

# Pattern-Aware Chain-of-Thought Prompting in Large Language Models

Anonymous submission

## Abstract

Chain-of-thought (CoT) prompting can guide language models to engage in complex multi-step reasoning. The quality of provided demonstrations significantly impacts the success of downstream inference tasks. While existing automated methods prioritize accuracy and semantics in these demonstrations, we show that the underlying reasoning patterns play a more crucial role in such tasks. In this paper, we propose Pattern-Aware CoT, a prompting method that considers the diversity of demonstration patterns. By incorporating patterns such as step length and reasoning process within intermediate steps, PA-CoT effectively mitigates the issue of bias induced by demonstrations and enables better generalization to diverse scenarios. We conduct experiments on nine reasoning benchmark tasks using two open-source LLMs. The results show that our method substantially enhances reasoning performance and exhibits robustness to errors. The code will be made publicly available.

## 1 Introduction

Large language models (LLMs) have been proven highly effective in solving complex reasoning tasks. One technique contributing to their success is the chain-of-thought (CoT) prompting (Wei et al., 2022b), which motivates the LLMs to perform multi-step reasoning instead of providing direct answers. This approach can significantly enhance the model’s ability to handle challenging tasks such as arithmetic and symbolic questions.

Generally, the overall effectiveness of CoT relies on the quality of the demonstrations provided. When confronted with no examples but only the prompt “Let’s think step by step”, known as **Zero-Shot-CoT** (Kojima et al., 2022), LLMs struggle with reasoning and encounter hallucination-related issues. While manually designing demonstrations for each question can alleviate such problems (Wei et al., 2022b), it comes with a significant labour

### Question

Mary has 9 yellow marbles. Joan has 3 yellow marbles. How many yellow marbles do they have in all ?

### Rationale

We know that Mary has 9 yellow marbles, and Joan has 3 yellow marbles. So together, they have  $9 + 3 = 12$  yellow marbles. The answer is 12.



Figure 1: Example of the chain-of-thought reasoning process: This comprises a question accompanied by a rationale. The rationale serves as a depiction of how LLMs navigate the reasoning process to arrive at the answer to the given question.

cost. To address such challenges, Zhang et al. (2023) propose **Auto-CoT**, which can automatically construct demonstrations as prompts. It initially partitions questions from a given dataset into clusters and then selects a representative question from each cluster. The selected questions are answered using Zero-Shot-CoT to obtain their **rationales** (the intermediate reasoning chain). The performance of this automated method is comparable to that of Manual-CoT.

Despite the efficacy of the automated method, how to develop a sound and complete set of demonstrations remains an area for further exploration. Several studies advocate for incorporating external knowledge to ensure the accuracy of the intermediate reasoning chain (Zhao et al., 2023; Weng et al., 2023; Li et al., 2024). Others suggest generating multiple CoT paths, complemented by a verification process to maintain self-consistency (Wang et al., 2023b; Yao et al., 2023; Liu et al., 2023).

062 However, most prior research focuses on the preci- 112  
063 sion of demonstrations, with limited exploration of 113  
064 the distributional power inherent in these demon- 114  
065 strations. Enlightened by Min et al. (2022) and 115  
066 Madaan et al. (2023), LLMs perform CoT through 116  
067 a counterfactual approach: it does not necessitate 117  
068 precise example results but rather learns from the 118  
069 underlying **patterns** (e.g. equations, templates) 119  
070 exhibited by the examples. 120

071 In this paper, we introduce a novel approach 121  
072 called Pattern-Aware Chain-of-Thought (**PA-CoT**) 122  
073 and demonstrate that LLMs can achieve improved 123  
074 reasoning performance by embracing the diversity 124  
075 inherent in demonstration patterns. Following the 125  
076 Auto-CoT schema, we automatically generate ques- 126  
077 tion clusters and select representative questions 127  
078 from each cluster. However, instead of relying 128  
079 solely on question embeddings for clustering, we 129  
080 explore multiple methods to enrich the diversity 130  
081 of rationale patterns. We contend that the conven- 131  
082 tional embedding-based clustering focuses solely 132  
083 on question semantics, lacks reflection on the ra- 133  
084 tionale, and consequently fails to encompass the 134  
085 full spectrum of demonstrations, as shown in Fig- 135  
086 ure 1. To quantify the diversity of patterns, we 136  
087 introduce three metrics: (i) the length or steps of 137  
088 the rationale, where a shorter rationale implies a 138  
089 simpler solution, while a longer one indicates more 139  
090 complex reasoning requirements; (ii) the processes 140  
091 within the rationale, where distinct equations or 141  
092 logics represent different solving approaches; and 142  
093 (iii) a combination of rationale steps and processes, 143  
094 providing a comprehensive perspective that consid- 144  
095 ers both aspects simultaneously. 145

096 We evaluate the performance of PA-CoT across 146  
097 six arithmetic and three non-arithmetic reasoning 147  
098 tasks. The experimental results consistently demon- 148  
099 strate that the combination strategy outperforms 149  
100 other methods across two LLMs. This suggests 150  
101 that LLMs derive substantial benefits from the di- 151  
102 verse patterns presented in demonstrations. Further 152  
103 experiments are conducted to examine the impact 153  
104 of rationale step and process aspects. We empiri- 154  
105 cally find that PA-CoT introduces less bias to the 155  
106 generated answer and exhibits error robustness, at- 156  
107 tributed to our strategy emphasizing diversity. 157

## 108 2 Related Work 158

109 This section reviews how chain-of-thought (CoT) 159  
110 prompting works and introduces various advanced 160  
111 approaches.

