

# 000 001 *SynAdapt*: LEARNING ADAPTIVE REASONING IN 002 LARGE LANGUAGE MODELS VIA SYNTHETIC CONTIN- 003 UOUS CHAIN-OF-THOUGHT 004 005

006 **Anonymous authors**  
007 Paper under double-blind review  
008  
009  
010

## 011 ABSTRACT 012

013 While Chain-of-Thought (CoT) reasoning improves model performance, it incurs  
014 significant time costs due to the generation of discrete CoT tokens (DCoT). Continuous  
015 CoT (CCoT) offers a more efficient alternative, but existing CCoT methods  
016 are hindered by indirect fine-tuning, limited alignment, or inconsistent targets. To  
017 overcome these limitations, we propose *SynAdapt*, an innovative efficient reasoning  
018 framework. Specifically, *SynAdapt* generates the synthetic CCoT to serve as a  
019 precise and effective alignment target for LLMs. This synthetic CCoT explicitly  
020 guides the LLM to learn CCoT and derive accurate answers directly. Furthermore,  
021 relying solely on CCoT is insufficient for solving hard questions. To address this,  
022 *SynAdapt* integrates a difficulty classifier that leverages both question context and  
023 CCoT to identify hard questions. CCoT can effectively help identify hard ques-  
024 tions after some brief reasoning. We then adaptively prompt the LLM to re-think  
025 these hard questions for improved performance. Extensive experimental results  
026 across various benchmarks from different difficulty levels strongly demonstrate the  
027 effectiveness of our method, achieving the best accuracy-efficiency trade-off.<sup>1</sup>  
028

## 029 1 INTRODUCTION 030

031 Chain-of-Thought (CoT) reasoning (Kojima et al., 2022; Wei et al., 2022; Zhou et al., 2022) has  
032 shown remarkable potential in enhancing the problem-solving capabilities of Large Language Models  
033 (LLMs) for complex tasks (Guo et al., 2025; Yang et al., 2025; OpenAI, 2025). By decomposing  
034 problems into sequential steps, CoT allows LLMs to derive correct answers step-by-step. However,  
035 a major drawback of CoT is its high computational cost due to the generation of numerous tokens,  
036 which leads to substantial time consumption (Yu et al., 2024; Yeo et al., 2025). While this cost is often  
037 acceptable in **accuracy-sensitive scenarios**, such as AI for Science (AI4S) (Lu et al., 2024) where  
038 accuracy is paramount, it becomes problematic in **efficiency-sensitive scenarios**. For instance, in  
039 embodied intelligence, real-time human-computer interaction necessitates highly efficient reasoning  
040 to ensure a satisfactory user experience (Li et al., 2024a). Consequently, a critical challenge emerges:  
041 how to reduce the length of generated CoT while preserving its effective reasoning capabilities.

042 Existing efficient reasoning approaches mainly involve fine-tuning or direct prompting LLMs to  
043 reduce the number of COT steps (Arora & Zanette, 2025; Munkhbat et al., 2025; Xu et al., 2025a).  
044 However, the remaining CoT steps still involve numerous discrete natural language tokens, which we  
045 refer to as **DCoT**. As noted by Li et al. (2024b) and Lin et al. (2024), most of these verbalized tokens  
046 are mainly for communication and carry unnecessary linguistic details that do not contribute to the  
047 core reasoning process. One promising approach is fine-tuning LLM to replace DCoT with a more  
048 compact and continuous CoT representation, known as **CCoT** (Pfau et al., 2024; Goyal et al., 2023).  
049 During reasoning, CCoT retains the hidden state of the LLM and skips generating the one-hot token  
050 ID, allowing it to store more information than just a single token (Zhu et al., 2025).

051 Nonetheless, fine-tuning LLM to learn CCoT reasoning effectively remains challenging. Coconut  
052 (Hao et al., 2024) gradually fine-tunes the LLM to replace DCoT with CCoT using a curriculum  
053 learning strategy (Deng et al., 2024). However, as shown in Figure 1, it lacks explicit alignment

<sup>1</sup>We have released all our code and dataset in the supplementary materials for better review.




Figure 1: Comparisons between our *SynAdapt* and the other CCoT-based baselines. These baselines either train CCoT indirectly, provide only single-position alignment, or apply full alignment with incoherent targets.

between DCoT and CCoT, which limits its ability to effectively learn from the original DCoT. CODI (Shen et al., 2025b) introduces explicit alignment between the last token hidden state of DCoT and the final hidden state of CCoT, but ignores alignment for other intermediate tokens. CompressCoT (Cheng & Van Durme, 2024) attempts to identify a subset of important tokens from DCoT, whose length matches CCoT, and aligns the full CCoT with the hidden states of these tokens. However, selecting only several isolated DCoT tokens leads to incoherence in the reasoning process. This leads to significant performance degradation in CCoT learning.

