

GraphReason: Enhancing Reasoning Capabilities of Large Language Models through A Graph-Based Verification Approach

Lang Cao

University of Illinois Urbana-Champaign
Department of Computer Science
langcao2@illinois.edu

Abstract

Large Language Models (LLMs) have showcased impressive reasoning capabilities, particularly when guided by specifically designed prompts in complex reasoning tasks such as math word problems. These models typically solve tasks using a chain-of-thought approach, which not only bolsters their reasoning abilities but also provides valuable insights into their problem-solving process. However, there is still significant room for enhancing the reasoning abilities of LLMs. Some studies suggest that the integration of an LLM output verifier can boost reasoning accuracy without necessitating additional model training. In this paper, we follow these studies and introduce a novel graph-based method to further augment the reasoning capabilities of LLMs. We posit that multiple solutions to a reasoning task, generated by an LLM, can be represented as a reasoning graph due to the logical connections between intermediate steps from different reasoning paths. Therefore, we propose the Reasoning Graph Verifier (GraphReason) to analyze and verify the solutions generated by LLMs. By evaluating these graphs, models can yield more accurate and reliable results. Our experimental results show that our graph-based verification method not only significantly enhances the reasoning abilities of LLMs but also outperforms existing verifier methods in terms of improving these models' reasoning performance.

1 Introduction

Large Language Models (LLMs) have demonstrated exceptional capabilities in a variety of human tasks (Zhao et al., 2023). Among the many abilities LLMs possess, their reasoning capacity is of paramount importance (Kojima et al., 2023; Huang and Chang, 2023). This has been substantiated by recent progresses (Wei et al., 2022; Zhou et al., 2023; Lampinen et al., 2022a). Equipped with the ability to reason, especially in a multi-step manner, LLMs can decompose complex prob-

Chain-of-thought Reasoning in Math Word Problem	
Exemplars with Step-by-step Solution	<p>Q: Janet's ducks lay 16 eggs per day. She eats three for breakfast every morning and bakes muffins for her friends every day with four. She sells the remainder at the farmers' market daily for \$2 per fresh duck egg. How much in dollars does she make every day at the farmers' market?</p> <p>A: Step 1: Janet sells $16 - 3 - 4 = 9$ duck eggs a day. Step 2: She makes $9 * 2 = 18$ every day at the farmer's market. Step 3: ##### 18 (More Exemplars)</p>
Current Question	<p>Q: A robe takes 2 bolts of blue fiber and half that much white fiber. How many bolts in total does it take?</p>
LLM's Generated Solution	<p>A: Step 1: It takes $2/2 = 1$ bolt of white fiber. Step 2: So the total amount of fabric is $2+1=3$ bolts of fabric. Step 3: ##### 3</p>

Figure 1: An example of chain-of-thought reasoning in a math word problem, using data from the GSM8K dataset. Large language models learn from exemplars that provide step-by-step solutions, subsequently generating their reasoning path for the current question.

lems into simpler tasks, thereby facilitating their resolution. In everyday life, many complex tasks typically require multi-step solutions. A prime example of a reasoning task is arithmetic reasoning, also known as solving math word problems (Zhang et al., 2019). These math word problems represent simplified versions of complex real-life situations.

The reasoning ability is inherent in Large Language Models (LLMs), but it necessitates specific methods for manifestation. To activate the robust reasoning capability of LLMs, the use of specially designed prompts should be considered. Numerous methods have been proposed to tap into this potential, among which chain-of-thought reasoning (Wei et al., 2022) and in-context learning (Lampinen et al., 2022b) are two notable approaches. Chain-of-thought reasoning can elucidate the reasoning paths during the process. In-context learning fur-

nishes LLMs with exemplary cases, thereby enabling them to learn from and simulate these examples for improved results. In the arithmetic reasoning scenario, GPT-4 can achieve an accuracy of 92% on the GSM8K dataset using 5-shot chain-of-thought prompts (Cobbe et al., 2021a). This represents a level of difficulty that a bright middle school student should be capable of handling. As depicted in Figure 1, this illustrates a multi-step arithmetic reasoning process in LLMs.

In addition to further training of LLMs and prompt design, some methods have been proposed to enhance the reasoning capabilities of LLMs from the perspective of output verification. The primary idea is to have LLMs generate reasoning paths multiple times, and then design a verifier to evaluate these paths and deliver the final results. (Wang et al., 2023) introduces the concept of *self-consistency*, based on the intuition that a complex reasoning problem usually allows for multiple thought processes, all leading to a unique correct answer. (Li et al., 2023) also proposes *All Roads Lead to Rome*, which introduces a step-aware verifier to analyze reasoning paths not just through the entire path, but at every step. However, both methods treat each reasoning path as an independent entity and do not consider the potential interrelation and interaction between different reasoning paths. Once reasoning paths are disassembled into steps, intermediate steps from one path may bear reasoning relations to other reasoning paths. These methods do not perceive all LLM outputs for a given input as a collective entity, thereby failing to analyze the internal relations of all candidate paths in depth.