## 2.1 Chain-of-Thought Prompting 112

113 Large language models have demonstrated signifi- 114  
115 cant ability in comprehending context and respond- 116  
117 ing to prompts (Brown et al., 2020; Ouyang et al., 118  
119 2022). Recent studies highlight that LLMs can 120  
121 achieve improved task completion without fine- 122  
123 tuning, particularly on reasoning tasks, when pro- 124  
125 vided with few-shot demonstrations (Wei et al., 126  
127 2022b). For instance, when presented with an ex- 128  
129 ample like *Q: Mary has 9 yellow marbles. John* 130  
131 *has 3 yellow marbles. How many yellow marbles* 131  
132 *do they have in all? A: They have  $9 + 3 = 12$  yel-* 132  
133 *low marbles. The answer is 12,* LLMs are expected 133  
134 to emulate such a format, deconstruct the ques- 134  
135 tion, engage in multi-step reasoning, and refrain 136  
137 from generating random answers in subsequent 137  
138 tasks. This process is commonly referred to as 138  
139 chain-of-thought prompting or in-context learning 139  
140 (Wei et al., 2022a; Xie et al., 2022). However, im- 140  
141 plementing this practice often involves the manual 141  
142 design of prompts at a labour cost. Consequently, 142  
143 researchers are exploring more efficient example 143  
144 selection strategies to streamline this process. 144

## 2.2 Example Selection and Refinement 135

136 Several CoT studies are directed towards automat- 136  
137 ing the generation of demonstrations, such as 137  
138 retrieval-based (Rubin et al., 2022), zero-shot (Ko- 138  
139 jima et al., 2022), clustering-based (Zhang et al., 139  
140 2023), and self-prompt (Shao et al., 2023). How- 140  
141 ever, many of these approaches encounter chal- 141  
142 lenges in achieving performance comparable to 142  
143 Manual-CoT, primarily due to the absence of super- 143  
144 vision in example selection. In another branch of re- 144  
145 search, efforts are focused on enhancing the quality 145  
146 of CoT demonstrations. They incorporate elements 146  
147 such as knowledge-infusion (Zhao et al., 2023; 147  
148 Weng et al., 2023; Li et al., 2024), self-consistency 148  
149 (Wang et al., 2023b), complexity-based (Fu et al., 149  
150 2022), contrastive-based (Chia et al., 2023), and 150  
151 progressive-hint (Zheng et al., 2023). The primary 151  
152 goal of these strategies is to ensure that LLMs ad- 152  
153 here to the correct prompt and avoid being misled. 153

## 2.3 Role of Example Pattern 154

155 To understand the underlying mechanism of CoT, 155  
156 Min et al. (2022) and Madaan et al. (2023) employ 156  
157 counterfactual prompting methods. These methods 157  
158 involve substituting question-answer mapping, to- 158  
159 ken distributions, answer patterns, and many other 159  
160 factors. Their findings consistently show that the 160

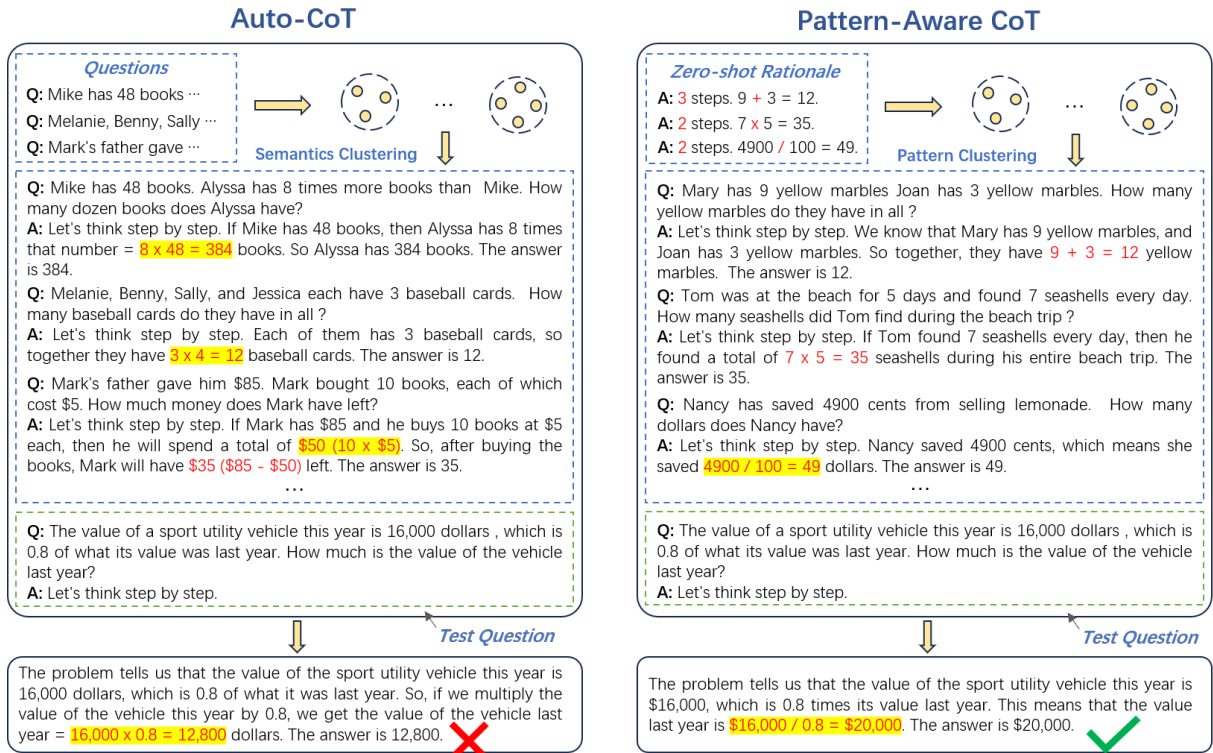


Figure 2: Example of Auto-CoT and PA-CoT. The upper part comprises selected demonstrations and a test question, and the lower part displays the corresponding answer generated by the same LLM.

161 correctness of examples is not the most crucial  
 162 factor, but rather the distribution or pattern (e.g.  
 163 equations, templates, sentence structure) of the ex-  
 164 amples. In this paper, we continue to uncover the  
 165 power of CoT patterns and show how they can im-  
 166 prove the reasoning process.