To overcome these limitations, we propose a novel efficient reasoning framework called *SynAdapt*, which helps LLM learn **Adaptive** reasoning through **Synthetic** CCoT. Our approach begins by generating a synthetic CCoT to serve as a comprehensive alignment target. Specifically, we initialize a random CCoT, fix the LLM, and iteratively optimize the random CCoT into a synthetic CCoT to guide the LLM towards correct answers. The synthetic CCoT thereby serves as a better alignment target than only using several isolated and incoherent tokens from the original DCoT. During fine-tuning, we apply the full alignment using the synthetic CCoT, as shown by Figure 1. This strategy helps LLM learn the full CCoT rather than only the last one. Notably, we fine-tune the LLM to iteratively refine a meaningless draft to obtain the CCoT, rather than generating CCoT autoregressively. This approach is more efficient (Jiang et al., 2025) and can boost the reasoning ability of LLM by iterative refinement (Saunshi et al., 2025; Yu et al., 2025).

Moreover, according to the information theory (Nalewajski, 2011), compressing DCoT into the dense CCoT inevitably leads to information loss and increases the complexity of solving hard questions (Koehn & Knowles, 2017). We provide an example in Figure 5 of the Appendix. To address this, we train a difficulty classifier that assesses question difficulty based on both the question itself and the CCoT. And then prompt the LLM to re-think hard questions using discrete CoT tokens for improved accuracy. While CCoT may not be sufficient to solve these hard questions, it can help the classifier effectively identify them. Some hard questions resemble simpler ones and can only be distinguished through the brief reasoning captured by CCoT. We also present an illustrative example in Figure 6 of the Appendix.

We evaluate our method across various benchmarks with different difficulty levels, including GSM8K, MATH500, AMC23, AIME24, and AIME25. By dynamically adjusting the ratio of re-think hard questions, our method demonstrates adaptability in both accuracy-sensitive and efficiency-sensitive scenarios. Comprehensive experimental results demonstrate that our method outperforms other baselines in both scenarios, achieving an optimal accuracy-efficiency trade-off. We further assess identification performance of our difficulty classifier, showing its superior performance compared to other baselines. In addition, we evaluate the generalization capacity of our method across broader domains, such as scientific QA and coding, as well as under different LLM backbones. The main contributions of this paper are as follows:

- We propose a novel efficient reasoning framework that generates synthetic CCoT, providing a better full alignment target to help LLMs learn CCoT more effectively.
- We introduce a difficulty classifier that more effectively distinguishes hard questions by considering both the question and the CCoT, enabling adaptive re-thinking for improved accuracy.
- Extensive experimental results strongly demonstrate the effectiveness of our framework, achieving the best accuracy-efficiency trade-off.

108  
109  

## 2 RELATED WORK

110  
111 In this section, we introduce the mainstream related work on efficient reasoning in the LLMs, which  
112 can be mainly categorized into three types: SFT-based methods, RL-based methods, Prompt-based  
113 methods, and CCoT-based methods.114 **SFT-based methods** either discard the CoT entirely or dynamically compress the CoT in the training  
115 data. And then they apply supervised fine-tuning (SFT) on these compressed data to help LLM learn  
116 to reduce generation length. While these methods are effective in shortening the generated output,  
117 they may ignore some crucial details of the original CoT during fine-tuning, leading to significant  
118 performance degradation (Yu et al., 2024; Ma et al., 2025b; Munkhbat et al., 2025; Xia et al., 2025;  
119 Kang et al., 2025). **RL-based methods** primarily design length penalties to prevent the model  
120 from generating excessively long CoT. While these methods can reduce reasoning length without  
121 sacrificing LLM performance, they require substantial resources for repeated data sampling to LLM  
122 training. Additionally, the reduction in length is limited and may not be suitable for efficiency-  
123 sensitive scenarios, where minimizing generation length is crucial (Arora & Zanette, 2025; Luo et al.,  
124 2025; Yeo et al., 2025; Aggarwal & Welleck, 2025; Shen et al., 2025a). **Prompt-based methods**  
125 explicitly add length constraint instructions in the prompt for guiding LLM to reduce generation  
126 length. Although these approaches are low-cost, their impact on length reduction is limited. LLMs  
127 still tend to generate long, redundant reasoning CoTs, especially for those hard questions (Renze &  
128 Guven, 2024; Xu et al., 2025a; Lee et al., 2025; Han et al., 2024).129 Instead of reasoning by numerous redundant tokens, **CCoT-based methods** aim to compress the  
130 reasoning steps by replacing the original discrete CoT (DCoT) with Continuous CoT (CCoT) in  
131 the latent space. However, these methods often suffer from significant performance drops. This is  
132 mainly because they either don't explicitly align CCoT with DCoT or only use parts of the DCoT  
133 to conduct alignment. These weak alignment signals can not effectively help LLM learn CCoT  
134 reasoning, leading to the performance degradation (Hao et al., 2024; Xu et al., 2025b; Shen et al.,  
135 2025b; Cheng & Van Durme, 2024). Due to the limited space, a detailed introduction of the above  
136 related works are shown in Appendix B.137  
138  