Inspired by these observations, we propose **Reasoning Graph Verifier** (GraphReason) in this paper. We posit that reasoning paths of one question can form reasoning graphs, where similar intermediate reasoning steps can be merged into the same node. With a graph structure, we can more effectively model and capture the reasoning logic between intermediate steps from different reasoning paths. Specifically, we first construct a reasoning graph based on all outputs from LLMs, and then train a verifier to learn the relationship between the graph structure and the final answer. During the prediction stage, we process the data in the same way as in the training stage, and use the verifier to evaluate each reasoning graph. We then select the reasoning graph with the highest score, using its answer as the final answer. To the best of our

knowledge, we are the first to approach reasoning logic of LLMs from a graph perspective. We conduct extensive experiments to demonstrate the improvements over the original LLMs, and show that our method outperforms other verifiers.

In summary, our contributions are as follows:

- We propose a graph-based verification method, GraphReason, aimed at significantly enhancing the reasoning capabilities of large language models without the need for additional training of LLMs.
- We establish an arithmetic reasoning benchmark using three Math Word Problem datasets to illustrate the fundamental reasoning performance of large language models, and to provide a fair comparison of the performance of various existing verifiers.
- Our experimental results indicate that the method proposed in this paper outperforms other enhancement methods. We also provide an extensive analysis of the limitations and future potential of GraphReason.

2 Related Works

Reasoning of Fine-tuning Models has been extensively studied. It focuses on addressing reasoning tasks using a general sequence-to-sequence approach, enhanced by reasoning-aware pre-training or fine-tuning of language models. (Cobbe et al., 2021a) proposed training a verifier to rank solutions sampled from fine-tuned language models. (Yoran et al., 2022; Wang et al., 2022) suggested equipping language models with reasoning abilities by generating training examples with human-designed templates. (Pi et al., 2022) proposed injecting reasoning capabilities into language models by continually pre-training on program execution data.

Several studies have focused on imbuing PLM with reasoning ability for specific tasks, such as arithmetic reasoning (Cobbe et al., 2021a; Miao et al., 2020; Patel et al., 2021), commonsense reasoning (Talmor et al., 2019), and inductive reasoning (Sinha et al., 2019). For instance, various strategies have been proposed to improve language models’ performance on arithmetic reasoning tasks, often referred to as math word problems. (Xie and Sun, 2019) proposed a tree-structured decoder to generate an equation tree, while (Zhang et al., 2020) applied graph convolutional networks to extract relationships

of quantities in math problems. (Li et al., 2022) used contrastive learning to better learn patterns in math word problems. However, (Valmeekam et al., 2023; Rae et al., 2022) suggested that reasoning, particularly multi-step reasoning, is often a weakness in language models and other NLP models.

Reasoning of Large Language Models has garnered significant attention and demonstrated immense potential. Recent advancements in LLMs suggest that the ability for multi-step reasoning is already embedded within these large-scale models (Kojima et al., 2023; Huang and Chang, 2023), such as PaLM (Chowdhery et al., 2022), GPT-4 (OpenAI, 2023). Therefore, providing an adequate prompt is sufficient to utilize this reasoning ability. For example, the prompting method proposed by (Kojima et al., 2023; Wei et al., 2022), which is based on a chain-of-thought, could aid LLMs in generating text with arithmetic reasoning and common factual knowledge. Following (Wei et al., 2022), experiments on current language models demonstrated that chain-of-thought prompting could enhance the accuracy of solving math problems from 18% to 57%. (Lampinen et al., 2022b) included explanations in the in-context examples and tested the influence of explanations by evaluating the score between *explain-then-predict* and *predict-then-explain*. Moreover, (Zhou et al., 2023) suggested a two-stage prompting strategy, *least-to-most* prompting, which breaks down a complex problem into a series of subproblems and solves them step-by-step. (Li et al., 2023) proposed sampling multiple times from diverse prompts to enhance the variety of responses.

In addition to designing prompts, adopting additional strategies like verifier has contributed to enhancing the performance of reasoning abilities of large language models. For instance, (Wang et al., 2023) proposes *self-consistency*, which involves sampling different reasoning paths from the language model, and then returning the most consistent final answer via majority voting. (Li et al., 2023) used a step-aware voting verifier to enhance the reasoning ability of LLMs from two perspectives. These methods strive to augment the reasoning abilities or yield superior reasoning results without necessitating additional training of LLMs. Our work continues this research direction, with a specific focus on developing a novel graph-based verifier to boost the reasoning capabilities of LLMs.

3 Methodology

3.1 GraphReason Framework

Problem 1 (Reasoning to Solve Problems)

Given a set of n math word problems $\mathbf{Q} = \{Q_1, Q_2, \dots, Q_n\}$, where each Q_i is represented by the text description of a single math word problem, the goal of reasoning to solve math word problems is to generate the answers $\mathbf{A} = \{A_1, A_2, \dots, A_n\}$ for these problems. Here, each A_i represents the generated text of the corresponding answer. During the process of large language models generating answers, a set of n reasoning paths for solutions $\mathbf{S} = \{S_1, S_2, \dots, S_n\}$ is also produced. Each solution S_i is represented as $S_i = \{Q, \text{Step}_1, \text{Step}_2, \dots, \text{Step}_l, A\}$, where each Step_i denotes the intermediate steps in the step-by-step solutions.

