

Achieving >97% on GSM8K: Deeply Understanding the Problems Makes LLMs Perfect Reasoners

Anonymous ACL submission

Abstract

Chain of Thought prompting strategy has enhanced the performance of Large Language Models (LLMs) across various NLP tasks. However, it still has shortcomings when dealing with complex reasoning tasks, following Wei et al. (2022), including understanding errors, calculation errors and process errors (e.g. missing-step and hallucinations). Subsequently, Our in-depth analysis of various error types has found that deeply understanding the whole problem is critical in addressing complicated reasoning tasks. In this paper, we proposed a novel prompt strategy called Deeply Understanding the Problems (DUP) prompting, inspired by how humans solve complex reasoning problems, designed to enhance the comprehensive understanding of problems by LLMs. It consists of three stages: 1) extract the core question; 2) find out problem-solving information based on the core question; 3) generate and extract answers by LLMs. We evaluate the performance of DUP prompting on ten diverse reasoning datasets. Experimental results suggest that DUP prompting significantly outperforms Zero-Shot CoT (Kojima et al., 2022) across all datasets. Notably, DUP achieves **state-of-the-art on SVAMP (90.4% to 94.2%) and GSM8K (94.6% to 97.1%)**.

1 Introduction

Despite the impressive performance of Large Language Models (LLMs) in diverse NLP tasks (Brown et al., 2020; Thoppilan et al., 2022; Chowdhery et al., 2022), they often suffer from limited reasoning abilities, a challenge that cannot be adequately addressed by simply scaling up the model sizes (Wang et al., 2023b). The limitation highlights the importance of further improving LLMs’ reasoning capabilities. To this end, Wei et al. (2022) proposes a few-shot chain of thought prompting, by allowing the language model to perform natural language reasoning before giving the final answer,

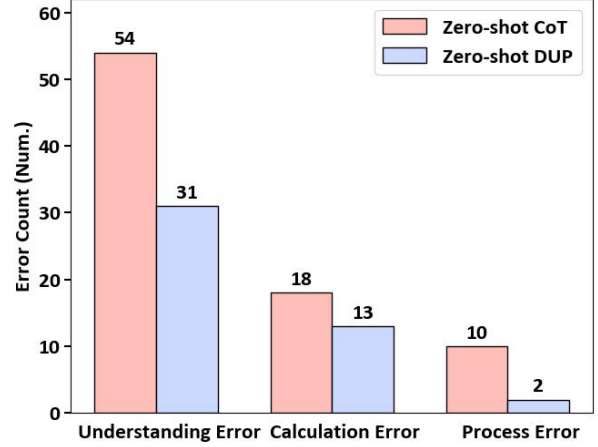


Figure 1: **Zero-shot CoT and DUP Prompting error analysis of failures examples** return by GPT-3.5-turbo LLM. Compared to Zero-shot CoT, DUP Prompting reduced 23 “Understanding Error” (from 54 to 31), 5 “Calculation Error” (from 18 to 13), and 8 “Process Error” (from 10 to 2).

which is demonstrated significant performance in various reasoning tasks.

In the current landscape of prompt methods, the main focus of researchers lies in modifying the prompt strategy to guide language models towards improved reasoning step quality, such as Zero-shot CoT (Kojima et al., 2022), Tree of Thought (ToT) (Gao et al., 2023), Plan and Solve (PS) (Wang et al., 2023a), and Complex CoT (Fu et al., 2023). However, these methodologies often fail to address a crucial aspect: **whether the LLMs fully understand the whole problem?** As shown in Figure 1 Zero-shot CoT strategy yielded 82 incorrect responses returned by GPT-3.5-turbo LLM from random 300 samples on the GSM8K dataset. Notably, **understanding errors accounted for 65.8% of total errors**. The results show that LLMs are insufficient for comprehensively understanding the whole problem. Unfortunately, the issue has not been given adequate attention.

To overcome the problem, in this paper, we propose DUP Prompting. Different from the previous

methods, DUP prompting improves reasoning ability by helping LLMs better understand the question. Specifically, DUP consists of three stages: ❶ Extract the core question from the original input by LLMs; ❷ Extract problem-solving information required to solve the core question; ❸ LLMs generates responses by combining core questions with problem-solving information. Subsequently, leveraging LLMs separates answers from the complex text generated. The first two stages enable LLMs to gain a clear and comprehensive understanding of the problem, including the goal of the question and the conditions required to solve it. The last stage is designed to generate more accurate reply text and answers by LLMs.

To evaluate the performance of the DUP prompting strategy, we conducted a series of experiments on ten datasets across arithmetic, commonsense, and symbolic reasoning benchmarks. The results of GPT-3.5-Turbo (Ouyang et al., 2022) and GPT-4 (OpenAI, 2023) show that DUP prompting consistently outperforms Zero-shot CoT across all reasoning datasets by a large margin. Furthermore, DUP prompting is also exceeds Zero-Shot PS+ (Wang et al., 2023a) and Least-to-Most (Zhou et al., 2023) prompting. Surprisingly, DUP prompting modestly improved compared to Few-shot manual CoT (Wei et al., 2022) and auto CoT prompting strategy (Zhang et al., 2023b) on most reasoning datasets, despite DUP prompting does not require any manual demonstration examples. Overall, DUP prompts is capable of generating higher-quality reasoning steps and more accurate answers.

In summary, our **main contributions** include:

- We analyzed various types of errors in reasoning tasks and found that understanding the whole problem is key to solving it.
- Based on this error analysis, we propose DUP Prompting, designed to enhance the comprehensive understanding of problems.
- Evaluations across multiple reasoning benchmarks demonstrated significant improvements compared to previous zero-shot methods.

2 Related Work

2.1 Reasoning with Large Language Models.

With the revolutionary development of pre-training (Brown et al., 2020; Devlin et al., 2019; Qiu et al., 2020), language models demonstrated to confer

a range of reasoning abilities by scaling up the size. (OpenAI, 2023; Touvron et al., 2023; Chowdhery et al., 2022) show that large language models (LLMs) have remarkable success across a range of NLP tasks. However, LLMs still struggle to provide stable and accurate answers when faced with complex reasoning tasks (Zhang et al., 2023a), such as mathematical reasoning (Cobbe et al., 2021; Patel et al., 2021; Ling et al., 2017; Hosseini et al., 2014), commonsense reasoning (Talmor et al., 2019; Geva et al., 2021) and symbolic reasoning (Wei et al., 2022).

Recent work by (Yuan et al., 2023; Luo et al., 2023; Yu et al., 2023) shown that adopted supervised fine-tuning LLMs can improves reasoning capabilities by a large margin. However, even with such advancements, these models still perform poorly in complex reasoning problems. This may be due to the fact that fine-tuning is not enough to unlock LLMs’ reasoning capabilities.

2.2 Prompting Methods.

It is well known that the language model directly outputs the answer (without any intermediate reasoning process) to complex reasoning tasks, it would be pretty hard to get that question right. To overcome this problem, Wei et al. (2022) propose Few-shot Chain-of-Thought prompting, which elicits a series of intermediate natural language reasoning steps before giving the final answer. So far, the chain of thought prompting has been proven to significantly improve the reasoning capabilities of language models. For one thing, a large amount of work emerged to improve language model reasoning capabilities by improving prompting strategy (Zhou et al., 2023; Wang et al., 2023a; Yao et al., 2023; Brown et al., 2020; Kojima et al., 2022; Zhang et al., 2023b; Chen et al., 2023; Xu et al., 2023). For another, optimizing the rationales for better consistency and continuity (Wang et al., 2023b; Li et al., 2023; Sanh et al., 2022; Zheng et al., 2023). Specifically, Kojima et al. (2022) illustrated that LLMs are also zero-shot reasoners without needing extra exemplars. By only concatenating “Let’s think step by step” to the end of the prompt, it can consciously generate reasoning steps. Zhou et al. (2023) propose to decompose the prompt first into subproblems, and then sequentially solve the subproblems. Solving each subproblem is facilitated by the answers to previously solved subproblems. Wang et al. (2023a) pro-