## 3 METHODOLOGY

139  
140 In this section, we present the details of our *SynAdpat* framework, which consists of two stages: the  
141 fine-tuning stage and the inference stage, as shown in Figure 2. During the fine-tuning stage, we first  
142 **generate the synthetic CCoT** by optimizing a randomly initialized one. The optimization goal is to  
143 ensure that the LLM generates the correct answer when using the synthetic CCoT. After generation,  
144 we fine-tune the LLM to learn CCoT by **utilizing the synthetic CCoT as the alignment target**.  
145 Specifically, the LLM is trained to iteratively refine a draft CCoT until it aligns with the pre-generated  
146 synthetic CCoT. Additionally, we **train a difficulty classifier** that assesses a question's difficulty  
147 based on both the question itself and its corresponding CCoT.148 During the inference stage, the fine-tuned LLM generates the CCoT for the given question. This  
149 generated CCoT, along with the original question, is then fed into the difficulty classifier to **distinguish**  
150 **between easy and hard questions**. For easy questions, the LLM directly generates the answer based  
151 on the CCoT, ensuring high efficiency. For hard questions, we discard the CCoT and prompt the  
152 LLM to re-think the question step by step, ensuring higher accuracy. More details of the training  
153 stage and the inference stage are presented in Section 3.1 and Section 3.2, respectively.154  
155 

### 3.1 TRAINING STAGE

156 **Synthetic CCoT Generation.** To provide a more effective alignment target to learn CCoT repre-  
157 sentation during fine-tuning LLM, we firstly generate the synthetic CCoT before fine-tuning.  
158159 As shown in the upper-left part of Figure 2, for each question  $Q$ , we randomly initialize a synthetic  
160 CCoT  $Z_{\text{syn}}$  with a fixed length  $m$ . We then concatenate  $Q$  with  $Z_{\text{syn}}$  and an end-of-think token to  
161 form  $[Q, Z_{\text{syn}}, \text{eot}]$ . Given that a well-constructed CCoT should guide the LLM to predict the correct  
162 answer based on the question and CCoT, we make  $Z_{\text{syn}}$  trainable and optimize it by minimizing the




Figure 2: Our *SynAdapt* framework consist of two stage. (1). In **Synthetic CCoT Generation**, we first generate the synthetic CCoT  $Z_{\text{syn}}$  for each question. And then in **Synthetic CCoT Enhanced Fine-tuning**,  $Z_{\text{syn}}$  serves as the full alignment target. By using  $Z_{\text{syn}}$ , we fine-tune the LLM  $\phi$  to effectively learn CCoT, enabling iterative refinement of a randomly initialized draft CCoT. Additionally, we train a **difficulty classifier**  $\delta$  to assess question difficulty based on both the question and the generated CCoT. (2). During the inference stage, we use the fine-tuned LLM  $\phi$  to iteratively refine and generate the final CCoT, while the difficulty classifier  $\delta$  determines the question difficulty. **For easy questions**, the LLM directly generates the output, and **for hard questions**, it is prompted to re-think in order to generate the correct answer.

following loss:

$$\mathcal{L}_{\text{ans}} = -\frac{1}{L_a} \sum_{i=1}^{L_a} \log \mathcal{P}_{\theta}(A_i | Q, Z_{\text{syn}}, \text{eot}, A_{<i}), \quad (1)$$

where  $L_a$  is the length of the answer  $A$ ,  $A_i$  denotes the  $i$ -th token of  $A$ , and  $\theta$  represents the parameters of the LLM.