### 167 3 Pattern-Aware Chain-of-Thought

168 We now explore the impact of diverse demonstra-  
 169 tion reasoning patterns on chain-of-thought prompt-  
 170 ing. According to Min et al. (2022), the precision  
 171 of demonstrations is not crucial when LLMs en-  
 172 gage in CoT. Even if all the demonstrations pro-  
 173 vided are incorrect, it would only marginally im-  
 174 pede performance. This aligns with the insight de-  
 175 rived from Auto-CoT (Zhang et al., 2023): cluster-  
 176 ing zero-shot question-answer pairs (Kojima et al.,  
 177 2022) without emphasizing accuracy can still yield  
 178 valuable examples. Consequently, our focus shifts  
 179 to a more nuanced factor - the underlying reason-  
 180 ing pattern that harbours more informative content  
 181 (Madaan et al., 2023) - to evaluate its potential  
 182 benefits for the CoT process.

183 We argue that demonstrations function as tem-  
 184 plates, and they provide accessible reasoning for-  
 185 mats for LLMs to emulate. The homogeneity in

186 demonstrations poses a risk of introducing bias  
 187 into the generated answers (Wang et al., 2023a).  
 188 Conversely, maintaining diverse demonstrations  
 189 enables a broader exploration of new reasoning  
 190 inferences. Although Auto-CoT claims to cluster  
 191 based on diversity, it predominantly clusters by  
 192 question semantics, providing limited assistance  
 193 in problem-solving. In light of this, we propose  
 194 **Pattern-Aware Chain-of-Thought (PA-CoT)** that  
 195 refines the example selection process to enhance  
 196 the variety of reasoning chains. This approach en-  
 197 sures that selected examples contribute to a broader  
 198 range of cases, fostering more generalizable out-  
 199 comes.

200 In particular, we choose to experiment with arith-  
 201 metic and symbolic problems since the process pat-  
 202 terns are relatively intuitive. Given a dataset, each  
 203 question is first answered by adding the phrase  
 204 “Let’s think step by step” (zero-shot). Then we se-  
 205 lect  $k$  questions along with their rationales to serve  
 206 as a general demonstration prompt for the entire  
 207 dataset (Wei et al., 2022b; Zhang et al., 2023). We  
 208 design a rationale-based demonstration selection  
 209 method followed by three simple yet efficient vari-  
 210 ants to form our testbed:

- **Cluster based on rationale semantics.** This

approach involves a straightforward shift from question embeddings to rationale embeddings. The goal is to determine if the underlying pattern can be discovered through this minor alteration. However, our experiment indicates that this method can still be distracted from irrelevant elements such as characters or scenes, hindering its ability to generate diverse demonstrations.

- **Cluster based on rationale step length.** This approach is inspired by the notion of reasoning complexity (Fu et al., 2022; Zhou et al., 2022), where a simple task typically involves a few steps, and a complex task requires longer reasoning chains. Our aim is for the demonstrations to encompass both aspects simultaneously. For instance, if all demos are complex, the test question may involve an unnecessarily lengthy reasoning process, and vice versa. To validate this hypothesis, we include two comparative studies in our experiment.
- **Cluster based on rationale reasoning process.** This approach is designed to extract patterns that guide the task towards reaching its objectives (Madaan et al., 2023). Empirically, we choose mathematical symbols for arithmetic tasks and keywords for symbolic ones. For more details, please see Appendix A. In these problems, a process can effectively represent a solution for a particular question type. For example, an equation like  $2 + 3 = 5$  can evoke the association of addition, but it provides little assistance in understanding multiplication. Our findings demonstrate that diverse process patterns can significantly mitigate bias in the rationale, as illustrated in Figure 2.
- **Combination of step length and process.** Given that the previously mentioned methods focus on distinct dimensions of rationale patterns, this approach seeks to integrate them, offering a comprehensive perspective. As semantics may introduce irrelevant distractions, it is not considered in this method. There are various ways to combine the step length and the process, and we opt for the straightforward concatenation of the two dimensions. We also test additional variants in subsequent experiments.

In summary, we adopt the aforementioned methods as our demonstration clustering strategy. We explicitly extract patterns for each question-rationale pair and encode them into vector representations using Sentence-BERT<sup>1</sup> (Reimers and Gurevych, 2019). For instance, we encode “3” if the step length is 3 (split by “. ” or “\n”), encode “+” if the process appears in the rationale (concatenate if there are multiple processes), and encode “3+” for our combination strategy. These representations undergo processing by the  $k$ -means clustering algorithm, similar to Auto-CoT. Within each cluster, we sort the distances and select the example closest to the centre. It is important to note that Wei et al. (2022b) and Zhang et al. (2023) both impose restrictions on the chosen example, requiring it to be simple (question less than 60 tokens and rationale less than 5 steps). In contrast, we do not impose such restrictions to preserve variety. The  $k$  selected question-rationale pairs are then assembled as the final prompt for inference.

## 4 Experiments

In this section, our objective is to evaluate the effectiveness of our proposed PA-CoT and assess whether the introduced variety yields benefits.

### 4.1 Experimental Setup

**Datasets.** We adopt nine representative datasets for our reasoning tasks: MultiArith (Roy and Roth, 2015), GSM8K (Cobbe et al., 2021), AddSub (Hosseini et al., 2014), AQUA-RAT (Ling et al., 2017), SingleEq (Koncel-Kedziorski et al., 2015), SVAMP (Patel et al., 2021), Coin-Flip (Wei et al., 2022b), BIG-bench Date Understanding, and BIG-bench Tracking Shuffled Objects (Srivastava et al., 2023). They require certain reasoning steps and are commonly used for CoT method comparisons (Wei et al., 2022b; Kojima et al., 2022; Zhang et al., 2023; Wang et al., 2023a; Fu et al., 2022).