We propose GraphReason to verify the solutions generated by LLMs in order to improve the final answer accuracy. This method is a graph-based verification technique that analyzes reasoning paths from generated solutions from a graph perspective. The final answer is obtained without modifying the original LLMs, functioning much like a plugin. As illustrated in Figure 2, there are two steps in the training stage: *Graph Construction* and *Graph Classification*. In the *Graph Construction* step, we obtain the generated solution from LLMs with the specific designed prompt and group them according to their final answers. We split reasoning paths by steps and then merge intermediate steps with identical expression to form reasoning graphs. In the *Graph Classification* step, we classify these reasoning graphs with the additional feature of the sum of scores from the *base verifier* to train the integrated verifier model. In the prediction stage, the candidate solutions are first generated by LLMs. We process them in the same manner as in the training stage, then we use trained verifier to evaluate the scores of each candidate solution. The best solution, denoted by the highest score, is selected as the final predicted answer. We will now provide a detailed introduction to the entire process.

3.2 Prompt Design

To improve the output of Language Models (LLMs) in providing solutions, it is essential to design effective prompts. We incorporate chain-of-thought and in-context learning to enable LLMs to generate step-by-step answers for math word problems. The language models generate output y based on the

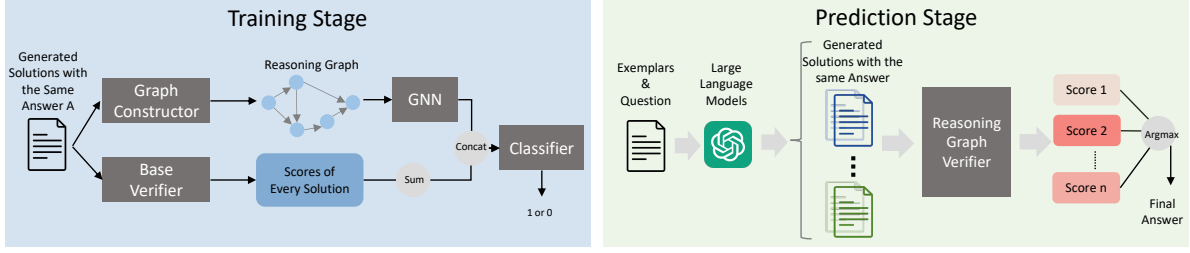


Figure 2: The framework of GraphReason. In the training stage, GraphReason processes generated solutions from LLMs to construct reasoning graphs, and then trains a verifier to judge them according to graph classification. In the prediction stage, GraphReason evaluates candidate solutions to assign a score, and selects the solution with the highest score as the final answer.

input \mathbf{x} using the following equation:

$$p(\mathbf{y}|C, \mathbf{x}) = \prod_{t=1}^{|\mathbf{y}|} p_{LM}(\mathbf{y}_t|C, \mathbf{x}, \mathbf{y} < t), \quad (1)$$

where, C represents the input provided to the LLMs prior to the current math word problem’s question. C is a concatenation of k exemplars, denoted as:

$$C = [(Q_1, S_1, A_1); (Q_2, S_2, A_2), \dots; (Q_k, S_k, A_k)], \quad (2)$$

where, Q_i represents the question, S_i represents the intermediate steps of the solution, and A_i represents the answer. We set k to five in this study, resulting in a prompt that consists of five question-answer pairs sampled from the training split of a math word problem dataset. Therefore, the prompt can be denoted as:

$$\text{Prompt} = [C; Q], \quad (3)$$

where Q represents the question of the current math word problem.

Using a greedy decoding approach to sample one output from LLMs may not be robust. It can lead to instability and occasional errors. To address this, (Wang et al., 2023) propose the concept of *self-consistency*. This approach involves sampling different reasoning paths from the language model and then selecting the most consistent final answer through majority voting. Instead of using greedy decoding to sample only once and verify, they utilize sampling decoding to sample N_1 times. We also follow the idea presented by (Li et al., 2023) in their work named *All Roads Lead to Rome*. This approach involves generating N_2 diverse prompts for LLMs to produce multiple outputs. By employing multiple sampling decodes on diverse prompts, we can obtain generated solutions from different sources. Specifically, we obtain $N = N_1 \times N_2$

diverse reasoning paths for each question. In our main experiments, we set $N_1 = 10$ and $N_2 = 3$. These solutions will be further processed and verified using our designed verifier.

3.3 Reasoning Graph Construction

After generating multiple solutions for a question, it becomes necessary to construct reasoning graphs based on the reasoning paths taken by these solutions.

As shown in Figure 3, we begin by grouping all the generated solutions for a particular question according to their final answer. Since these solutions originate from the same question, their reasoning paths will share the same starting point. Similarly, solutions with the same final answer will have the same endpoint, as their reasoning paths converge. Therefore, a group of generated solutions with the same final answer can form a reasoning graph with a uniform start node (question node) and end node (answer node). We define this division process as follows:

$$\mathbf{S} = \{S_{A_1}, S_{A_2}, \dots, S_{A_n}\}, \quad (4)$$

where \mathbf{S} represents the set of generated solutions for a question, and $S_{A_i} = \{S_1, S_2, \dots, S_m\}$ is the subset of generated solutions that all have the same final answer A_i .