Moreover, to prevent overfitting during CCoT optimization, we additionally align the hidden state of the eot token when using the synthetic CCoT with that obtained when using DCoT. Assuming  $\mathbf{h}_{\text{eot\_syn}}^l$  is the hidden state of the eot token at the  $l$ -th layer of the LLM when provided with synthetic CCoT  $Z_{\text{syn}}$  and  $\mathbf{h}_{\text{eot\_dcot}}^l$  is that when provided with DCoT, the alignment loss is defined as:

$$\mathcal{L}_{\text{dcot}} = \frac{1}{L} \sum_{l=1}^L \|\mathbf{h}_{\text{eot\_syn}}^l - \mathbf{h}_{\text{eot\_dcot}}^l\|_1, \quad (2)$$

where  $L$  is the total number of layers in the LLM. After optimizing using both  $\mathcal{L}_{\text{ans}}$  and  $\mathcal{L}_{\text{dcot}}$ , we obtain the high-quality synthetic CCoT  $Z_{\text{syn}}$ , which serves a similar function to DCoT but is represented in a denser, continuous format. These  $Z_{\text{syn}}$  can serve as valuable alignment targets during fine-tuning LLM to learn CCoT.

**Synthetic CCoT Enhanced Fine-tuning.** As demonstrated by Saunshi et al. (2025); Yu et al. (2025), iteratively looping an LLM can significantly enhance its reasoning capabilities and refine outputs. Inspired by this, we fine-tune the LLM to iteratively refine the CCoT from a draft in a looping manner instead of generating it autoregressively.

As shown in Figure 2, we concatenate the question  $Q$  with a draft CCoT  $Z_{\text{draft}}^0$ . The  $Z_{\text{draft}}^0$  is initialized as the embedding of a repeated meaningless token sequence (i.e., <T>..<T>), with a fixed length of  $m$ . We input the  $Z_{\text{draft}}^0$  into LLM and use the corresponding output hidden state as the refined one.

216 The iterative refinement process can be formulated as:  
 217  
 218

$$Z_{\text{draft}}^i = f_\phi(Q, Z_{\text{draft}}^{i-1})[L_q :], \quad (3)$$

219 where  $Z_{\text{draft}}^i$  is the CCoT after refining  $i$  iterations,  $L_q$  is the length of the question  $Q$ ,  $\phi$  represents  
 220 the fine-tuned LLM with a trainable LoRA module and  $f_\phi(\cdot)$  returns the output hidden state from  $\phi$ .  
 221 After  $k$  refining iterations, we obtain the final CCoT  $Z_{\text{final}} = Z_{\text{draft}}^k$ . We explicitly align the full  $Z_{\text{final}}$   
 222 with the synthetic CCoT  $Z_{\text{syn}}$  and compute the  $\mathcal{L}_{\text{align}}$  loss as:  
 223

$$\mathcal{L}_{\text{align}} = \|Z_{\text{final}} - Z_{\text{syn}}\|_1. \quad (4)$$

224 Moreover,  $Z_{\text{final}}$  should also guide the initial LLM to generate the correct answer. Therefore, we  
 225 compute an additional losses, similar to Equation 1, as shown below:  
 226

$$\mathcal{L}'_{\text{ans}} = -\frac{1}{L_a} \sum_{i=1}^{L_a} \log \mathcal{P}_\theta(A_i | Q, Z_{\text{final}}, \text{eot}, A_{<i}), \quad (5)$$

$$\mathcal{L}_{\text{refine}} = \mathcal{L}_{\text{align}} + \mathcal{L}'_{\text{ans}}, \quad (6)$$

227 where  $\theta$  represents the initial LLM without the LoRA module. The  $\mathcal{L}_{\text{refine}}$  loss fully utilizes the  
 228 alignment information from  $Z_{\text{align}}$ . After training using  $\mathcal{L}_{\text{refine}}$ , the fine-tuned LLM  $\Phi$  effectively  
 229 learns to iteratively refine the draft CCoT, ultimately generating the final CCoT to replace the original  
 230 redundant DCot.  
 231

232 **Difficulty Classifier Training.** Additionally, we train a difficulty classifier  $\delta$ , composed of two  
 233 MLP layers, to distinguish between hard and easy questions. It takes both the question itself and the  
 234 CCoT as input. Specifically, we construct question pairs  $\langle Q_c, Q_r \rangle$  based on existing difficulty labels  
 235 from the DeepMath dataset (He et al., 2025).  $Q_c$  is a hard question and  $Q_r$  is an easy question. Next,  
 236 we input  $Q_c$  and  $Q_r$  to the fine-tuned LLM  $\phi$  to obtain the corresponding CCoT  $Z_{\text{final}}^c$  and  $Z_{\text{final}}^r$ . Then  
 237 we concatenate  $Q_c$ ,  $Z_{\text{final}}^c$  and one eot token and input to the initial LLM to obtain the output hidden  
 238 state of eot as:  
 239

$$\mathbf{h}_{\text{eot\_final}}^c = f_\theta(Q_c, Z_{\text{final}}^c, \text{eot})[-1], \quad (7)$$