**Language Models.** We consider open-source large language models as our inference engine. Specifically, we choose LLaMA-2-7b-chat-hf (Touvron et al., 2023) and qwen-7b-chat (Bai et al., 2023) models, as they have been reported to be comparable to GPT-3.5<sup>2</sup> in terms of arithmetic ability and possess chain-of-thought reasoning capabilities.

<sup>1</sup>We use all-MiniLM-L6-v2 as the embedding encoder. <https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2>

<sup>2</sup><https://platform.openai.com/docs/models>



	Model	MultiArith	GSM8K	AddSub	AQuA	SingleEq	SVAMP	Coin	Date	Tracking
LLaMA-2-7b-chat-hf	Zero-Shot-CoT	72.33	21.00	57.97	24.01	57.67	41.90	44.60	39.29	30.80
	Auto-CoT	76.00	<u>27.36</u>	58.48	24.01	64.96	43.80	51.20	44.71	<b>32.53</b>
	PA-CoT-semantic	74.83	26.76	63.29	24.80	66.92	47.19	48.00	43.08	31.66
	PA-CoT-step	76.16	24.41	<b>67.59</b>	<u>29.13</u>	66.14	47.59	48.00	44.44	<u>33.33</u>
	PA-CoT-process	<b>79.66</b>	25.39	65.06	25.19	<b>71.85</b>	<u>48.50</u>	<b>59.40</b>	<b>47.96</b>	32.26
	PA-CoT-concat	<u>76.67</u>	<b>28.05</b>	<u>66.83</u>	<b>29.92</b>	<u>71.06</u>	<b>50.10</b>	<u>58.40</u>	<u>46.07</u>	<b>32.53</b>
qwen-7b-chat	Zero-Shot-CoT	87.33	42.83	54.93	<b>35.03</b>	69.09	55.70	45.40	50.13	<u>32.40</u>
	Auto-CoT	<u>90.66</u>	47.01	62.53	30.31	80.31	60.19	45.40	48.78	29.73
	PA-CoT-semantic	<b>91.33</b>	44.80	65.06	31.88	78.74	59.00	43.20	52.38	31.00
	PA-CoT-step	90.33	46.85	<b>74.17</b>	33.07	78.14	<u>62.00</u>	38.00	49.32	30.46
	PA-CoT-process	90.50	<u>47.16</u>	67.59	29.52	<u>82.08</u>	61.50	<b>52.60</b>	<b>55.72</b>	<b>32.53</b>
	PA-CoT-concat	<b>91.33</b>	<b>48.14</b>	<u>72.40</u>	<u>33.46</u>	<b>83.85</b>	<b>62.30</b>	<u>47.40</u>	<u>53.13</u>	31.60

Table 1: Accuracy (%) on nine reasoning datasets. We present the mean value obtained from five runs.

ities. These LLMs are deployed on our local server equipped with 4x NVIDIA GeForce RTX 3090. We use the inference function of these models and the process does not involve training or finetuning. We set the hyperparameter temperature as 0.4 to regulate the model’s randomness (Xu et al., 2022).

It is noteworthy that, as highlighted by Wei et al. (2023), larger models are more susceptible to the influence of examples. We observe that these 7B models can also be impacted. Thus, PA-CoT is expected to be effective in enhancing their performance.

**Baselines.** We primarily compare our methods with Zero-Shot-CoT (Kojima et al., 2022) and Auto-CoT (Zhang et al., 2023). To clarify the different variations of our proposed PA-CoT method, we note each pattern at the end of its name. For example, PA-CoT-semantic for clustering based on rationale semantics, and similarly for PA-CoT-step, PA-CoT-process, and PA-CoT-concat.

## 4.2 Main Results

Table 1 displays the overall performance of various methods on two LLMs. Since our primary focus is on evaluating the effectiveness of PA-CoT, we are not concerned with determining which LLM outperforms the other. Based on the results, we have the following observations:

- Auto-CoT consistently outperforms Zero-Shot-CoT, indicating that the cluster-sample strategy is effective across different LLMs. With the guidance of demonstrations, LLMs exhibit an enhanced capability to generate improved results.
- Simply switching from question embeddings (Auto-CoT) to rationale embeddings (PA-CoT-

semantic) does not yield significant benefits, as they generally perform at a similar level. We attribute this phenomenon to the inherent similarity between the two embeddings. As the embedding encoder considers the entire sentence as input, it unavoidably incorporates numerous irrelevant elements, such as characters and scenes. Consequently, this approach does not effectively address the fundamental problem.

- Considering naive rationale patterns (PA-CoT-step and PA-CoT-process) can notably enhance performance in various scenarios, with some instances even ranking first among all methods. This observation suggests that by incorporating diverse patterns into demonstrations, LLMs can effectively learn from this variability and generalize better across the entire dataset. However, given the inherent characteristics of different datasets, a single pattern may not universally adapt to every case, leading to occasional failures.
- Concatenating step length and process patterns (PA-CoT-concat) consistently produces the most favourable results across various scenarios compared to alternative methods. This finding implies that LLMs derive substantial benefits from incorporating multiple dimensions in the demonstration. The inclusion of both step length and process patterns encompasses a broader spectrum of the data distribution. Consequently, they are less prone to sampling similar examples, contributing to improved overall performance.

In summary, we present different approaches and evaluate their performance on various reason-

Model		MultiArith	GSM8K	AddSub	AQuA	SingleEq	SVAMP	Coin	Date	Tracking
qwen-7b-chat	PA-CoT-step	90.33	46.85	74.17	33.07	78.14	62.00	38.00	49.32	30.46
	CoT-simple	84.50	43.82	70.37	27.55	80.31	62.00	47.59	52.74	32.73
	CoT-complex	81.50	41.16	74.43	OOM	78.14	59.40	38.20	39.83	31.13

Table 2: Comparison between methods with various demonstration lengths.