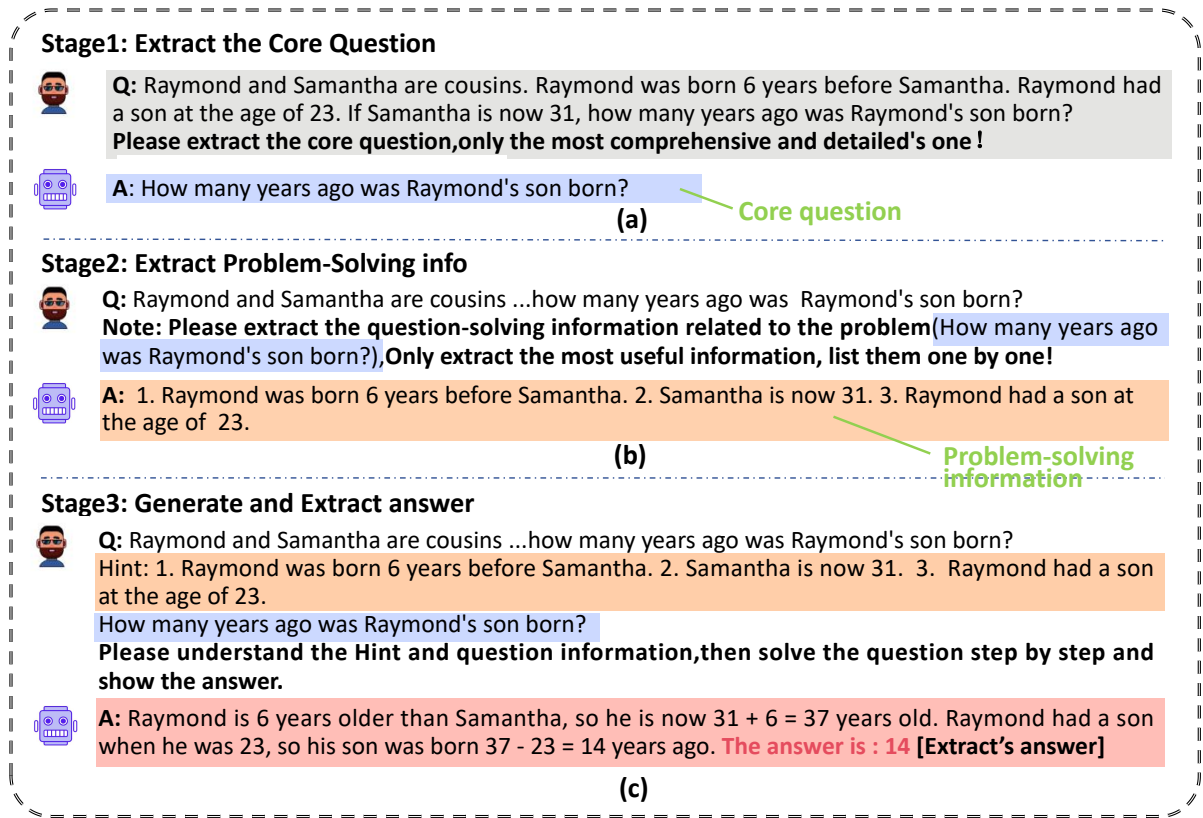


Figure 2: **Detailed illustration of DUP prompting strategy.** ❶ Extracting the core question from the original input; ❷ Subsequently, extracting problem-solving information based on the core question; ❸ Requiring LLMs to understand the core question and the problem-solving information to generate an answer text. Then, isolating the final answer from the generated complex text.

pose to guide LLMs in extracting variables related to the problem and their corresponding numbers, formulating a complete plan, and solving the problem step by step, which can be easily customized to solve various reasoning problems. Zhang et al. (2023b) propose to divide the problems of a given dataset into several clusters. Then, a representative problem is selected from each cluster and its reasoning chain is generated using zero-shot CoT.

Different from the above works, our focus is on LLMs to fully understand the whole problem to avoid various errors in the reasoning steps. It consists of three stages, where the first two stages are designed to enhance LLMs' comprehensive understanding of the whole problem to avoid understanding errors and process errors. At the end of the stage, it aims to avoid calculation errors and extract answers accurately.

3 DUP Prompting

Overview. We introduce a new zero-shot CoT prompting approach, called DUP prompting. Specifically, in stage 1, extracts the **core question** from complex and lengthy problem descriptions.

In stage 2, DUP prompting further extracts the **problem-solving information** that is crucial for solving the core question from the same description. Given the core question and problem-solving information. At the end of the stage, DUP incorporates them into the origin question to generate the reasoning process as well as the answer to the question. Then, LLMs extract the final answer from the generated text.

3.1 Stage 1: Extracting the Core Question

Understanding the goal of a question is the first step to solving it, even for humans. Unfortunately, LLMs may be confused by lengthy descriptions of complex reasoning questions, leading to inaccurate understanding and an inability to solve the goal. To deal with this problem, we use LLMs to explicitly extract the core question from the origin input before reasoning. Specifically, we design a core question extraction template "Please extract core question, only extract the most comprehensive and detailed one!" It appends to the end of question and use GPT-3.5-turbo (Ouyang et al., 2022) to extract the core question from the input. Subsequently, the

output of this step will be a shorter, clearer question that will be used to help LLMs focus on the goal of input questions in subsequent steps.

3.2 Stage 2: Extracting Problem-Solving Information

In addition to clarifying the goal, it is also important to find the conditions required to solve the problem. Without fully understanding and utilizing the conditions provided by the question, reasoning cannot be correctly completed. LLMs also have trouble taking full advantage of these conditions. Therefore, we design a problem-solving information extraction template to help solve this problem. Specifically, the template “\nNote: Please extract the problem-solving information related to the core question [Core Question info], Only extract the most useful information, list them one by one.” to extract all useful information required to solving the core question. The input slot contains the core question extracted in Stage 1. The output of this step will be a list of conditions useful in reasoning.

3.3 Stage 3: Generate and Extract Answers

Given the core question and problem-solving information extracted in previous stages, we incorporate them into the original input by the template “\nHint: [Problem-Solving Info]\n[Core Question]\n Please understand the Hint and question information, then solve the problem step by step and show the answer.”, where the input slots refer to the corresponding outputs in previous steps. This prompt is beneficial to improve LLM’s understanding of the question by explicitly pointing out the goal and necessary conditions to solve the question.

Similar to previous work in Plan and Solve (Wang et al., 2023a), we use LLMs to extract the final numerical answer from a long reasoning text generated. Compared with rule-based matching methods, using LLMs to extract the final answers is more robust in practice. More template details for separated answers can be found in Appendix A.1.

4 Experimental Setup

4.1 Benchmarks

To evaluate the performance of the DUP prompting strategy across ten reasoning benchmark datasets from three categories of reasoning tasks:

Dataset	Avg. words	# Samples	Answer Format
GSM8K	46.9	1319	Number
MultiArith	31.8	600	Number
AddSub	31.5	395	Number
SVAMP	31.8	1000	Number
SingleEq	27.4	508	Number
AQuA	51.9	254	Option
Last Letters	15.0	500	String
Coin Flip	37.0	500	Yes / No
StrategyQA	9.6	2290	Yes / No
CSQA	27.8	1221	Option

Table 1: **Details of datasets being evaluated.** CSQA mean is CommonsenseQA.

Arithmetic Reasoning, we consider the following six arithmetic reasoning benchmarks: (1) the GSM8K (Cobbe et al., 2021), (2) the SVAMP (Patel et al., 2021), (3) the MultiArith (Roy and Roth, 2015), (4) the AddSub (Hosseini et al., 2014), (5) the AQuA (Ling et al., 2017), and (6) the SingleEq (Koncel-Kedziorski et al., 2015);

Commonsense Reasoning, we include two datasets: (7) the CommonsenseQA (Talmor et al., 2019); and (8) the StrategyQA (Geva et al., 2021);

Symbolic Reasoning, we include two datasets: (9) the Last Letter (Wei et al., 2022), and (10) the Coin Flip (Wei et al., 2022). Table 1 shows dataset statistics.

4.2 Baselines

We introduce five baselines including Zero-shot and Few-shot to compare with Zero-shot DUP prompting:

- **Zero-Shot-CoT** (Kojima et al., 2022): It appends “Let’s think step by step.” to the end of prompt.
- **Plan-and-Solve** (Wang et al., 2023a): It enables LLMs to explicitly devise a plan for solving a given problem. We adopted the Plan-and-Solve Prompting strategy with more detailed instructions(PS+).
- **Least-to-Most** (Zhou et al., 2023): The initial input is prefixed with “Let’s break down this problem:”. Subsequently, the initial response from the LLMs is appended to this prompt, followed by a newline and “The answer is:”.
- **Manual-CoT** (Wei et al., 2022): Design of eight hand-crafted examples as demonstrations.