240 where  $f_\theta$  represents the output hidden state from the initial LLM  $\theta$  and  $\mathbf{h}_{\text{eot\_final}}^c$  denotes the output  
 241 hidden state of the eot token. Considering the attention mechanism of LLM,  $\mathbf{h}_{\text{eot\_final}}^c$  can fully capture  
 242 the information in  $Q_c$  and  $Z_{\text{final}}^c$ . Similarly, we compute the  $\mathbf{h}_{\text{eot\_final}}^r$  for the easy question  $Q_r$ . We  
 243 train the difficulty classifier  $\delta$  according to the following loss:  
 244

$$\mathcal{L}_{\text{diff}} = -\log \sigma(f_\delta(\mathbf{h}_{\text{eot\_final}}^c) - f_\delta(\mathbf{h}_{\text{eot\_final}}^r)), \quad (8)$$

245 where  $f_\delta(\cdot)$  denotes the difficulty level predicted by  $\delta$ .  $\mathcal{L}_{\text{diff}}$  encourages the classifier to give higher  
 246 scores for hard question  $Q_c$  and lower scores to easy ones  $Q_r$ . By utilizing additional information  
 247 from the CCoT, the classifier  $\delta$  can more effectively distinguish between hard and easy questions.  
 248

### 249 3.2 INFERENCE STAGE

250 **Adaptive Reasoning via CCoT.** During the inference stage, we concatenate the question with a  
 251 draft CCoT and utilize the fine-tuned LLM  $\phi$  to iteratively refine the draft CCoT to obtain the final  
 252 CCoT. And then we utilized the difficulty classifier to assign the difficulty score based on both the  
 253 question and the CCoT. Questions with a difficulty score below the threshold  $\tau$  are considered easy,  
 254 while those above are regarded as hard.  
 255

256 **For easy questions**, we just append a eot token after the CCoT and prompt the base LLM  $\theta$  to directly  
 257 output answer. The generated CCoT effectively replaces the original discrete CoT reasoning process,  
 258 which often contains numerous tokens and is time-consuming to generate, thereby achieving higher  
 259 efficiency. However, compressing DCot into CCoT inevitably leads to information loss (Nalewajski,  
 260 2011). And as shown by Hao et al. (2024), relying solely on CCoT is insufficient **for hard questions**  
 261 and may even lead to incorrect answer. Therefore, we discard the generated CCoT and prompt the  
 262 LLM to re-think the question via discrete CoT, using a more detailed reasoning process to generate  
 263 the correct answer. Additionally, inspired by Xu et al. (2025a), we explicitly prompt the LLM to  
 264 condense each reasoning step, achieving a better trade-off between accuracy and efficiency.  
 265

266 Moreover, we can dynamically adjust the threshold  $\tau$  to control the ratio of re-thinking. This allows  
 267 our method to simultaneously adapt to both accuracy-sensitive and efficiency-sensitive scenarios  
 268 according the specific requirements of the real application. All our used prompts are provided in  
 269 Appendix H.

270 

## 4 EXPERIMENTS

271  
272 In this section, we conduct comprehensive experiments to demonstrate the effectiveness of our  
273 *SynAdapt* and address the following four key research questions:  
274275 

- **RQ1:** Can *SynAdapt* offer a better accuracy-efficiency trade-off compared to other efficient rea-  
276 soning baselines in both accuracy-sensitive and efficiency-sensitive scenarios? (see section 4.1)
- **RQ2:** Does our difficulty classifier, which uses both the question and CCoT, can effectively  
277 distinguish between hard and easy questions? (see section 4.2)
- **RQ3:** How about the training efficiency of *SynAdapt*? (see section 4.3)
- **RQ4:** How well does *SynAdapt* generalize across more domains (Scientific QA/Coding), LLM  
278 backbones, and hyperparameters? (see Section section 4.4)

  
282283 

### 4.1 EVALUATION OF ACCURACY-EFFICIENCY TRADE-OFF

  
284285 **Experimental Settings** We use DeepMath (He et al., 2025) as the training set and evaluate our  
286 method and baselines on five widely adopted math-related benchmarks: AIME25, AIME24, AMC23,  
287 MATH500 (Lightman et al., 2023) and GSM8K (Cobbe et al., 2021). These datasets cover a diverse  
288 range of math questions across varying difficulty levels. As for the evaluation metrics, we report  
289 accuracy (**Acc**) and generation length (**Len**) to assess both performance and efficiency. Additionally,  
290 we introduce the **Relative Gain** metric (**Rel-G**) defined as:  
291