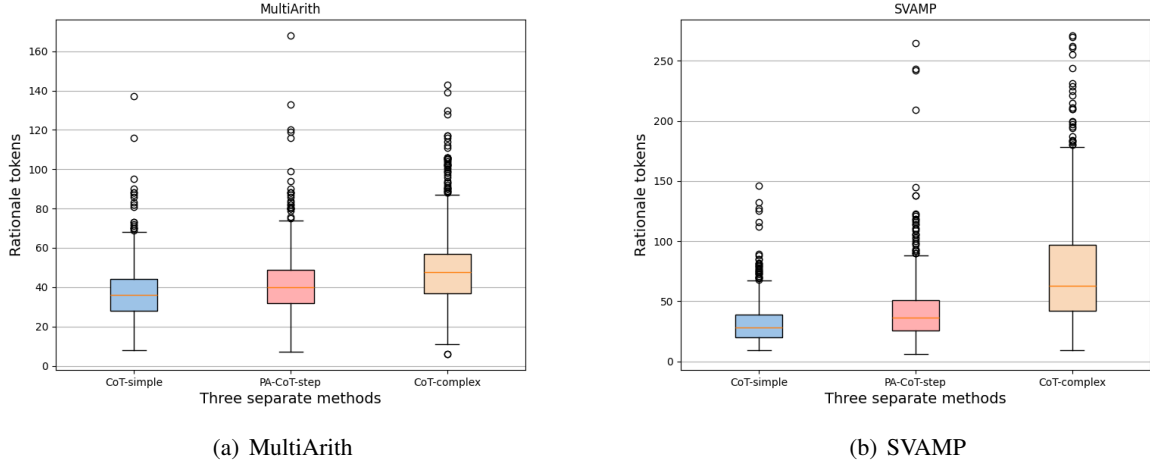


Figure 3: The box plot of generated rationale length across CoT-simple (pink), PA-CoT-step (blue), CoT-complex (green). The x-axis represents method names, and the y-axis represents the number of sentence tokens. The box in the middle represents where half of the numbers are. Extending from the box are whiskers that reach out to the minimum and maximum values within a specific range. Circles denote outliers, and the line splitting the box represents the median.

ing tasks. The results indicate the significance of demonstration patterns.

### 4.3 Impact of Step Length

To explore the influence of step length on LLMs' inference, we conduct additional experiments on this factor. In particular, we introduce two comparison methods: CoT-simple and CoT-complex. CoT-simple involves selecting examples with the fewest rationale steps, while CoT-complex involves selecting examples with the most (Fu et al., 2022). We aim to assess whether our PA-CoT-step method outperforms these two comparison methods.

Table 2 illustrates the performance of PA-CoT-step alongside two comparison methods. Overall, PA-CoT-step demonstrates advantages over the other two methods in most scenarios. We observe that CoT-complex tends to generate more errors during long intermediate steps and faces an out-of-memory (OOM) issue when the input becomes excessively long. While CoT-simple yields decent results in specific cases, it struggles with tasks requiring intricate reasoning.

We further visualize the distribution of gener-

ated answer length as in Figure 3. The box in the middle represents the interquartile range (IQR) and encapsulates the middle 50% of the data, with its lower and upper boundaries marked by the first quartile (Q1) and third quartile (Q3), respectively (Williamson et al., 1989; Kampstra, 2008). Inside the box, a line denotes the median (Q2) and indicates the dataset's central tendency. Extending from the box are whiskers that reach out to the minimum and maximum values within a specific range. Individual points beyond the whiskers signify potential outliers in the dataset.

The plot illustrates the correlation between the length of demonstrations and the number of generated tokens. With predominantly short demonstrations, CoT-simple tends to produce concise answers, resulting in a lower average value. Conversely, CoT-complex encourages longer answers, with most taking an extended route to complete the task. PA-CoT-step, on the other hand, maintains a moderate average rationale length. It covers a wider range from simple to complex reasoning. This adaptability allows it to perform well in more general situations.

Model		MultiArith	GSM8K	AddSub	AQuA	SingleEq	SVAMP	Coin	Date	Tracking
LLaMA-2-7b-chat-hf	PA-CoT-concat	76.67	28.05	66.83	29.92	71.06	50.10	58.40	46.07	32.53
	PA-CoT-sep	76.16	26.09	66.58	25.19	68.91	49.70	59.40	47.96	32.26
	PA-CoT-mean	75.83	27.67	68.86	24.01	70.86	48.19	54.80	41.73	31.85

Table 3: Comparison between methods with various combination strategies.

#### 4.4 Impact of Reasoning Process

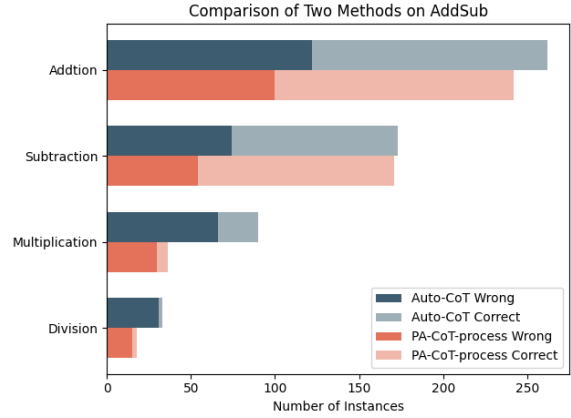
To investigate the role of process patterns in demonstrations, we also perform additional experiments on this aspect. Specifically, we categorize answers from Auto-CoT and PA-CoT-process based on basic arithmetic symbols: Addition, Subtraction, Multiplication, and Division. We then tally the number of correct and incorrect instances within each group. Figure 4 presents a comparison of the results on datasets AddSub and SingleEq, where the tasks are relatively straightforward.

Our observations reveal that Auto-CoT produces more incorrect arithmetic equations, leading to a higher error rate within each symbol group. This indicates a higher likelihood of being misled by the demonstrations. For instance, as depicted in Figure 2, the selected demos for Auto-CoT exhibit an overemphasis on multiplication. This trend is reflected in the results of Figure 4, where Auto-CoT generates instances solved using multiplication even when it is not appropriate. In contrast, PA-CoT-process exhibits a better ability to select the correct solving approach, resulting in fewer errors within each group.