Model	Method	Arithmetic Reasoning						Score	
		SVAMP	GSM8K	AddSub	MultiArith	AQuA	SingleEq	Avg.	Δ
Performance of Zero-shot Strategy									
GPT-3.5-Turbo	Zero-shot CoT	79.3	78.9	85.8	95.3	53.0	93.5	80.9	-
	Least-to-Most	80.9	77.5	91.3	95.5	57.4	93.5	82.6	+1.7
	Zero-shot PS+	80.7	79.3	86.5	92.0	55.9	93.0	80.4	+0.3
	Zero-shot DUP	82.5	82.3	92.1	97.8	60.2	94.9	84.9	+4.0
GPT-4	Zero-shot CoT	90.4	94.6	92.4	97.8	72.8	95.0	90.6	-
	Least-to-Most	90.3	92.1	92.1	97.1	71.6	95.0	89.7	-0.9
	Zero-shot PS+	92.6	94.3	93.1	98.1	75.5	95.3	91.4	+0.8
	Zero-shot DUP	94.2	97.1	95.1	98.1	77.1	96.0	92.9	+2.3
Performance of Few-shot Strategy									
GPT-3.5-Turbo	Manual CoT	78.5	81.6	90.6	95.6	55.9	94.2	82.6	+1.7
	Auto CoT	82.9	80.2	89.9	99.0	54.3	94.6	83.4	+2.5

Table 2: **Results for various Zero-shot and Few-shot reasoning approaches** on six math reasoning datasets. The best results are in **bold**. “ Δ ” denotes the **improvement** or **decline** of various methods compared to Zero-shot CoT.

Method	CSQA	StrategyQA	Avg.
Zero-shot-CoT	72.3	66.1	69.2
Zero-shot-PS+	68.8	62.8	65.8
Least-to-Most	71.9	61.5	66.7
Zero-shot DUP	74.5	68.5	71.5
Few-shot-CoT (Manual)	76.5	64.8	70.8
Few-shot-CoT (Auto)	74.2	62.5	68.3

Table 3: **Results for various approaches using GPT-3.5-turbo LLM on commonsense reasoning datasets.**

Method	Last Letter	Coin Flip	Avg.
Zero-shot CoT	60.8	94.4	77.6
Zero-shot PS+	60.6	95.4	78.0
Least-to-Most	83.2	82.8	83.0
Zero-shot DUP	81.2	97.6	89.4
Few-shot CoT (Manual)	74.4	98.2	86.3
Few-shot CoT (Auto)	81.2	98.6	89.9

Table 4: **Results for various approaches using GPT-3.5-turbo LLM on symbolic reasoning datasets.**

- Auto-CoT (Zhang et al., 2023b): It samples questions with diversity and generates reasoning chains to construct demonstrations.

4.3 Implementations

We use the public GPT-3.5-Turbo (0613) (Ouyang et al., 2022) and GPT-4 (0613) (OpenAI, 2023) as the backbone language models. In this paper, all models are employed via OpenAI’s API¹ key, and we adopt the greedy decoding strategy with a temperature setting of 0 across all experiments. In the Few-Shot baselines, we adhere to the recommended number of demonstration examples specified in the original paper, Manual-CoT (Wei et al., 2022) and Auto-CoT (Zhang et al., 2023b).

¹<https://beta.openai.com/docs/models>

5 Experimental Results

5.1 Main Results

Arithmetic Reasoning. In Table 2 presents the main results, which include all Zero-shot and Few-shot baselines. Our findings reveal that DUP prompting significantly enhances reasoning performance compared to Zero-shot CoT on arithmetic reasoning datasets. Specifically, in the GPT-3.5-turbo settings, DUP prompting improves accuracy by an average of 4% over Zero-shot CoT. In the GPT-4 settings, DUP prompting achieves state-of-the-art results on SVAMP (94.2%) and GSM8K (97.1%). Additionally, DUP prompting consistently outperforms other Zero-shot baselines across all arithmetic reasoning datasets.

We also reports the few-shot methods results. Due to the high cost of the GPT-4 API, we opted to use the more affordable GPT-3.5-turbo as a responder for few-shot methods, including Manual CoT (Wei et al., 2022) and Auto CoT (Zhang et al., 2023b). Generally, the performance of zero-shot methods tends to be lower than that of few-shot methods. However, within the GPT-3.5-turbo settings, DUP prompting demonstrates superior performance with an average accuracy of 84.9%, surpassing both Auto-CoT (82.7%) and Manual-CoT (81.4%). The results suggest that deeply understanding the whole problem enables enhancing the response quality of LLMs’ reasoning.

Commonsense and Symbolic Reasoning. Table 3 shows the performance on commonsense reasoning datasets. Considering the experimental cost, we only used the GPT-3.5-turbo as the backbone language model. Compared to zero-shot meth-

Stage 1	Stage 2	Stage 3	GSM8K	AQuA	Avg.
✗	✗	✗	76.5	51.2	63.8
✓	✗	✗	78.9	53.1	66.0
✗	✓	✗	80.6	55.1	67.8
✗	✗	✓	80.3	54.7	67.5
✓	✓	✗	79.9	57.0	68.4
✓	✗	✓	80.8	56.2	68.5
✗	✓	✓	81.7	58.2	69.9
✓	✓	✓	82.3	60.2	71.2

Table 5: **Ablation study for different variations of DUP prompting** using GPT-3.5-turbo LLMs on GSM8K and AQuA Datasets: Stage 1 involves extracting core questions, Stage 2 focuses on extracting problem-solving information, and Stage 3 entails solving the problem step by step.

ods, DUP prompting consistently outperforms all other baseline zero-shot methods. In comparison with Few-shot methods, the accuracy of zero-shot DUP prompting (74.5%) is only lower than that of Manual-CoT (76.5%) on the CommonsenseQA dataset. However, the average accuracy of DUP prompting is still better than few-shot methods

Table 4 reveals the results on Symbolic reasoning datasets. On Last Letters, zero-shot DUP (81.2%) is marginally worse than Zero-shot-Least-to-Most (83.2%) but significantly exceeds other Zero-shot approaches, including Manual-CoT (74.4%) and is on par with Auto-CoT (81.2%). on Coin Flip, zero-shot DUP (97.6%) is slightly below Manual-CoT (98.2%) and Auto-CoT (98.6%) but better than other zero-shot baseline methods by a good margin. More examples from the experiment results can be found in Appendix A.3.

5.2 Ablation Study

We observed a significant enhancement in the accuracy of reasoning tasks when using DUP prompting. To investigate the impact of each stage and inference cost reduction on performance, we conducted a series of ablation experiments across various reasoning tasks, as described below.

Each stage is important for solving reasoning problems. We conducted an ablation study on each stage to explore the effectiveness of DUP prompting, as shown in Table 5, which presents eight variations of the DUP prompting strategy. The results show that each prompting strategy is consistently better than Baseline (*no prompt strat-*

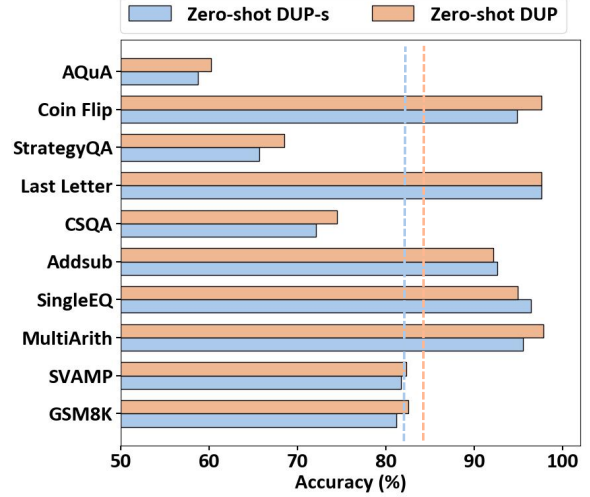


Figure 3: **Performance of DUP and DUP-s methods across various reasoning tasks on GPT-3.5-Turbo.** Orange dashline and Blue dashline represent the average accuracy on DUP and DUP-s respectively.

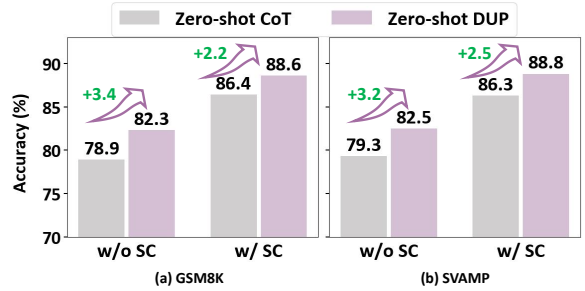


Figure 4: **Results of DUP Prompting with and without self-consistency(SC)** using GPT-3.5-turbo LLM on GSM8K and SVAMP.

egy). On the other hand, it works best only when three stages are combined on GSM8K and AQuA, further verifying that each stage is indispensable.