For each subset of generated solutions S_{A_i} , we construct a reasoning graph. This construction is motivated by the understanding that each step in the reasoning path of a generated solution does not exist in isolation from the other solutions. The steps from one solution’s reasoning path can impact the steps from another solution, enhancing the overall reasoning process. We utilize the graph structure to model and capture these relationships between steps from different solutions. As the different reasoning paths can benefit each other, we construct




Figure 3: The graph constructor in GraphReason. We detail the process of transforming ‘Generated Solutions with the Same Answer A1’ to ‘Graph 1’.

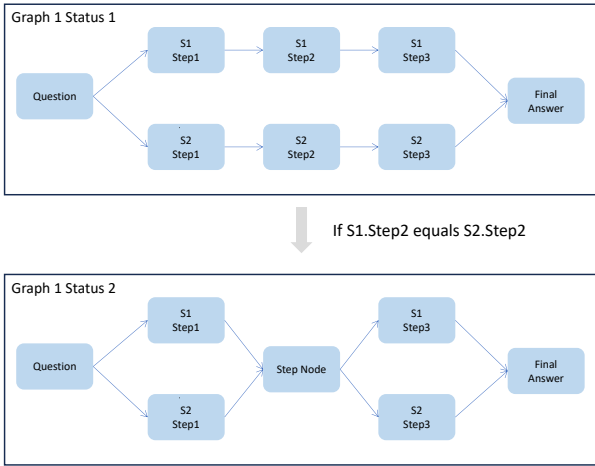


Figure 4: The process of reasoning graph construction. The primary operation here is the merging of identical intermediate steps in reasoning paths into a single graph node.

a reasoning graph to link these paths together. As shown in Figure 4, the primary operation here is the merging of identical intermediate nodes in reasoning paths into a single graph node. We first compare the reasoning steps from any two solution reasoning paths. If they have the same intermediate steps of arithmetic expression, we merge them into the same node, and if they differ, we do not. For reasoning math word problems here, we define reasoning steps as the current arithmetic expression without other language text in the current reasoning step for clarity. It can help us simplify construction of reasoning graphs in the reasoning task. The detailed algorithm for constructing a reasoning graph is shown in Algorithm 1.

The generated solutions, divided by their final

answers $\{S_{A_1}, S_{A_2}, \dots, S_{A_n}\}$, can be transformed into n reasoning graphs of generated solutions $\{G_{A_1}, G_{A_2}, \dots, G_{A_n}\}$.

Regarding the node features in the graph, we select the score from the *Base Verifier* and the node degree. We believe the score from the *Base Verifier* encapsulates the semantic information of solutions, and the node degree contains information about the graph structure. The *Base Verifier* is trained independently from the whole framework. It is designed to judge whether a single reasoning path of one solution is correct, which is a binary text classification task. After training, it can be used to verify any single solution and assign a $score \in (0, 1)$ to evaluate the likelihood of the solution being correct, where $score = 0.99$ suggests a 99% probability of the solution being correct. We use the score from the *Base Verifier* to better incorporate solution semantic information because, according to our experiments, it is challenging to model semantic information while modeling reasoning logic information. The score of a step is the same as its solution score. Therefore, for one step node V , it has many scores $\{score^a, score^b, \dots, score^e\}$ from different solutions. The feature of one node V_i in the graph is then concatenated by the selected feature, which can be represented as:

$$\mathbf{V} = [score_i^{mean}, score_i^{max}, score_i^{min}, score_i^{num}, in_degree_i], \quad (5)$$

where $\mathbf{V} \in \mathbb{R}^5$, $score_i^{mean}$ is the mean of all scores of one step V_i , $score_i^{max}$ is the maximum score, $score_i^{min}$ is the minimum score, $score_i^{num}$ is the number of scores, and in_degree_i is the in-degree of the step node V_i .

Algorithm 1 Reasoning graph construction algorithm

Input: generated solutions S_{A_i} which have the same final answers

Output: a reasoning graph G_{A_i}

```
1:  $node\_num \leftarrow 0$ 
2:  $node2id \leftarrow dict()$ 
3:  $edges \leftarrow list()$ 
4: for each  $reason\_path$  in  $S_{A_i}$  do
5:   for each  $step$  in  $reason\_path$  do
6:     if  $step$  not in  $node2id.keys()$  then
7:        $node2id[step] \leftarrow node\_num$ 
8:        $node\_num \leftarrow node\_num + 1$ 
9:     end if
10:  end for
11: end for
12: for each  $reason\_path$  in  $S_{A_i}$  do
13:   for each  $step$  in  $reason\_path$  do
14:      $start\_node \leftarrow node2id[last\_step]$ 
15:      $end\_node \leftarrow node2id[step]$ 
16:     if  $(start\_node, end\_node)$  not in  $edges$  then
17:        $edges.add((start\_node, end\_node))$ 
18:     end if
19:      $last\_step \leftarrow step$ 
20:   end for
21: end for
22:  $G_{A_i} \leftarrow graph(node2id, edges)$ 
```

In this way, we can obtain multiple reasoning graphs to represent all generated solutions from LLMs for a single math word problem question.