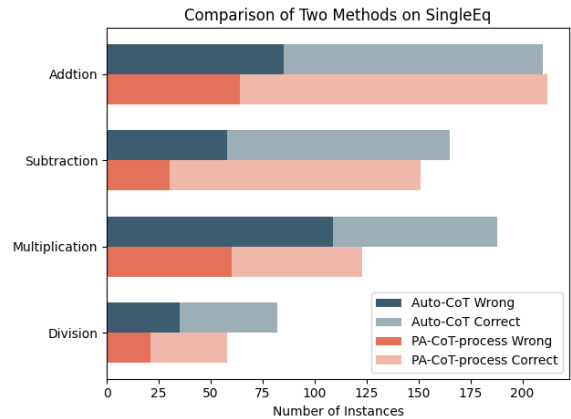
#### 4.5 Combination Strategy

The preceding sections showcase the impact of different pattern aspects. We now turn our attention to exploring the optimal way to combine them. We initially devise PA-CoT-concat to encode the concatenation of step length and process strings. Considering the potential limitations of this approach, we introduce two alternative methods to explore potential improvements. The first approach involves concatenating separate vector representations encoded from step length and process strings, denoted as PA-CoT-sep. The second approach employs mean pooling over the separate vector representations, denoted as PA-CoT-mean. All other settings remain constant as we conduct experiments on LLaMA-2-7b-chat-hf.

Table 3 presents the comparison results of these combination strategies. Overall, the performance of PA-CoT-concat slightly exceeds that of PA-CoT-sep and PA-CoT-mean. We attribute this outcome



(a) AddSub



(b) SingleEq

Figure 4: The distribution of the number of correct and wrong instances regarding different arithmetic symbols.

to the different practices of semantics encoding. PA-CoT-concat takes the entire pattern string as input, where the encoded vector reflects an integration of information. In contrast, the other two approaches separate the two patterns into distinct vectors, which creates a gap between their distributions.

In conclusion, our exploration of PA-CoT and its combination strategies sheds light on the importance of considering diverse demonstration patterns in enhancing language models' reasoning capabilities. Despite slight variations in performance among the approaches, our findings underscore the




Figure 5: Visualization of clustering on six reasoning tasks. Cluster centres are noted as stars. The scatter of PA-CoT-concat clusters shows its superiority in example differentiation.

Dataset	Demos	Incorrect	Error Rate
MultiArith	8	2	25.0%
GSM8K	8	5	62.5%
AddSub	8	3	37.5%
AQuA	4	4	100%
SingleEq	8	2	25.0%
SVAMP	8	3	37.5%
Coin	8	4	50.0%
Date	8	3	37.5%
Tracking	8	4	50.0%

Table 4: The number of demonstrations and their error rate for each dataset.

significance of integrating multiple pattern aspects for improved reasoning outcomes.

#### 4.6 Error Robustness

It is noteworthy that we do not enforce accuracy constraints on demonstrations. We proceed to count the incorrect instances within our selected demonstrations, as illustrated in Table 4.

It is intriguing to notice that the majority of our provided prompts are imperfect, with AQuA even exhibiting a 100% error rate. This phenomenon suggests that LLMs struggle to discern incorrect examples from correct ones. Instead, they learn from how the example approaches problem-solving, which we refer to as “pattern”. PA-CoT

encourages LLMs to follow the most probable reasoning chain towards the final answer and thus provides a significant improvement.

#### 4.7 Visualization

Figure 5 visualizes the  $k$  clusters of PA-CoT-concat on six reasoning tasks through PCA projection. The plot depicts that there is an apparent divergence between each cluster. The scatter implies that the step length and the process can effectively differentiate the patterns. With such diversities, LLMs can more effectively learn from demonstrations to generalize reasoning scenarios.

### 5 Conclusion

This paper introduces a novel pattern-aware chain-of-thought prompting method, which significantly enhances the reasoning performance of language models. Our experiments reveal that incorporating a variety of rationale step lengths prevents LLMs from taking excessively long or short steps, thereby maintaining a balanced inference chain. Similarly, diverse process patterns instruct LLMs to select appropriate reasoning routes and reduce bias from singular patterns. We also introduce a combination strategy that considers both aspects simultaneously. Further investigations show the effectiveness of our proposed strategy. Apart from performance gains, our method offers additional advantages such as ease of use and error robustness.



## 523 Limitations

524 Due to the shutdown of OpenAI code-davinci-002  
525 and text-davinci-002 API, we are unable to perform  
526 experiments on their models. Since most previous  
527 works choose to experiment on these models, we  
528 seek alternative LLMs as our inference engine. The  
529 two LLMs used in this paper are open-source, CoT-  
530 capable, and comparable to code-davinci-002. We  
531 hope such a practice can help future researches.

532 Another limitation is that our method has only  
533 been tested on datasets with explicit reasoning  
534 paths, such as arithmetic and symbolic tasks, where  
535 patterns are intuitive and easily extractable. When  
536 applied to datasets with implicit reasoning paths, it  
537 may be necessary to identify the inherent reason-  
538 ing processes. For more discussions, please see  
539 Appendix A.

## 540 References

541 Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang,  
542 Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei  
543 Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin,  
544 Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu,  
545 Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren,  
546 Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong  
547 Tu, Peng Wang, Shijie Wang, Wei Wang, Sheng-  
548 guang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang,  
549 Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu,  
550 Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingx-  
551 uan Zhang, Yichang Zhang, Zhenru Zhang, Chang  
552 Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang  
553 Zhu. 2023. [Qwen technical report](#). *arXiv preprint*  
554 *arXiv:2309.16609*.

555 Tom Brown, Benjamin Mann, Nick Ryder, Melanie  
556 Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind  
557 Neelakantan, Pranav Shyam, Girish Sastry, Amanda  
558 Askell, et al. 2020. [Language models are few-shot](#)  
559 [learners](#). *Advances in Neural Information Processing*  
560 *Systems*, 33:1877–1901.