Reduce inference cost but sacrifice accuracy.

To measure whether the performance can be consistent after reducing the inference cost. We merge the three stages of the DUP prompting strategy, called DUP-s prompting, and conduct experiments using across all reasoning datasets. Surprisingly, DUP-s has better accuracy than DUP in AddSub and SingleEq datasets (92.6% vs. 92.1%) and (96.4% vs. 94.9%). But overall, the average performance of DUP-s prompting dropped slightly compared to DUP prompting (83.6% vs. 84.8%).

6 Analysis

Compatibility with Self-consistency. We employ an innovative decoding strategy, with self-consistency (SC) (Wang et al., 2023b), as a substitute for the conventional greedy decoding approach,

Model	Method	GSM8K	AddSub	Avg.
Llama-2-Chat 13b	Zero-shot CoT	35.1	70.6	52.8
	Zero-shot DUP	35.9(+0.8)	79.7(+9.1)	57.8(+5.0)
Llama-2-Chat 70b	Zero-shot CoT	53.9	75.6	64.7
	Zero-shot DUP	56.4(+2.5)	87.8(+12.2)	72.1(+7.4)
CodeLlama-Instruct 13b	Zero-shot CoT	24.2	73.1	48.6
	Zero-shot DUP	28.1(+3.9)	74.6(+1.5)	51.3(+2.7)
CodeLlama-Instruct 34b	Zero-shot CoT	39.1	81.2	60.1
	Zero-shot DUP	43.5(+4.4)	86.0(+4.8)	64.7(+4.1)

Table 6: **Results of various large language models (LLMs) on GSM8K and Addsub.** Compared with Zero-shot CoT, DUP prompting still has outstanding performance.

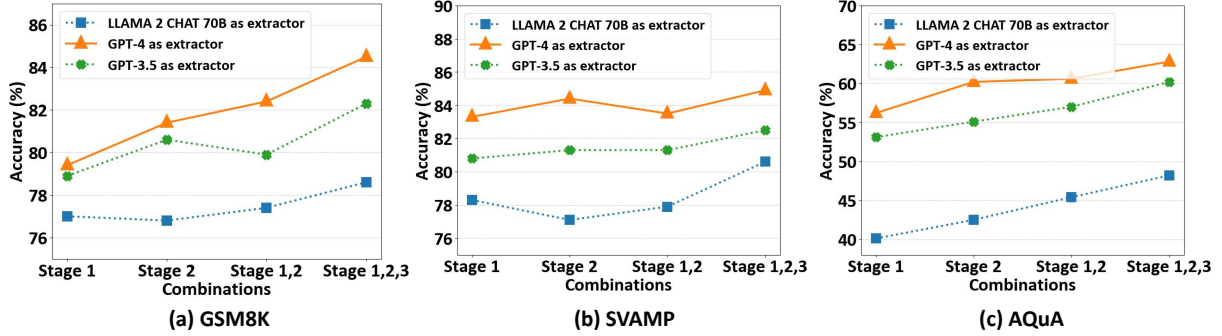


Figure 5: **Results of each stage of DUP prompting strategy on GSM8K, AQuA and SVAMP** using GPT-4, GPT-3.5-turbo and Llama-2-Chat 70b as core question and problem-solving information extractor. Finally, GPT-3.5-turbo llms as responder.

which initially samples a diverse array of reasoning paths rather than only opting for the greedy approach. Subsequently, choosing the most consistent answer as the answer. Existing research (Wang et al., 2023a; Xu et al., 2023) indicates that adopted SC notably enhances the performance of chain-of-thought prompting. We are also curious whether using SC can further enhance the performance of DUP prompting. Hence, to evaluate the impact of integrating DUP prompting and self-consistency on GSM8K and SVAMP using GPT-3.5-Turbo, setting the temperature to 0.7 and N to 10 for experiments with SC. When combined with SC, the results revealed in Figure 4 demonstrate a significant enhancement in the performance of both zero-shot CoT and DUP prompting. Notably, DUP prompting with SC(88.6% and 88.8%) consistently outperforms Zero-shot CoT with SC(86.4% and 86.3%), indicating its compatibility with the SC approach.

Performance of DUP prompting on other LLMs.

To further evaluate the performance of DUP prompting in LLMs with relatively smaller parameters, we adopted Llama-2-Chat 13b and Llama-2-Chat 70b (Touvron et al., 2023), CodeLlama-

Instruct 13b and CodeLlama-Instruct 34b (Roziere et al., 2023) as backbone LLMs evaluated on GSM8K and AddSub. As shown in Table 6, the performance of DUP prompting is significantly ahead of Zero-shot CoT method no matter which one LLMs. Moreover, the average performance of DUP prompting increased by 7.4% in Llama-2-Chat 70b.

The accuracy of Core questions and Problem-solving information is beneficial to improving DUP prompt performance.

To examine the impact of core question and problem-solving information quality on the DUP prompting strategy. Our tests spanned across AQuA, GSM8K, and SVAMP datasets for each stage of DUP prompting. The results, as illustrated in 5, indicate a significant performance boost in DUP prompting when using GPT-4 as the extractor. On the other hand, using Llama-2-Chat 70b for information extraction leads to a significant decrease in accuracy. This outcome highlights the critical impact of extracted information quality on overall performance. More detailed analysis are in Appendix A.2.

Error Analysis. To further analyze the effect of DUP prompting on Understanding errors, Process

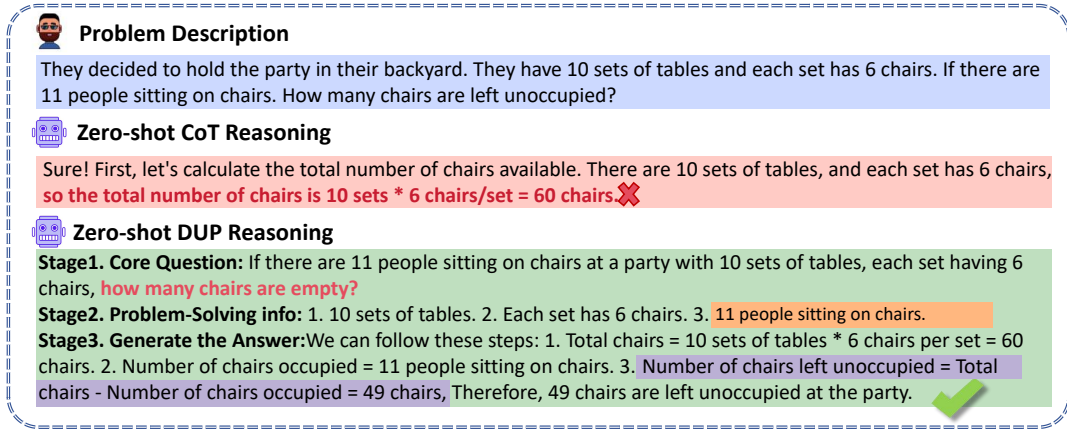


Figure 6: **Case study about the comparison between Zero-shot CoT and Zero-shot DUP prompting.** Examples from the SVAMP dataset, with solutions generated by Zero-shot CoT and DUP. Zero-shot CoT will generate answers directly through a chain of thought. However, DUP generates answers that combine problem-solving information with the core question.