292 
$$\text{Rel-G} = \frac{\text{Acc}/\text{Acc}_{\text{raw}}}{\text{Len}/\text{Len}_{\text{raw}}}, \quad (9)$$

293 where  $\text{Acc}_{\text{raw}}$  and  $\text{Len}_{\text{raw}}$  denote the accuracy and generation length of the raw model, respectively.  
294 A higher Rel-G indicates a better trade-off between accuracy and efficiency. We also further evaluate  
295 our method on additional domains, including scientific QA and coding, in Section 4.4.296 We adopt DeepSeek-R1-Distill-Qwen-7B as our raw model. We set the length of the synthetic CCoT  
297 to  $m = 512$ , and refining iterations for the draft CCoT to  $k = 4$ . The difficulty score ranges from 0 to  
298 1. In accuracy-sensitive scenarios, we set threshold  $\tau = 0.5$  to route difficult questions for re-thinking.  
299 In efficiency-sensitive scenarios, we set  $\tau = 1.0$  to prompt the LLM to directly generate answers  
300 based on CCoT for higher efficiency. Further details on the datasets and implementation details are  
301 provided in Appendix C.1 and Appendix C.3 respectively.302 **Compared Methods** Here, we consider a broad range of existing efficient reasoning baselines, not  
303 limited to CCoT-based methods. We categorize these baselines into two scenarios, **accuracy-sensitive**  
304 **scenario** and **efficiency-sensitive scenario**, based on their different focuses.  
305306 In the accuracy-sensitive scenario, **CoT-FT** directly uses the full discrete CoT from the training data  
307 to perform supervised fine-tuning (SFT) for improving performance. **TokenSkip** (Xia et al., 2025)  
308 compresses the discrete CoT based on token importance and then applies SFT on the compressed  
309 CoT. **NoThinking** (Ma et al., 2025a) skips the SFT process and directly prompts the model to skip  
310 reasoning and directly generate the answer. **CoD** (Xu et al., 2025a) prompts the model to condense  
311 each reasoning step rather than skipping the reasoning process entirely. **TokenBudget** (Han et al.,  
312 2024) let the LLM to predict a token budget for each question in advance and prompts the model do  
313 not exceed the token budget during further generation.314 In the efficiency-sensitive scenario, **NoCoT-FT** (Yu et al., 2024) discards the discrete CoT and  
315 performs SFT using only the answer to improve efficiency. **SelfTraining** (Munkhbat et al., 2025)  
316 applies best-of- $n$  sampling to extract the shortest correct CoT from the LLM and then fine-tunes the  
317 LLM on these CoT. **Coconut** (Hao et al., 2024), **CompressCoT** (Cheng & Van Durme, 2024), and  
318 **CODI** (Shen et al., 2025b) all belongs to CCoT-based methods, utilizing the CCoT to replace the  
319 DCot for better efficiency. Coconut adopts a curriculum learning strategy to gradually internalize  
320 DCot into CCoT. CompressCoT identifies key tokens in the DCot and aligns the CCoT with their  
321 hidden states. CODI employs self-distillation, aligning the last token hidden state of CCoT with that  
322 of DCot during training. More details of these compared method are provided in Appendix C.2.323 **Main Results** For the **accuracy-sensitive scenario**, as shown in the upper part of Table 1, our  
324 method with  $\tau = 0.5$  outperforms all other baselines by achieving the second-highest average