3.4 Verifier Design

Our designed verifier GraphReason, is used to evaluate the answer of a generated solutions group, which is also represented as a reasoning graph. This verifier has two inputs: the graph and the sum of solution scores. We employ the Graph Isomorphism Network (GIN) (Xu et al., 2019) to perform node feature propagation, thereby encoding the information from the reasoning graphs we obtained. The node feature is propagated and aggregated as follows:

$$h_v^{(k)} = MLP^{(k)}((1 + \varepsilon^{(k)}) \cdot h_v^{(k-1)} + \sum_{u \in N(v)} h_u^{(k-1)}), \quad (6)$$

where $h_v^{(k)}$ represents the state of node v after the k^{th} update. $MLP^{(k)}$ refers to a multi-layer perceptron in the k^{th} layer. $N(v)$ denotes all the neighbors of node v and ε is a learnable parameter. Then,

we perform a sum readout to obtain the representation of the reasoning graph:

$$h_G = \sum_{v \in G} h_v^{(k)}, \quad (7)$$

where $h_G \in \mathbb{R}^5$. We set k to 3, signifying the application of three layers of GIN. Concurrently, the sum of the scores of solutions with the same final answer, A_i , denoted as $score_{A_i}$, is represented as follows:

$$score_A = \sum_{i \in S_A} score_i. \quad (8)$$

Then a reasoning graph can then be represented as:

$$\mathbf{G} = [h_G, score_A], \quad (9)$$

where $\mathbf{G} \in \mathbb{R}^6$.

The target label of the graph $\mathbf{y} \in \{0, 1\}$ indicates whether the final answer matches the correct final answer. We compute the loss and train the verifier model by:

$$\mathcal{L} = \sum_{i=1}^n \mathcal{L}_{BCE}(label_i, f(\mathbf{G}_i)), \quad (10)$$

where i represents the number of solution subset among all n subsets after grouping solutions. The corresponding reasoning graph for this subset is denoted by \mathbf{G}_i , and $f()$ is a linear classifier.

3.5 Answer Verification

During the prediction stage, all generated solutions are processed in the same way as in the training stage. The trained verifier is then used to evaluate the scores of each reasoning graph, each of which represents a group of solutions that yield the same final answer. The final answer associated with the highest score is selected as our final predicted answer:

$$\hat{\mathbf{y}} = \mathbf{Answer}[\arg \max_i score_i], \quad (11)$$

where $score_i$ denotes the score of the reasoning graph \mathbf{G}_i , as determined by our verifier. **Answer** represents the list of all candidate final answers. By predicting the number of the optimal reasoning graph, we can determine the final predicted result of the current reasoning task.

	GSM8K	SVAMP	ASDiv-a	StrategyQA
Fine-tuning SOTA	57 ^a	57.4 ^b	75.3 ^c	73.9 ^d
9–12 year olds	60	-	-	-
gpt-3.5-turbo:				
Greedy Decode	72.7	78.7	93.0	65.0
Self-Consistency (Voting)	82.3	82.9	95.6	66.0
Verifier	66.9	73.1	92.8	69.3
Voting Verifier	85.4	84.8	96.9	70.7
DIVERSE (Step-aware Voting Verifier)	85.0	85.1	96.8	66.9
Reasoning Graph Verifier (Ours)	85.7	85.4	97.0	71.2

Table 1: The comparison experiment results of GraphReason, other verifiers, and other baselines. We primarily compare GraphReason with other verifiers which are all based on the same generated solutions from *gpt-3.5-turbo*

4 Experiments

In this section, we conducted extensive experiments to demonstrate the performance of GraphReason, along with a more in-depth analysis. Universally, we reproduced all types of verifiers to report their results based on the same generated solutions. Our experiments are conducted in two settings: Arithmetic Reasoning and Commonsense Reasoning. We ensured a fair comparison by setting the same random seed, using the same hardware environment, and applying similar hyperparameters. We used accuracy as the metric to evaluate the ability of solving math word problems, which determines whether the final answer is correct or not.

4.1 Training Details

For LLMs sampling, we use *gpt-3.5-turbo* as our base LLMs and set the temperature t to 1. All verifiers use the same LLMs’ output. Regarding verifier training, we fine-tune on *bert-base-uncased* (Devlin et al., 2019). We employ the AdamW optimizer (Loshchilov and Hutter, 2019) to optimize the model parameters during training. We apply differential learning rates, setting the learning rate of the final linear classifier to $4e-2$, while the other graph neural network layers are set to $4e-3$. The activation layer between them is ReLU (Agarap, 2019). The batch size in each training step is set to 2. The batch size is small because the verifier needs to verify multiple reasoning graphs for a single question.

To ensure a fair comparison between the Voting Verifier, Simple Verifier, and GraphReason, we use the same trained base verifier for all three approaches.

The details of the datasets and baselines are provided in Appendix A and Appendix B, respectively.

4.2 Main Results

We present the main results in Table 1. As can be seen from the table, GraphReason significantly enhances the original *gpt-3.5-turbo*’s reasoning abilities across all three datasets, for instance, improving accuracy by 13.0% (72.7% \rightarrow 85.7%) on GSM8K. It is also evident that our method surpasses other verifier methods with the same output from LLMs and achieves the state-of-the-art on all three datasets.