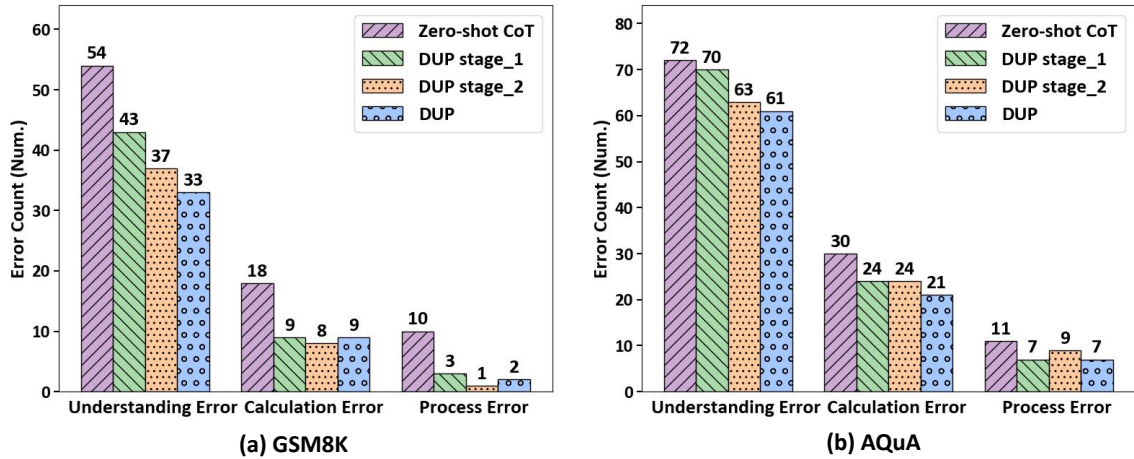


Figure 7: **Quantitative analysis of various types of errors between CoT and variant methods of DUP prompting.** Randomly select 300 examples on GSM8K and AQuA. Subsequently, use GPT-3.5-Turbo LLMs to generate responses and count failed answers. We can see that our method reduces the frequency of various error types compared with Zero-shot CoT.

errors and Calculation errors. We first separate the core problem and problem-solving information from DUP prompting, and called DUP stage_1 and DUP stage_2 methods respectively. Subsequently, the total number of three major types of errors was counted. As illustrated in Figure 7, the variant method of DUP prompting significantly reduces the frequency of various error types compared with Zero-shot CoT. For error statistics and further analysis, see Appendix. A.2.

7 Conclusion

In this work, we discovered that a deep understanding of the whole problem is crucial for tackling complex reasoning tasks. Consequently, we introduce the DUP prompting, which consists of

three stages focused on assisting LLMs in gaining a more accurate and comprehensive understanding of the problem. A series of experiments on arithmetic, commonsense, and symbolic reasoning tasks demonstrated that DUP prompting significantly improves the reasoning abilities of Large Language Models (LLMs). It consistently outperforms previous Zero-shot methods and modestly surpasses Few-shot Manual CoT and Auto CoT across most reasoning datasets. Further analysis confirms the efficiency and generalization of our method. Overall, our approach is a general strategy which can be applied to various reasoning tasks. Moreover, considering whether LLMs fully understand the whole problem maybe also be beneficial to non-reasoning tasks. We leave them for future work.

Limitations

There are two limitations to this work. First, DUP prompting necessitates three visit to LLMs, which significantly increases the inference cost. Although we attempted to merge the three stages of the DUP strategy into a single stage, this approach unfortunately led to a decrease in accuracy. Secondly, DUP prompting are helpful in reducing the frequency of various types of errors. However, understanding errors remain prevalent in across all reasoning tasks. We will further explore how to solve understanding errors and reduce the cost of reasoning.

Ethic Statements

We take ethical considerations very seriously and strictly adhere to the ACL Ethics Policy. This paper aims to improve the large language model reasoning through a novel prompt strategy. All used model (or APIs) and datasets in this paper are publicly available and have been widely adopted by researchers. All experimental results upon these open models and datasets are reported accurately and objectively. Thus, we believe that this research will not pose any ethical issues.

References

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *NeurIPS*.

Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W. Cohen. 2023. [Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks](#). *TMLR*.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. [PaLM: Scaling language modeling with pathways](#). *arXiv preprint*.

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

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

Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. 2023. [Complexity-based prompting for multi-step reasoning](#). In *ICLR*.

Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. [Pal: Program-aided language models](#). *arXiv preprint*.

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](#). *TACL*.

Mohammad Javad Hosseini, Hannaneh Hajishirzi, Oren Etzioni, and Nate Kushman. 2014. [Learning to solve arithmetic word problems with verb categorization](#). In *EMNLP*.

Takeshi Kojima, Shixiang (Shane) Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). In *NeurIPS*.

Rik Koncel-Kedziorski, Hannaneh Hajishirzi, Ashish Sabharwal, Oren Etzioni, and Siena Dumas Ang. 2015. [Parsing algebraic word problems into equations](#). *ACL*.

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 *ACL*.

Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. 2017. [Program induction by rationale generation: Learning to solve and explain algebraic word problems](#). In *ACL*.

Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng, Qingwei Lin, Shifeng Chen, and Dongmei Zhang. 2023. [Wizardmath: Empowering mathematical reasoning for large language models via reinforced evol-instruct](#). *arXiv preprint*.

OpenAI. 2023. [Gpt-4 technical report](#).

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *NeurIPS*.

Arkil Patel, Satwik Bhattamishra, and Navin Goyal. 2021. [Are NLP models really able to solve simple math word problems?](#) In *NAACL*.

571	XiPeng Qiu, TianXiang Sun, YiGe Xu, YunFan Shao,	Xiaohan Xu, Chongyang Tao, Tao Shen, Can Xu,	625
572	Ning Dai, and XuanJing Huang. 2020. Pre-trained	Hongbo Xu, Guodong Long, and Jian-guang Lou.	626
573	models for natural language processing: A survey.	2023. Re-reading improves reasoning in language	627
574	<i>Science China Technological Sciences.</i>	models. <i>arXiv preprint.</i>	628
575	Subhro Roy and Dan Roth. 2015. Solving general arith-	Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran,	629
576	metic word problems. In <i>EMNLP</i> .	Thomas L. Griffiths, Yuan Cao, and Karthik	630
577	Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten	Narasimhan. 2023. Tree of thoughts: Deliberate	631
578	Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi,	problem solving with large language models.	632
579	Jingyu Liu, Tal Remez, J��r��my Rabin, et al. 2023.	Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu,	633
580	Code llama: Open foundation models for code. <i>arXiv</i>	Zhengying Liu, Yu Zhang, James T. Kwok, Zhenguo	634
581	<i>preprint.</i>	Li, Adrian Weller, and Weiyang Liu. 2023. Meta-	635
582	Victor Sanh, Albert Webson, Colin Raffel, Stephen	math: Bootstrap your own mathematical questions	636
583	Bach, Lintang Sutawika, Zaid Alyafeai, Antoine	for large language models. <i>arXiv preprint.</i>	637
584	Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey,	Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting	638
585	M Saiful Bari, Canwen Xu, Urmish Thakker,	Dong, Keming Lu, Chuanqi Tan, Chang Zhou, and	639
586	Shanya Sharma Sharma, Eliza Szczechla, Taewoon	Jingren Zhou. 2023. Scaling relationship on learning	640
587	Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti	mathematical reasoning with large language models.	641
588	Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han	<i>arXiv preprint.</i>	642
589	Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong,	Yifan Zhang, Jingqin Yang, Yang Yuan, and Andrew	643
590	Harshit Pandey, Rachel Bawden, Thomas Wang, Tr-	Chi-Chih Yao. 2023a. Cumulative reasoning with	644
591	ishala Neeraj, Jos Rozen, Abheesht Sharma, An-	large language models. <i>arXiv preprint.</i>	645
592	drea Santilli, Thibault Fevry, Jason Alan Fries, Ryan	Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex	646
593	Teehan, Teven Le Scao, Stella Biderman, Leo Gao,	Smola. 2023b. Automatic chain of thought prompt-	647
594	Thomas Wolf, and Alexander M Rush. 2022. Multi-	ing in large language models. In <i>ICLR</i> .	648
595	task prompted training enables zero-shot task gener-	Chuanyang Zheng, Zhengying Liu, Enze Xie, Zhenguo	649
596	alization. In <i>ICLR</i> .	Li, and Yu Li. 2023. Progressive-hint prompting	650
597	Alon Talmor, Jonathan Herzig, Nicholas Lourie, and	improves reasoning in large language models.	651
598	Jonathan Berant. 2019. Commonsenseqa: A question	Denny Zhou, Nathanael Sch��rli, Le Hou, Jason Wei,	652
599	answering challenge targeting commonsense knowl-	Nathan Scales, Xuezhi Wang, Dale Schuurmans,	653
600	edge. In <i>NAACL</i> .	Claire Cui, Olivier Bousquet, Quoc V Le, and Ed H.	654
601	Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam	Chi. 2023. Least-to-most prompting enables com-	655
602	Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng,	plex reasoning in large language models. In <i>ICLR</i> .	656
603	Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al.		
604	2022. Lamda: Language models for dialog applica-		
605	tions. <i>arXiv preprint.</i>		
606	Hugo Touvron, Louis Martin, Kevin Stone, Peter Al-		
607	bert, Amjad Almahairi, Yasmine Babaei, Nikolay		
608	Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti		
609	Bhosale, et al. 2023. Llama 2: Open foundation and		
610	fine-tuned chat models. <i>arXiv preprint.</i>		
611	Lei Wang, Wanyu Xu, Yihuai Lan, Zhiqiang Hu,		
612	Yunshi Lan, Roy Ka-Wei Lee, and Ee-Peng Lim.		
613	2023a. Plan-and-solve prompting: Improving zero-		
614	shot chain-of-thought reasoning by large language		
615	models. In <i>ACL</i> .		
616	Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le,		
617	Ed H. Chi, Sharan Narang, Aakanksha Chowdhery,		
618	and Denny Zhou. 2023b. Self-consistency improves		
619	chain of thought reasoning in language models. In		
620	<i>ICLR</i> .		
621	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten		
622	Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022.		
623	Chain of thought prompting elicits reasoning in large		
624	language models. In <i>NeurIPS 2022</i> .		