Methods	AIME25		AIME24		AMC23		MATH500		GSM8K		Average		
	Acc	Len	Acc	Len	Acc	Len	Acc	Len	Acc	Len	Acc ↑	Len ↓	Rel-G ↑
Raw Model	36.7	13348.6	53.3	14071.4	92.5	6315.7	93.2	4087.4	90.7	1110.8	73.3	7786.84	1.00
<i>Accuracy-Sensitive Scenario</i>													
CoT-FT	40.0	16427.3	40.0	15560.6	87.5	7049.3	88.6	3694.0	83.0	700.7	67.8	8686.4	0.83
TokenSkip	30.0	17811.3	36.7	14385.0	70.0	10030.8	78.4	16542.8	81.1	17165.5	59.2	15187.1	0.41
NoThinking	30.0	10623.6	40.0	11099.7	75.0	4143.6	82.4	1355.4	85.7	229.5	62.6	5490.4	1.21
CoD	40.0	10498.0	56.7	8488.5	80.0	2894.3	81.8	1591.1	84.2	286.2	68.5	4751.6	1.53
TokenBudget	36.7	15235.0	53.3	14897.7	82.5	5006.5	90.2	3186.8	86.9	573.0	69.9	7779.8	0.95
<i>SynAdapt</i> ( $\tau=0.5$ )	40.0	10198.3	56.7	8288.1	80.0	2881.6	82.4	1547.7	85.7	258.6	69.0	<b>4694.8</b>	<b>1.58</b>
<i>Efficiency-Sensitive Scenario</i>													
NoCoT-FT	13.3	637.0	10.0	1680.1	50.0	513.1	74.8	478.9	87.1	209.5	47.0	703.7	7.13
SelfTraining	10.0	671.6	10.0	772.7	55.0	627.0	71.6	397.0	85.1	207.6	46.3	<b>535.2</b>	<u>9.10</u>
Coconut	6.7	647.2	13.3	1692.5	52.5	548.0	76.2	426.4	89.3	232.6	47.6	709.3	7.13
CompressCoT	10.0	623.1	6.7	1673.7	52.5	1356.1	75.0	445.8	88.2	207.7	46.5	861.3	5.73
CODI	13.3	2798.7	6.7	613.5	50.0	518.6	72.4	537.5	87.2	238.1	45.9	941.3	5.18
<i>SynAdapt</i> ( $\tau=1.0$ )	13.3	718.8	16.7	620.7	57.5	591.9	75.6	739.4	88.5	253.5	<b>50.3</b>	<u>584.9</u>	<b>9.14</b>
- Synthetic CCoT	10.0	1743.9	16.7	475.9	52.5	510.2	73.2	599.8	87.8	266.7	<u>48.0</u>	719.3	7.10
- Iterative Refine	6.7	767.6	10.0	700.2	50.0	1073.9	76.0	993.8	85.4	728.9	45.6	852.9	5.68

Table 1: Comparison between our *SynAdapt* and efficient reasoning baselines for both Accuracy-Sensitive Scenario and Efficiency-Sensitive Scenario. **For the accuracy-sensitive scenario**, we set the threshold  $\tau = 0.5$  for our method, meaning that questions with a difficulty score greater than 0.5 are routed to re-thinking, while others directly generate an answer based on the CCoT. **For the efficiency-sensitive scenario**, we set  $\tau = 1.0$ , meaning all questions are answered directly using the CCoT to achieve high efficiency. **Bold** and underlined numbers represent the best and second-best average accuracy, generation length and Rel-G score for each scenario.

accuracy while maintaining the shortest average generation length. CoT-FT fine-tunes directly on the full DCoT, improving accuracy on hard questions but also increasing generation length. TokenSkip selects parts of DCoT for fine-tuning, resulting in inconsistent CoT and performance degradation. NoThinking can skip CoT for reducing length, but often causes accuracy drops. CoD condenses each CoT step but cannot skip the unnecessary CoT in simple questions, resulting in a suboptimal accuracy-efficiency trade-off. TokenBudget dynamically allocates more tokens to harder questions, preserving accuracy but not reducing generation length effectively. In contrast, our method identifies hard questions and dynamically re-thinks them while directly generating answers for simple ones. It maintains similar accuracy compared to the raw model while reducing generation length, achieving the highest Rel-G score of 1.55 in the accuracy-sensitive scenario.

For the **efficiency-sensitive scenario**, our method with  $\tau = 1.0$  significantly reduces the average generation length to just 584.9 tokens, while maintaining competitive accuracy compared to other baselines, as shown in the bottom part of Table 1. NoCoT-FT, which fine-tunes only on answers without CoT, leads to the accuracy drop. SelfTraining allows the LLM to search for the shortest correct CoT via best-of- $n$  sampling. But it struggles with harder questions and also results in a substantial drop in accuracy.

The three CCoT-based methods, Coconut, CompressCoT, and CODI, attempt to replace DCoT with CCoT. However, these methods only use a portion of DCoT or the last token as the alignment target when fine-tuning the LLM to learn CCoT. Due to the limited alignment signals, especially for hard questions, they achieve unsatisfactory accuracy. In contrast, our method introduces a more effective alignment target, the synthetic CCoT. By fully leveraging the alignment information from it, we enable more effective fine-tuning. Consequently, our method achieves the highest accuracy and the second shortest generation length in average, yielding the best trade-off with a Rel-G score of 9.14. We also present a representative case study in Figure 4 of Appendix.

Moreover, we evaluate our method under various  $\tau$  values. As shown in Figure 3(a), our method consistently outperforms all other baselines, achieving the best accuracy-efficiency trade-off. As shown in the bottom of Table 1, we observe a significant performance decline when either Synthetic CCoT or Iterative Refinement is removed, which further highlights the importance of both components.