Additionally, the Step-aware Voting Verifier improves upon the Voting Verifier by recognizing that not all steps in an incorrect reasoning path are equally erroneous, and some steps may still be useful for reasoning. We believe this hypothesis is overly simplistic and cannot describe complex logical relationships among steps. According to Table 1, it leads to some metric decline, and the same finding also observed in the original paper. Furthermore, it does not perform well in the StrategyQA task, because there are no gold reasoning paths for the training of this commonsense reasoning task. In this task, the reasoning paths are generated and pseudo, indicating a requirement for gold labels at each step of the reasoning process. However, our paper consistently improves upon the Voting Verifier by considering complex relationship between different reasoning paths through reasoning graphs. We enhance the previous method, which did not consider relations in steps between different solutions, by 0.3% (85.4% \rightarrow 85.7%), 0.3% (85.1% \rightarrow 85.4%), 0.1% (96.9% \rightarrow 97.0%), and 0.5% (70.7% \rightarrow 71.2%) across the four datasets.

Moreover, GraphReason yields only a slight improvement in performance on ASDiv-a, and the results are nearly identical. One reason for this is that the math word problems from ASDiv-a are sim-

	GSM8K	▽	SVAMP	▽	ASDiv-a	▽
Reasoning Graph Verifier (Ours)	85.7	-	85.4	-	97.0	-
w/o solution semantic from base verifier	81.2	-4.5	83.1	-2.3	94.3	-2.7
w/o solution scores sum	82.8	-2.9	83.2	-2.2	95.6	-1.4
w/o reasoning graphs	85.4	-0.3	84.8	-0.6	96.9	-0.1

Table 2: The ablation experiment results of GraphReason. Missing each component leads to a decline in the final result.

pler compared to those in the other two datasets, based on our observations. In most cases, these problems do not require complex reasoning from a graph perspective to generate a satisfactory answer. It demonstrates that our method is particularly well-suited for such situations. We believe that GraphReason can offer more substantial improvements in the more complex scenario.

4.3 Ablation Study

We conducted an ablation study to evaluate the impact of each component on the overall performance of our method. Table 2 presents the results of this study, highlighting how these modules contribute to the improvement of the base model in distinct ways. It can be observed that the omission of any component leads to a decline in the final result. The solution semantics from the base verifier appear to be most crucial to GraphReason. The current method still relies on semantic information, which is reasonable since reasoning steps from different solutions require semantic information for better reasoning. We also notice that reasoning graphs bring a slight improvement to the entire method, thereby proving effectiveness of graph structure. The improvement is not substantial because we do not model the graph structure and semantic information simultaneously, and create a training gap here. Another essential factor is the complexity of graph classification, compounded by the presence of noise and limitations in our training data.

4.4 GraphReason with Different LLMs

To evaluate the compatibility of GraphReason and its effectiveness across various models, we additionally include *gpt-4* (OpenAI, 2023) and *PaLM-2* (Google, 2023) in our experiments. Given our limited computing resources, we utilize the same training data previously sampled from *gpt-3.5-turbo*. For testing in the GSM8K task, we select samples from 100 pieces of data from *gpt-4* and *PaLM-2*

	gpt-3.5-turbo	gpt-4	PaLM-2
Greedy Decode	72.7	87.0	53.0
Voting	82.3	94.0	71.0
Simple Verifier	66.9	89.0	36.0
Voting Verifier	85.4	97.0	77.0
DIVERSE	85.0	97.0	75.0
Ours	85.7	94.0	78.0

Table 3: The experimental results of GraphReason with different LLMs.

respectively. We conduct the sampling 10 times using three types of five exemplars, maintaining the same settings as in our previous experiments. Our method aims to enhance the original reasoning capabilities. Therefore, we do not include small-sized LMs, which typically exhibit weaker reasoning abilities.

From Table 3, it is evident that our method enhances the original reasoning performance of both *GPT-4* and *PaLM-2*. However, there is a performance decline in *gpt-4* when compared with the best baselines. The performance of GraphReason is comparable to that of the voting method. We hypothesize that this is because the reasoning patterns of *GPT-4* differ from those of *GPT-3.5-Turbo*, and our verifier is trained specifically on *GPT-3.5-Turbo* samples in this setting.

5 Conclusion

In this paper, we propose GraphReason, a novel and general method to enhance the reasoning abilities of large language models. Our method is the first to approach reasoning logic of large language models from a graph perspective and verifies candidate reasoning paths accordingly. We demonstrate the superiority of GraphReason through extensive experiments.