A Appendix

A.1 Prompt details.

In this selection, We show the prompts used in this study, covering the inference prompts, extract answer prompts and error analysis prompts. Firstly, table 7 shows the inference template for all reasoning tasks. Secondly, tables 8, 9 and 10 list the extract answer prompts we have tried for each dataset. Finally, As shown in the Figure 8, this prompt is designed to categorize failure examples.

A.2 More analysis.

In this selection, we further analyze why DUP prompting is efficient. It mainly covers the following three aspects: ❶ Compared to Zero-shot CoT, our methods introduce fewer errors or solve more problems? ❷ The impact of the quality of extracted core question and problem-solving information on the performance of DUP prompting. ❸ Perform statistics and analysis of error types on more reasoning dataset.

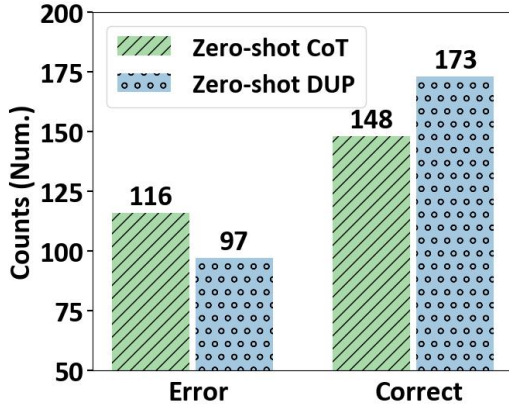


Figure 9: Distribution of the number of new failure examples (Error) and solved problems (Correct) introduced by Zero-shot CoT and DUP prompting compared to Baseline (no prompt strategy).

Introduce fewer errors while solving more problems. What causes the performance difference between DUP and CoT prompting strategies compared to the baseline: **Solving more problems or reducing more errors?** As shown in 9, the CoT strategy successfully solved an additional 148 problems, but introduced 116 new errors. In comparison, the DUP strategy not only surpassed CoT in the number of problems solved (173 vs 148) but also introduced fewer bugs, at only 97. This finding highlights the importance of considering error introduction when optimizing a prompting strategy.

Evaluate the quality of the extracted core question and problem-solving information. Furthermore, we randomly sampled 100 examples on GSM8K and used GPT-4 and GPT-3.5-Turbo to extract core questions and problem-solving information respectively. Subsequent manual evaluation, as shown in Figure 10, shows GPT-4 achieving a 20% improvement in core question and a 13% increase in problem-solving information over GPT-3.5-Turbo. The result shows that more accurate information is beneficial to improve DUP prompting performance.

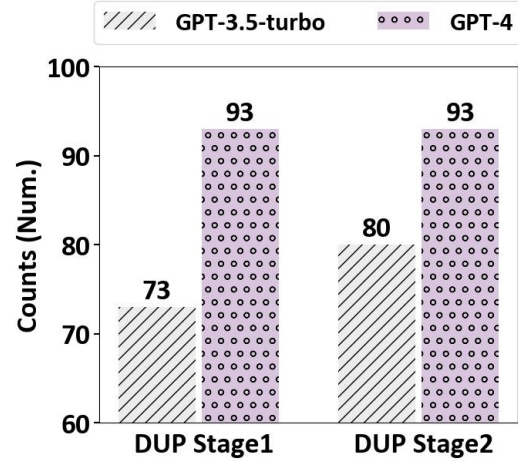


Figure 10: Compare the accuracy of DUP prompting stage1 and stage 2 information extracted by GPT-3.5-Turbo and GPT-4 respectively on GSM8K.

Quantitatively analyze more reasoning datasets

As shown in the Figure 11, understanding errors still account for the largest proportion of all error categories. Specifically, understanding errors accounts for the largest proportion of error types, and our method also focuses on resolving these errors. However, the reduction in calculation errors and process errors should be attributed to a deep understanding of the problem. For the CSQA and Coin Flip datasets, which respectively belong to Commonsense reasoning and Symbolic reasoning tasks, this leads to the absence of calculation errors and seldom process errors. We can see that our approach benefits more from reducing understanding errors.

A.3 Example Outputs by Zero-shot DUP

Tables 11 to 20 list example outputs generated by Zero-shot DUP for each dataset.

Error Analysis Prompt using GPT-4-1106:Question: *[Question Descripte]*.Wrong Response: *[Wrong Answer]*.Correct Response: *[Correct Answer]*.Please judge which type of error it belongs to based on the above information *[question, wrong answer, correct answer]*:

1. Errors in understanding: e.g., Errors in semantic understanding and lack of commonsense concepts

2. Calculation error: An error occurred while performing a basic operation

3. Process errors: e.g., missing step, hallucination

Finally, please explain why this error falls into the category you selecte

Figure 8: **Prompt for error analysis.** Question descripte is origin problem input. wrong answers represent incorrect text generated by the LLMs. Correct answer indicates the correct label for the question.

No.	Templates	Reasoning tasks
1	Extract core question: Please extract core question, only the most comprehensive and detailed's one! Extract problem-solving information : Please extract the most useful information related to the core question(<i>[Core Question]</i>), Only extract the most useful information, and list them one by one! Generate the answer: Hint: <i>[Problem-solving Info]</i> , \n <i>[Core Question]</i> . \n Please understand core question and problem-solving information, then solve thequestion step by step and show the answer.	GSM8K, AddSub, SVAMP, MultiArith, SingleEq, AQuA, CSQA, StrategyQA, Coin Flip
2	Prompt: Please accurately understand the question useful information and solve the question step by step.	Last Letter

Table 7: Reasoning prompt templates include all reasoning tasks. *Red* indicates extracted problem-solving information to the problem. *Blue* indicates extracted core question.

No.	Template	Arithmetic Reasoning
1	Here is a math question and a model's answer about this question. Please extract the EXACT number from the answer text as the final answer for question. QUESTION: {}. \nANSWER: {} Final format should be a legal 'number' without any suffix such as '\$'. The final answer is:	GSM8K, AddSub, SVAMP, MultiArith, SingleEq
2	Here is a math question and a model's answer about this question. Please extract the EXACT choice from the answer text as the final answer for question. QUESTION: {}. \nANSWER: {} Final format should be a legal 'options', If you can't find the right choice , just answer Z. The final answer is:	AQUA

Table 8: Extract answer template with GPT-3.5-turbo on Arithmetic Reasoning.

No.	Template	Commonsense Reasoning
1	<p>Here is a Commonsense question and a model’s answer about this question. Please extract the EXACT one choice from the answer text as the final answer for question.</p> <p>QUESTION: {}. \nANSWER: {}</p> <p>Final format should be a legal 'choice'(eg. (A) or (b)),If you can’t find the correct choice, just answer the one that is closest to the answer. The final answer is:</p>	CommonsenseQA
2	<p>Here is a Commonsense question and a model’s answer about this question. Please extract the EXACT one choice from the answer text as the final answer for question.</p> <p>QUESTION: {}. \nANSWER: {}</p> <p>Final format should be a legal 'string'(Yes or No), If you Uncertain or unknow, Please understand that the question and answer information outputs the closest answer,you can only output Yes or No. The final answer is:</p>	StrategyQA

Table 9: Extract answer template with GPT-3.5-turbo on Commonsense Reasoning.

