *Journal of Artificial Intelligence Research* 61 (2018), 1–64.
- [11] Lise Getoor, Nir Friedman, Daphne Koller, and Avi Pfeffer. 2001. Learning probabilistic relational models. *Relational data mining* (2001), 307–335.
- [12] Claire Glanois, Zhaohui Jiang, Xuening Feng, Paul Weng, Matthieu Zimmer, Dong Li, Wulong Liu, and Jianye Hao. 2022. Neuro-symbolic hierarchical rule induction. In *International Conference on Machine Learning*. PMLR, 7583–7615.
- [13] Alex Graves, Greg Wayne, Malcolm Reynolds, Tim Harley, Ivo Danihelka, Agnieszka Grabska-Barwińska, Sergio Gómez Colmenarejo, Edward Grefenstette, Tiago Ramalho, John Agapiou, et al. 2016. Hybrid computing using a neural network with dynamic external memory. *Nature* 538, 7626 (2016), 471–476.
- [14] Graeme S Halford, William H Wilson, and Steven Phillips. 1998. Processing capacity defined by relational complexity: Implications for comparative, developmental, and cognitive psychology. *Behavioral and brain sciences* 21, 6 (1998), 803–831.
- [15] Soha Hassoun and Tsutomu Sasao. 2001. *Logic synthesis and verification*. Vol. 654. Springer Science & Business Media.
- [16] Kai Huang, Xiangxin Meng, Jian Zhang, Yang Liu, Wenjie Wang, Shuhao Li, and Yuqing Zhang. 2023. An Empirical Study on Fine-Tuning Large Language Models of Code for Automated Program Repair. In *2023 38th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE Computer Society, 1162–1174.
- [17] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems* 35 (2022), 22199–22213.
- [18] Daphne Koller, Nir Friedman, Sašo Džeroski, Charles Sutton, Andrew McCallum, Avi Pfeffer, Pieter Abbeel, Ming-Fai Wong, Chris Meek, Jennifer Neville, et al. 2007. *Introduction to statistical relational learning*. MIT press.
- [19] Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691* (2021).
- [20] Belinda Z Li, Maxwell Nye, and Jacob Andreas. 2021. Implicit representations of meaning in neural language models. *arXiv preprint arXiv:2106.00737* (2021).
- [21] Robin Manhaeve, Sebastijan Dumančić, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. 2021. Neural probabilistic logic programming in DeepProbLog. *Artificial Intelligence* 298 (2021), 103504.
- [22] Stephen Muggleton. 1991. Inductive logic programming. *New generation computing* 8 (1991), 295–318.
- [23] Arkil Patel, Satwik Bhattamishra, and Navin Goyal. 2021. Are NLP models really able to solve simple math word problems? *arXiv preprint arXiv:2103.07191* (2021).
- [24] Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. Winogrande: An adversarial winograd schema challenge at scale. *Commun. ACM* 64, 9 (2021), 99–106.
- [25] Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al. 2022. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615* (2022).
- [26] Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2018. Commonsenseqa: A question answering challenge targeting commonsense knowledge. *arXiv preprint arXiv:1811.00937* (2018).
- [27] Aalok Thakkar, Aaditya Naik, Nathaniel Sands, Rajeev Alur, Mayur Naik, and Mukund Raghothaman. 2021. Example-guided synthesis of relational queries. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*. 1110–1125.
- [28] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shrutu Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023).
- [29] Karthik Valmeekam, Matthew Marquez, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. 2023. PlanBench: An Extensible Benchmark for Evaluating Large Language Models on Planning and Reasoning about Change. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- [30] Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned Language Models are Zero-Shot Learners. In *International Conference on Learning Representations*.
- [31] 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. *Advances in Neural Information Processing Systems* 35 (2022), 24824–24837.
- [32] Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhjanjan Kambadur, David Rosenberg, and Gideon Mann. 2023. Bloombergpt: A large language model for finance. *arXiv preprint arXiv:2303.17564* (2023).
- [33] Hongyang Yang, Xiao-Yang Liu, and Christina Dan Wang. 2023. FinGPT: Open-Source Financial Large Language Models. *arXiv preprint arXiv:2306.06031* (2023).
- [34] Matthieu Zimmer, Xuening Feng, Claire Glanois, Zhaohui JIANG, Jianyi Zhang, Paul Weng, Dong Li, Jianye HAO, and Wulong Liu. 2023. Differentiable Logic Machines. *Transactions on Machine Learning Research* (2023). <https://openreview.net/forum?id=mXfkKtu5JA>