Limitations

There are several limitations in the current research that contribute to performance that is not as good as expected:

- **Computing Resources.** Despite the impressive performance it achieves, our framework requires large language models like GPT3.5. Inference with these models is more time-consuming and costly than fine-tuning models like BERT (Devlin et al., 2019). Some experiments, such as hyperparameter analysis, have already been conducted in related previous work and are not replicated here. Furthermore, due to limited computing resources, we have not conducted experiments with additional LLMs. We have chosen solely to use the representative LLM, GPT3.5, to compare the performance of the verifiers.
- **Labeled CoT data.** GraphReason is a complex verifier method that builds on graph classification, which requires more labeled data with well-annotated chain-of-thought reasoning paths for training. In the training of GraphReason, we use reasoning paths from LLMs’ output which may introduce significant noise. If the training data included labeled reasoning graphs, the performance would improve significantly.
- **Other Reasoning Tasks.** There are many types of reasoning tasks beyond math word problems, such as Commonsense Reasoning (Talmor et al., 2019), Inductive Reasoning (Sinha et al., 2019), etc. Given that graph construction is a complex process, we have focused mainly on solving math word problems (Arithmetic Reasoning). This focus allows for a more convenient implementation of the merging of intermediate steps. In other cases, identifying similar steps can be challenging. On the other hand, a math word problem typically presents a greater variety of potential solutions.

Nevertheless, we believe that future studies, conducted by us or others, can overcome these limitations and further improve upon our approach.

References

Abien Fred Agarap. 2019. [Deep learning using rectified linear units \(relu\)](#). *Preprint*, arXiv:1803.08375.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts,

Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pilla, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. [Palm: Scaling language modeling with pathways](#). *Preprint*, arXiv:2204.02311.

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. 2021a. [Training verifiers to solve math word problems](#). *arXiv preprint arXiv:2110.14168*.

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. 2021b. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). *Preprint*, arXiv:1810.04805.

Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. [Did Aristotle Use a Laptop? A Question Answering Benchmark with Implicit Reasoning Strategies](#). *Transactions of the Association for Computational Linguistics (TACL)*.

Google. 2023. [Palm 2 technical report](#). *Preprint*, arXiv:2305.10403.

Jie Huang and Kevin Chen-Chuan Chang. 2023. [Towards reasoning in large language models: A survey](#). *Preprint*, arXiv:2212.10403.

Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2023. [Large language models are zero-shot reasoners](#). *Preprint*, arXiv:2205.11916.

Andrew Lampinen, Ishita Dasgupta, Stephanie Chan, Kory Mathewson, Mh Tessler, Antonia Creswell, James McClelland, Jane Wang, and Felix Hill. 2022a. [Can language models learn from explanations in context?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 537–563,

- Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Andrew K Lampinen, Ishita Dasgupta, Stephanie CY Chan, Kory Matthewson, Michael Henry Tessler, Antonia Creswell, James L McClelland, Jane X Wang, and Felix Hill. 2022b. Can language models learn from explanations in context? *arXiv preprint arXiv:2204.02329*.
- Yifei Li, Zeqi Lin, Shizhuo Zhang, Qiang Fu, Bei Chen, Jian-Guang Lou, and Weizhu Chen. 2023. [Making language models better reasoners with step-aware verifier](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5315–5333, Toronto, Canada. Association for Computational Linguistics.
- Zhongli Li, Wenxuan Zhang, Chao Yan, Qingyu Zhou, Chao Li, Hongzhi Liu, and Yunbo Cao. 2022. [Seeking patterns, not just memorizing procedures: Contrastive learning for solving math word problems](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2486–2496, Dublin, Ireland. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). *Preprint*, arXiv:1711.05101.
- Shen-yun Miao, Chao-Chun Liang, and Keh-Yih Su. 2020. A diverse corpus for evaluating and developing english math word problem solvers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 975–984.
- OpenAI. 2023. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Arkil Patel, Satwik Bhattamishra, and Navin Goyal. 2021. [Are NLP models really able to solve simple math word problems?](#) In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2080–2094, Online. Association for Computational Linguistics.
- Xinyu Pi, Qian Liu, Bei Chen, Morteza Ziyadi, Zeqi Lin, Qiang Fu, Yan Gao, Jian-Guang Lou, and Weizhu Chen. 2022. [Reasoning like program executors](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 761–779, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jack W. Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, Eliza Rutherford, Tom Hennigan, Jacob Menick, Albin Cassirer, Richard Powell, George van den Driessche, Lisa Anne Hendricks, Mari-beth Rauh, Po-Sen Huang, Amelia Glaese, Johannes Welbl, Sumanth Dathathri, Saffron Huang, Jonathan Uesato, John Mellor, Irina Higgins, Antonia Creswell, Nat McAleese, Amy Wu, Erich Elsen, Siddhant Jayakumar, Elena Buchatskaya, David Budden, Esme Sutherland, Karen Simonyan, Michela Paganini, Laurent Sifre, Lena Martens, Xiang Lorraine Li, Adhiguna Kuncoro, Aida Nematzadeh, Elena Gribovskaya, Domenic Donato, Angeliki Lazaridou, Arthur Mensch, Jean-Baptiste Lespiau, Maria Tsim-poukelli, Nikolai Grigorev, Doug Fritz, Thibault Sotiaux, Mantas Pajarskas, Toby Pohlen, Zhitao Gong, Daniel Toyama, Cyprien de Masson d’Autume, Yujia Li, Tayfun Terzi, Vladimir Mikulik, Igor Babuschkin, Aidan Clark, Diego de Las Casas, Aurelia Guy, Chris Jones, James Bradbury, Matthew Johnson, Blake Hechtman, Laura Weidinger, Iason Gabriel, William Isaac, Ed Lockhart, Simon Osindero, Laura Rimell, Chris Dyer, Oriol Vinyals, Kareem Ayoub, Jeff Stanway, Lorraine Bennett, Demis Hassabis, Koray Kavukcuoglu, and Geoffrey Irving. 2022. [Scaling language models: Methods, analysis & insights from training gopher](#). *Preprint*, arXiv:2112.11446.
- Koustuv Sinha, Shagun Sodhani, Jin Dong, Joelle Pineau, and William L. Hamilton. 2019. [CLUTRR: A diagnostic benchmark for inductive reasoning from text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4506–4515, Hong Kong, China. Association for Computational Linguistics.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A question answering challenge targeting commonsense knowledge](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Karthik Valmeekam, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. 2023. [Large language models still can’t plan \(a benchmark for llms on planning and reasoning about change\)](#). *Preprint*, arXiv:2206.10498.
- Siyuan Wang, Wanjun Zhong, Duyu Tang, Zhongyu Wei, Zhihao Fan, Daxin Jiang, Ming Zhou, and Nan Duan. 2022. [Logic-driven context extension and data augmentation for logical reasoning of text](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1619–1629, Dublin, Ireland. Association for Computational Linguistics.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#). *Preprint*, arXiv:2203.11171.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022. [Chain of thought prompting elicits reasoning in large language models](#). *arXiv preprint arXiv:2201.11903*.