No.	Template	Symbolic Reasoning
1	<p>Here is a question related to concatenating last letters, along with a model’s answer to this question. Please extract the EXACT string from the answer text to determine the final answer to the question.</p> <p>QUESTION: {}. \nANSWER: {}</p> <p>Note: Please extract the string after (the answer is:), or the string that appears at the end of the sentence. The final answer is:</p>	Last letters
2	<p>Here is a Coin flip related question and a model’s answer about this question. Please extract the EXACT string from the answer text as the final answer for question.</p> <p>QUESTION: {}. \nANSWER: {}</p> <p>Note: If you find that the coin is still heads up, the answer is "yes". If you find that the coin is not still heads up or tails up. the answer is "no". The final answer is:</p>	Coin Flip

Table 10: Extract answer template with GPT-3.5-turbo on Symbolic Reasoning.

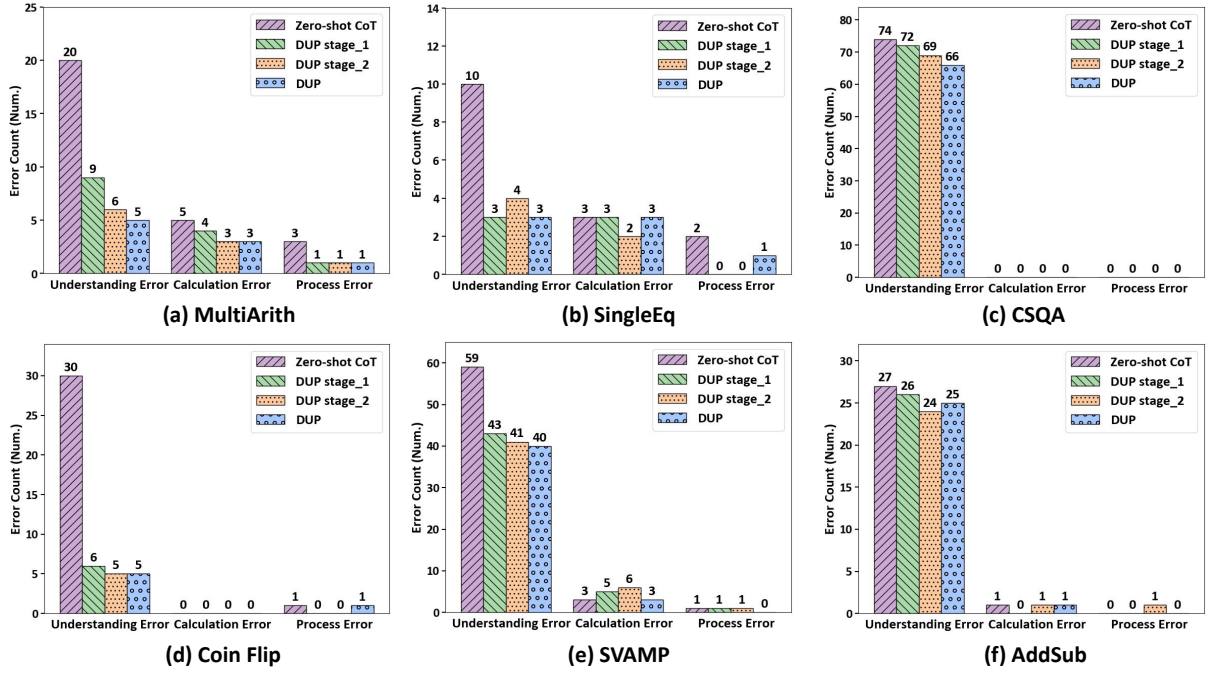


Figure 11: **Distribution for errors of more reasoning datasets.** Randomly select 300 examples from different reasoning task datasets (only 254 examples except AQuA), use GPT-3.5-Turbo LLM to generate responses and count failed answers to quantitatively analyze different error categories. In the figure, CSQA belongs to Commonsense reasoning, coin flip belongs to Symbolic reasoning, and the rest belong to Arithmetic reasoning.

Q: A number is as much greater than 36 as is less than 86. Find the Number. Answer Choices: A)60 B)56 C)51 D)61 E)41

Stage1: Extract the core question.

A: Please extract core question,only the most comprehensive and detailed one!

Output: A number is as much greater than 36 as is less than 86. Find the number.

Stage2: Extract problem-solving info.

A: Note: Please extract the most useful information related to the core question(**A number is as much greater than 36 as is less than 86. Find the number**), Only extract the most useful information , list them one by one!

Output: 1. The number is as much greater than 36 as it is less than 86.

Stage3: Generate the answers.

A: Hint: 1. The number is as much greater than 36 as it is less than 86.

A number is as much greater than 36 as is less than 86. Find the number.

Please understand the Hint and question information,then solve the question step by step and show the answer.

Output: The number that is as much greater than 36 as it is less than 86 can be found by calculating the average of 36 and 86. Average = $(36 + 86) / 2$ Average = 61. So, the number is 61, which is option D.

Table 11: Example outputs by Zero-shot DUP for AQuA.

Q: Morisette and Kael were asked to bring fruits. Morisette brought 5 apples and 8 oranges, while Kael brought twice the amount of apples and half the number of oranges than Morisette. How many fruits do they have in total?

Stage1: Extract the core question.

A: Please extract core question, only the most comprehensive and detailed's one!

Output: How many fruits do Morisette and Kael have in total?

Stage2: Extract problem-solving info.

A: Please extract the most useful information related to the core question (How many fruits do Morisette and Kael have in total?), Only extract the most useful information, list them one by one!

Output: -Morisette brought 5 apples and 8 oranges.\n- Kael brought twice the amount of apples and half the number of oranges than Morisette.

Stage3: Generate the answers.

A: -Morisette brought 5 apples and 8 oranges.\n- Kael brought twice the amount of apples and half the number of oranges than Morisette.\nHow many fruits do Morisette and Kael have in total? \nPlease understand the Hint and question information, then solve the question step by step and show the answer.

Output: Morisette brought 5 apples + 8 oranges = 13 fruits.\nKael brought $2 * 5$ apples = 10 apples.\nKael also brought $1/2 * 8$ oranges = 4 oranges. \nSo, Kael brought a total of 10 apples + 4 oranges = 14 fruits.\nTherefore, Morisette and Kael have 13 fruits + 14 fruits = 27 fruits in total.

Table 12: Example outputs by Zero-shot DUP for GSM8K.

Q: Each chocolate bar in a box cost \$3. If a box had 7 bars total and Olivia sold all but 4 bars, how much money would she have made?

Stage1: Extract the core question.

A: Please extract core question, only the most comprehensive and detailed's one!

Output: How much money would Olivia have made if she sold all but 4 bars from a box of 7 chocolate bars, where each bar costs \$3?

Stage2: Extract problem-solving info.

A: Please extract the most useful information related to the core question (How much money would Olivia have made if she sold all but 4 bars from a box of 7 chocolate bars, where each bar costs \$3?), Only extract the most useful information, list them one by one!

Output: 1. Each chocolate bar costs \$3. 2. There are 7 bars in a box. 3. Olivia sold all but 4 bars.

Stage3: Generate the answers.

A: Hint: 1. Each chocolate bar costs \$3. 2. There are 7 bars in a box. 3. Olivia sold all but 4 bars. How much money would Olivia have made if she sold all but 4 bars from a box of 7 chocolate bars, where each bar costs \$3?

Please understand the Hint and question information, then solve the question step by step and show the answer.

Output: Olivia sold all but 4 bars from a box of 7 chocolate bars, where each bar costs \$3. So, Olivia sold $7 - 4 = 3$ bars. The total amount of money Olivia would have made is $3 \text{ bars} * \$3 \text{ per bar} = \9 .

Table 13: Example outputs by Zero-shot DUP for MultiArith.

Q: Baker made 144 cakes. He sold 71 of them. Then he made 111 more cakes. How many more cakes did baker make than those he sold?

Stage1: Extract the core question.

A: Please extract core question, only the most comprehensive and detailed's one!

Output: How many more cakes did the baker make than those he sold?

Stage2: Extract problem-solving info.

A: Please extract the most useful information related to the core question (How many problems did Paige have to do for homework after she finished 44 of the problems at school, if she initially had 43 math problems and 12 science problems?), Only extract the most useful information, list them one by one!

Output: - Baker made 144 cakes.\n- He sold 71 of them.\n- He made 111 more cakes.

Stage3: Generate the answers.

A: Hint: - Baker made 144 cakes.\n- He sold 71 of them.\n- He made 111 more cakes. \nHow many more cakes did the baker make than those he sold? \nPlease understand the Hint and question information, then solve the question step by step and show the answer.

Output: The baker initially made 144 cakes and then made an additional 111 cakes, so in total, he made $144 + 111 = 255$ cakes. He sold 71 cakes. Therefore, the baker made $255 - 71 = 184$ more cakes than he sold.