561 Yew Ken Chia, Guizhen Chen, Luu Anh Tuan,  
562 Soujanya Poria, and Lidong Bing. 2023. [Con-](#)  
563 [trastive chain-of-thought prompting](#). *arXiv preprint*  
564 *arXiv:2311.09277*.

565 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian,  
566 Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias  
567 Plappert, Jerry Tworek, Jacob Hilton, Reiichiro  
568 Nakano, Christopher Hesse, and John Schulman.  
569 2021. [Training verifiers to solve math word prob-](#)  
570 [lems](#). *arXiv preprint arXiv:2110.14168*.

571 Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark,  
572 and Tushar Khot. 2022. [Complexity-based prompt-](#)  
573 [ing for multi-step reasoning](#). *arXiv preprint*  
574 *arXiv:2210.00720*.

Mohammad Javad Hosseini, Hannaneh Hajishirzi, Oren  
Etzioni, and Nate Kushman. 2014. [Learning to solve](#)  
[arithmetic word problems with verb categorization](#).  
In *Proceedings of the 2014 Conference on Empirical*  
*Methods in Natural Language Processing (EMNLP)*,  
pages 523–533, Doha, Qatar. Association for Com-  
putational Linguistics.

Peter Kampstra. 2008. [Beanplot: A boxplot alternative](#)  
[for visual comparison of distributions](#). *Journal of*  
*statistical software*, 28:1–9.

Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yu-  
taka Matsuo, and Yusuke Iwasawa. 2022. [Large lan-](#)  
[guage models are zero-shot reasoners](#). *Advances in*  
*Neural Information Processing Systems*, 35:22199–  
22213.

Rik Koncel-Kedziorski, Hannaneh Hajishirzi, Ashish  
Sabharwal, Oren Etzioni, and Siena Dumas Ang.  
2015. [Parsing algebraic word problems into equa-](#)  
[tions](#). *Transactions of the Association for Computa-*  
*tional Linguistics*, 3:585–597.

Xingxuan Li, Ruochen Zhao, Yew Ken Chia, Bosheng  
Ding, Shafiq Joty, Soujanya Poria, and Lidong Bing.  
2024. [Chain-of-knowledge: Grounding large lan-](#)  
[guage models via dynamic knowledge adapting over](#)  
[heterogeneous sources](#). In *International Conference*  
*on Learning Representations ICLR 2024*.

Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blun-  
som. 2017. [Program induction by rationale genera-](#)  
[tion: Learning to solve and explain algebraic word](#)  
[problems](#). In *Proceedings of the 55th Annual Meet-*  
*ing of the Association for Computational Linguistics*  
*(ACL)*, pages 158–167, Vancouver, Canada. Associa-  
tion for Computational Linguistics.

Tengxiao Liu, Qipeng Guo, Yuqing Yang, Xiangkun  
Hu, Yue Zhang, Xipeng Qiu, and Zheng Zhang. 2023.  
[Plan, verify and switch: Integrated reasoning with](#)  
[diverse X-of-thoughts](#). In *Proceedings of the 2023*  
*Conference on Empirical Methods in Natural Lan-*  
*guage Processing*, pages 2807–2822, Singapore. As-  
sociation for Computational Linguistics.

Aman Madaan, Katherine Hermann, and Amir Yazdan-  
bakhsh. 2023. [What makes chain-of-thought prompt-](#)  
[ing effective? a counterfactual study](#). In *Findings*  
*of the Association for Computational Linguistics:*  
*EMNLP 2023*, pages 1448–1535.

Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe,  
Mike Lewis, Hannaneh Hajishirzi, and Luke Zettle-  
moyer. 2022. [Rethinking the role of demonstrations:](#)  
[What makes in-context learning work?](#) In *Proceed-*  
*ings of the 2022 Conference on Empirical Methods in*  
*Natural Language Processing*, pages 11048–11064,  
Abu Dhabi, United Arab Emirates. Association for  
Computational Linguistics.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida,  
Carroll Wainwright, Pamela Mishkin, Chong Zhang,  
Sandhini Agarwal, Katarina Slama, Alex Ray, et al.