Zhipeng Xie and Shichao Sun. 2019. [A goal-driven tree-structured neural model for math word problems](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5299–5305. International Joint Conferences on Artificial Intelligence Organization.

Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. [How powerful are graph neural networks?](#) In *International Conference on Learning Representations*.

Ori Yoran, Alon Talmor, and Jonathan Berant. 2022. [Turning tables: Generating examples from semi-structured tables for endowing language models with reasoning skills](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6016–6031, Dublin, Ireland. Association for Computational Linguistics.

Dongxiang Zhang, Lei Wang, Luming Zhang, Bing Tian Dai, and Heng Tao Shen. 2019. [The gap of semantic parsing: A survey on automatic math word problem solvers](#). *Preprint*, arXiv:1808.07290.

Jipeng Zhang, Lei Wang, Roy Ka-Wei Lee, Yi Bin, Yan Wang, Jie Shao, and Ee-Peng Lim. 2020. [Graph-to-tree learning for solving math word problems](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3928–3937, Online. Association for Computational Linguistics.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. [A survey of large language models](#). *Preprint*, arXiv:2303.18223.

Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, and Ed Chi. 2023. [Least-to-most prompting enables complex reasoning in large language models](#). *Preprint*, arXiv:2205.10625.

A Datasets

We compared GraphReason with other methods on three different math word problem datasets: **GSM8K** (Cobbe et al., 2021a), **SVAMP** (Patel et al., 2021), and **ASDiv-a** (Miao et al., 2020) and one commonsense reasoning dataset: **StrategyQA** (Geva et al., 2021). We selected the subset ASDiv-a (arithmetic) from the original dataset ASDiv, which only involves arithmetic operations.

These three arithmetic reasoning datasets are more challenging than other math word problem datasets, making them more suitable for testing

the reasoning capability of LLMs with a verifier. As the GSM8K dataset is the only one providing step-by-step solutions as chain-of-thought exemplars, we chose exemplars from the GSM8K training dataset and tested them on all three datasets. Additionally, the training data for the verifier also used the GSM8K training data. In this setting, we could also demonstrate the transfer learning and generalization ability of our method. The size of the training split from GSM8k is 1000. The test data sizes for GSM8K, SVAMP, and ASDiv-a are 1319, 1000, and 1218, respectively.

In the StrategyQA commonsense reasoning task, we set the number of exemplars to 8 and select pseudo-exemplars from (Li et al., 2023). Additionally, we conduct five sampling iterations for each context of LLMs. From the entire dataset, we select a subset of 1,000 instances, allocating 700 for training and 300 for testing.

B Baselines

In our evaluation, we consider the following baselines:

- **Greedy Decode** is a simple method that uses a greedy decoding strategy to sample once.
- **Self-Consistency (Voting)** (Wang et al., 2023) samples multiple times and selects the final answers based on majority voting.
- **Simple Verifier** (Cobbe et al., 2021b), which is also known as the Sampling and Re-ranking strategy, uses a verifier to assign scores to sampled solutions and selects the final answer with the highest score.
- **Voting Verifier** (Li et al., 2023) combines the Voting and Verifier approaches. It assigns total scores to answers from scores of all candidate solutions and selects the final answer with the highest score.
- **DIVERSE (Step-aware Voting Verifier)** (Li et al., 2023), which is the state-of-the-art method, considers the reasoning steps throughout the entire reasoning path. It recognizes that not all steps in an incorrect reasoning path are equally wrong and that some steps may still be useful for reasoning.

We primarily compare GraphReason with other verifiers using the same generated solutions from *gpt-3.5-turbo*. Additionally, we include some previous Fine-tuning state-of-the-art methods to reflect

the strong reasoning ability of LLMs. The previous Fine-tuning SOTA methods are denoted as follows: a: (Cobbe et al., 2021b), b: (Pi et al., 2022), c: (Miao et al., 2020), d: (Chowdhery et al., 2022).