Table 14: Example outputs by Zero-shot DUP for SVAMP.

Q: Craig walked 0.2 mile from school to David's house and 0.7 mile from David's house to his own house. How many miles did Craig walk in all?

Stage1: Extract the core question.

A: Please extract core question, only the most comprehensive and detailed's one!

Output: How many miles did Craig walk in all?

Stage2: Extract problem-solving info.

A: Please extract the most useful information related to the core question (How many miles did Craig walk in all?), Only extract the most useful information, list them one by one!

Output: - Craig walked 0.2 mile from school to David's house.\n- Craig walked 0.7 mile from David's house to his own house.

Stage3: Generate the answers.

A: Hint: - Craig walked 0.2 mile from school to David's house.\n- Craig walked 0.7 mile from David's house to his own house. \nHow many miles did Craig walk in all? \nPlease understand the Hint and question information, then solve the question step by step and show the answer.

Output: Craig walked a total of $0.2 \text{ mile} + 0.7 \text{ mile} = 0.9 \text{ mile}$.

Table 15: Example outputs by Zero-shot DUP for AddSub.

Q: Dana earns \$13 per hour. She worked 9 hours on Friday, 10 hours on Saturday, and 3 hours on Sunday. How much money did Dana earn in all?

Stage1: Extract the core question.

A: Please extract core question, only the most comprehensive and detailed's one!

Output: How much money did Dana earn in total if she earns \$13 per hour and worked 9 hours on Friday, 10 hours on Saturday, and 3 hours on Sunday?

Stage2: Extract problem-solving info.

A: Please extract the most useful information related to the core question (How much money did Dana earn in total if she earns \$13 per hour and worked 9 hours on Friday, 10 hours on Saturday, and 3 hours on Sunday?), Only extract the most useful information, list them one by one!

Output: 1. Dana earns \$13 per hour.\n2. She worked 9 hours on Friday.\n3. She worked 10 hours on Saturday.\n4. She worked 3 hours on Sunday.

Stage3: Generate the answers.

A: Hint: 1. Dana earns \$13 per hour.\n2. She worked 9 hours on Friday.\n3. She worked 10 hours on Saturday.\n4. She worked 3 hours on Sunday. \nHow much money did Dana earn in total if she earns \$13 per hour and worked 9 hours on Friday, 10 hours on Saturday, and 3 hours on Sunday? \nPlease understand the Hint and question information, then solve the question step by step and show the answer.

Output: Dana worked a total of 9 hours on Friday + 10 hours on Saturday + 3 hours on Sunday = 22 hours. Since Dana earns \$13 per hour, for the 22 hours she worked, she earned 22 hours * \$13/hour = 286.

Table 16: Example outputs by Zero-shot DUP for SingleEq.

Q: The artist was sitting quietly pondering, then suddenly he began to paint when what struck him?
Answer Choices: (A)sadness (B)anxiety (C)inspiration (D)discomfort (E)insights

Stage1: Extract the core question.

A: Please extract core question, only the most comprehensive and detailed's one!

Output: What struck the artist and prompted him to start painting?

Stage2: Extract problem-solving info.

A: Please extract the most useful information related to the core question (What struck the artist and prompted him to start painting?), Only extract the most useful information, list them one by one!

Output: - The artist was sitting quietly pondering\n- Suddenly, he began to paint\n- Something struck him and prompted him to start painting

Stage2: Generate the answers.

A: Hint: - The artist was sitting quietly pondering\n- Suddenly, he began to paint\n- Something struck him and prompted him to start painting \n3. She worked 10 hours on Saturday.\n4. She worked 3 hours on Sunday. \nWhat struck the artist and prompted him to start painting? \nPlease understand the Hint and question information, then solve the question step by step and show the answer.

Output: Based on the information provided, the answer is (C) inspiration. The artist was sitting quietly pondering, and suddenly something struck him, prompting him to start painting. This suggests that he was inspired by something, perhaps an idea, a vision, or a feeling, which motivated him to express himself through painting.

Table 17: Example outputs by Zero-shot DUP for CommonsenseQA.

Q: Can the original name of the zucchini be typed on the top row of a QWERTY keyboard?

Stage1: Extract the core question.

A: Please extract core question, only the most comprehensive and detailed's one!

Output: Can the original name of the zucchini be typed on the top row of a QWERTY keyboard?

Stage2: Extract problem-solving info.

A: Please extract the most useful information related to the core question (Can the original name of the zucchini be typed on the top row of a QWERTY keyboard?), Only extract the most useful information, list them one by one!

Output: 1. The original name of the zucchini is "Cucurbita pepo." \n2. The top row of a QWERTY keyboard contains the letters "QWERTYUIOP." \n3. The original name of the zucchini cannot be typed using only the top row of a QWERTY keyboard.

Stage3: Generate the answers.

A: Hint: 1. The original name of the zucchini is "Cucurbita pepo." \n2. The top row of a QWERTY keyboard contains the letters "QWERTYUIOP." \n3. The original name of the zucchini cannot be typed using only the top row of a QWERTY keyboard. \nCan the original name of the zucchini be typed on the top row of a QWERTY keyboard? \nPlease understand the Hint and question information, then solve the question step by step and show the answer.

Output: No, the original name of the zucchini "Cucurbita pepo" cannot be typed on the top row of a QWERTY keyboard. The top row of a QWERTY keyboard contains the letters "QWERTYUIOP" and none of these letters are present in the original name of the zucchini. Therefore, it is not possible to type the original name of the zucchini using only the top row of a QWERTY keyboard.

Table 18: Example outputs by Zero-shot DUP for StrategyQA.

Q: Take the last letters of each words in "Whitney Erika Tj Benito" and concatenate them.

A: Please accurately understand the question useful information and solve the question step by step.

Output: To solve the question, we need to take the last letters of each word in "Whitney Erika Tj Benito" and concatenate them. \nStep 1: Identify the words in the given phrase: \n Whitney \n Erika \n Tj \n Benito \nStep 2: Take the last letter of each word: \n Whitney: y \n Erika: a \n Tj: j \n Benito: o \nStep 3: Concatenate the last letters: \n y + a + j + o = yajo \nTherefore, the result of concatenating the last letters of each word in "Whitney Erika Tj Benito" is "yajo".

Q: Take the last letters of each words in "Silvia Carolina Stan Chuck" and concatenate them.

A: Please accurately understand the question useful information and solve the question step by step.

A: To solve this question step by step, we need to take the last letters of each word in the given phrase and concatenate them. \n1. "Silvia" - The last letter is "a". \n2. "Carolina" - The last letter is "a". \n3. "Stan" - The last letter is "n". \n4. "Chuck" - The last letter is "k". \nNow, we concatenate these last letters together: "aank". \nTherefore, the answer is "aank".

Table 19: Example outputs by Zero-shot DUP for Last Letters.

Q: A coin is heads up. Walter does not flip the coin. Mikayla does not flip the coin. Larry does not flip the coin. Ryan does not flip the coin. Is the coin still heads up? Note that "flip" here means "reverse".

Stage1: Extract the core question.

A: Please extract core question, only the most comprehensive and detailed's one!

Output: Is the coin still heads up if no one flips it?

Stage2: Extract problem-solving info.

A: Please extract the most useful information related to the core question (Is the coin still heads up if no one flips it?), Only extract the most useful information, list them one by one!

Output: 1. The initial position of the coin is heads up.\n2. Walter does not flip the coin.\n3. Mikayla does not flip the coin.\n4. Larry does not flip the coin.\n5. Ryan does not flip the coin.

Stage3: Generate the answers.

A: Hint: 1. The initial position of the coin is heads up.\n2. Walter does not flip the coin.\n3. Mikayla does not flip the coin.\n4. Larry does not flip the coin.\n5. Ryan does not flip the coin. \nIs the coin still heads up if no one flips it? \nPlease understand the Hint and question information, then solve the question step by step and show the answer.

Output: Based on the given information, we can conclude the following: \n1. The initial position of the coin is heads up.\n2. Walter does not flip the coin. Therefore, the coin remains heads up.\n3. Mikayla does not flip the coin. Therefore, the coin remains heads up.\n4. Larry does not flip the coin. Therefore, the coin remains heads up.\n5. Ryan does not flip the coin. Therefore, the coin remains heads up. Since no one flips the coin, the coin remains heads up.

Table 20: Example outputs by Zero-shot DUP for Coin Flip.