631	2022. <a href="#">Training language models to follow instructions with human feedback</a> . <i>Advances in Neural Information Processing Systems</i> , 35:27730–27744.	
632		
633		
634	Arkil Patel, Satwik Bhattamishra, and Navin Goyal.	
635	2021. <a href="#">Are NLP models really able to solve simple math word problems?</a> In <i>Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 2080–2094, Online.	
636		
637		
638		
639		
640	Association for Computational Linguistics.	
641	Nils Reimers and Iryna Gurevych. 2019. <a href="#">Sentence-bert: Sentence embeddings using siamese bert-networks</a> .	
642	In <i>Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> .	
643		
644	Association for Computational Linguistics.	
645		
646	Subhro Roy and Dan Roth. 2015. <a href="#">Solving general arithmetic word problems</a> . In <i>Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 1743–1752, Lisbon, Portugal.	
647		
648		
649		
650	Association for Computational Linguistics.	
651		
652	Ohad Rubin, Jonathan Herzig, and Jonathan Berant.	
653	2022. <a href="#">Learning to retrieve prompts for in-context learning</a> . In <i>Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 2655–2671, Seattle, United States.	
654		
655		
656		
657		
658	Association for Computational Linguistics.	
659	Zhihong Shao, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, and Weizhu Chen. 2023. <a href="#">Synthetic prompting: Generating chain-of-thought demonstrations for large language models</a> . <i>arXiv preprint arXiv:2302.00618</i> .	
660		
661		
662		
663		
664	Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al. 2023. <a href="#">Beyond the imitation game: Quantifying and extrapolating the capabilities of language models</a> . <i>Transactions on Machine Learning Research</i> .	
665		
666		
667		
668		
669		
670		
671	Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu,	
672		
673		
674		
675		
676		
677		
678		
679		
680		
681		
682		
683		
684		
685		
686		
687		
688		
	Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. <a href="#">Llama 2: Open foundation and fine-tuned chat models</a> .	689
		690
		691
		692
		693
	Xinyi Wang, Wanrong Zhu, Michael Saxon, Mark Steyvers, and William Yang Wang. 2023a. <a href="#">Large language models are latent variable models: Explaining and finding good demonstrations for in-context learning</a> . In <i>Thirty-seventh Conference on Neural Information Processing Systems</i> .	694
		695
		696
		697
		698
		699
	Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023b. <a href="#">Self-consistency improves chain of thought reasoning in language models</a> . In <i>The Eleventh International Conference on Learning Representations ICLR 2023</i> .	700
		701
		702
		703
		704
		705
	Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022a. <a href="#">Emergent abilities of large language models</a> . <i>Transactions on Machine Learning Research</i> . Survey Certification.	706
		707
		708
		709
		710
		711
		712
		713
	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022b. <a href="#">Chain-of-thought prompting elicits reasoning in large language models</a> . <i>Advances in Neural Information Processing Systems</i> , 35:24824–24837.	714
		715
		716
		717
		718
	Jerry Wei, Jason Wei, Yi Tay, Dustin Tran, Albert Webson, Yifeng Lu, Xinyun Chen, Hanxiao Liu, Da Huang, Denny Zhou, et al. 2023. <a href="#">Larger language models do in-context learning differently</a> . <i>arXiv preprint arXiv:2303.03846</i> .	719
		720
		721
		722
		723
	Yixuan Weng, Minjun Zhu, Fei Xia, Bin Li, Shizhu He, Shengping Liu, Bin Sun, Kang Liu, and Jun Zhao. 2023. <a href="#">Large language models are better reasoners with self-verification</a> . In <i>Findings of the Association for Computational Linguistics: EMNLP 2023</i> , pages 2550–2575, Singapore. Association for Computational Linguistics.	724
		725
		726
		727
		728
		729
		730
	David F Williamson, Robert A Parker, and Juliette S Kendrick. 1989. <a href="#">The box plot: a simple visual method to interpret data</a> . <i>Annals of internal medicine</i> , 110(11):916–921.	731
		732
		733
		734
	Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. 2022. <a href="#">An explanation of in-context learning as implicit bayesian inference</a> . In <i>International Conference on Learning Representations ICLR 2022</i> .	735
		736
		737
		738
		739
	Frank F Xu, Uri Alon, Graham Neubig, and Vincent Josua Hellendoorn. 2022. <a href="#">A systematic evaluation of large language models of code</a> . In <i>Proceedings of the 6th ACM SIGPLAN International Symposium on Machine Programming</i> , pages 1–10.	740
		741
		742
		743
		744

Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. [Tree of thoughts: Deliberate problem solving with large language models](#).

Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2023. [Automatic chain of thought prompting in large language models](#). In *The Eleventh International Conference on Learning Representations ICLR 2023*.

Ruochen Zhao, Xingxuan Li, Shafiq Joty, Chengwei Qin, and Lidong Bing. 2023. [Verify-and-edit: A knowledge-enhanced chain-of-thought framework](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 5823–5840, Toronto, Canada. Association for Computational Linguistics.

Chuanyang Zheng, Zhengying Liu, Enze Xie, Zhenguo Li, and Yu Li. 2023. [Progressive-hint prompting improves reasoning in large language models](#). *arXiv preprint arXiv:2304.09797*.

Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, et al. 2022. [Least-to-most prompting enables complex reasoning in large language models](#). *arXiv preprint arXiv:2205.10625*.

## A Reasoning Process Patterns

The reasoning process pattern generally guides the task towards reaching its objectives and hence can manifest differently depending on the task (Madaan et al., 2023). Table 5 shows the process pattern identified in our experiments.

For arithmetic tasks, we utilize a standardized set of mathematical symbols. For symbolic tasks, we identify empirical keywords as patterns, which serve a similar function to symbols. For instance, in scenarios like the Coin Flip, consecutive patterns such as ‘head-head-head’ are different from ‘head-tail-tail’. PA-CoT incentivizes LLMs to learn from these reasoning pathways. Additionally, there are potential automated methods for pattern selection, such as chi-square testing and keyword extraction.

## B Heuristic Demonstration

We aim for a fair comparison with Auto-CoT, which utilizes embeddings for clustering. Our experiments are designed to showcase that while clustering question embeddings may not always identify the best examples, clustering pattern embeddings can achieve this to some extent. Acknowledging that embeddings might not be the optimal clustering approach, we introduce another variant

Dataset	Process Pattern
MultiArith GSM8K AddSub AQuA SingleEq SVAMP	+ , - , * , x , / , % , > , <
Coin Date Tracking	‘heads up’, ‘tails up’ ‘day’, ‘week’, ‘month’, ‘year’, ‘yesterday’, ‘tomorrow’ ‘trade’, ‘switch’, ‘exchange’, ‘swap’

Table 5: The process pattern identified for each dataset.

Model	MultiArith	GSM8K	AddSub	
LLaMA-2-7b-chat-hf	PA-CoT-step	76.16	24.41	67.59
	PA-CoT-heuristic	77.00	24.94	67.84
qwen-7b-chat	PA-CoT-step	90.33	46.85	74.17
	PA-CoT-heuristic	91.66	45.18	75.20

Table 6: Comparison between methods with and without embedding clustering.

of PA-CoT using heuristics. Specifically, we organize demonstrations into groups based on the number of reasoning steps (e.g., 1, 2, ..., k), selecting one random instance per group instead of creating embeddings and clusters. This approach allows us to evaluate the effectiveness of embeddings.

Table 6 shows the comparison between PA-CoT-step and PA-CoT-heuristic. We observed a marginal improvement with PA-CoT-heuristic. This suggests that using Sentence-BERT to encode information such as step length does not significantly hinder performance. However, there is potential to enhance performance by adjusting the method of demonstration selection.