#### 4.2 EVALUATION OF DIFFICULTY CLASSIFIER PERFORMANCE

**Experimental Settings** To evaluate the performance of our difficulty classifier, we use the **MATH500** dataset, treating questions with a difficulty level of 5 as hard and the rest as easy.




Figure 3: (a) Accuracy-efficiency trade-off comparison between our method and other efficient reasoning baselines. (b). Difficulty ratio (The ratio of hard questions) of our method under different  $\tau$  values across five benchmarks. (c/d). Accuracy under various difficulty ratios using different hard question identification methods on the MATH500 and MixD Datasets.

Additionally, we construct the **MixD** dataset by combining AIME25/AIME24/AMC23 and part of GSM8K. Questions from AIME25/AIME24/AMC23 are considered hard, while those from GSM8K are regarded as easy. We report macro precision (**Pre**), macro recall (**Rec**), and macro F1 (**F1**) of the hard question identification. We also report the accuracy of our method using different identification approaches, maintaining the same ratio of hard questions.

**Compared Methods** To demonstrate the effectiveness of our difficulty classifier, we consider several baselines for comparison: **Seq\_PPL** (Mahaut et al., 2024) computes the PPL score for each question, treating those with high PPL as hard and others as easy. **PromptLLM** (Han et al., 2024) directly prompts the LLM to assess question difficulty. **RouteLLM** (Ong et al., 2024) trains an additional BERT model to judge question difficulty. We directly use their released weights. **Probe\_Q** (Azaria & Mitchell, 2023) trains a simple classifier, consisting of two MLP layers, to assess difficulty based solely on the question. More details about the used datasets and the compared baselines are present in Appendix D.1 and D.2, respectively.

**Main Results** As shown in Table 2, our method, which identifies hard questions using both the question and CCoT, outperforms other baselines on both MATH500 and MixD datasets. Seq\_PPL relies solely on the PPL score, which does not strongly correlate with question difficulty. PromptLLM prompts the LLM to assess difficulty, but this approach is unreliable due to the model’s limitations in identifying hard questions. RouteLLM trains an additional BERT-based classifier, which incurs extra costs and struggles to effectively identify complex math questions requiring reasoning. Probe\_Q trains a classifier based only on the question, which can identify explicit hard questions but misses those that look simple but actually hard. In contrast, our method can effectively identify those hard questions by using the reasoning information in corresponding CCoT. As shown in Figure 3 (b), it accurately identifies most difficult questions, such as those in AIME25/24, and AMC23.

Moreover, we also report the impact of different identification methods on overall performance in Figure 3 (c/d). We evaluate the problem-solving accuracy when using these methods under different difficulty ratios. As shown in Figure 3 (c/d), at the same difficulty ratio, our method can more accurately identify hard questions, route them for re-thinking, and achieve the best accuracy on both the MATH500 and MixD datasets. However, we observe a decrease in accuracy when the difficulty ratio exceeds 0.6. This is because easy questions are also routed for re-thinking, and excessive reasoning for simple questions will confuse the model, leading to incorrect answers.

#### 4.3 ANALYSIS OF TRAINING EFFICIENCY

432 To evaluate training efficiency, we report the training  
 433 cost of our method and other CCoT-based  
 434 methods. As shown in Table 3, our method of-  
 435 fers comparable efficiency to the baselines. While  
 436 *SynAdapt* introduces additional synthetic CCoT  
 437 generation, this process is highly efficient, ac-  
 438 counting for **only 9.89%** of the total training cost.  
 439 **Single CCoT generation only requires 10 sec-  
 440 onds**, which is very fast.

441 CompressCoT and CODI require autoregressive  
 442 generation of CCoT during fine-tuning, leading to high training costs and low efficiency. Coconut  
 443 gradually internalizes DCoT, and since the initial CCoT length is small, the training cost is relatively  
 444 low. However, in the later stages, the cost still increases due to autoregressive generation. In contrast,  
 445 ***SynAdapt* iteratively refines a draft CCoT rather than generating it autoregressively, effectively**  
 446 **improving efficiency**. Therefore, our method achieves high training efficiency, demonstrating its  
 447 practicality.

Modules	Time (min)	Percentage
Coconut	740	-
CompressCoT	1192	-
CODI	1156	-
<i>SynAdapt</i>	1021	100%
LLM Training	920	90.11%
<b>Synthetic CCoT Generation (bs=16)</b>	<b>101</b>	<b>9.89%</b>
⇒ Single Synthetic CCoT Generation	10s	0.02%

Table 3: Training time costs for different CCoT-based methods. We use a batch size (bs) of 16 during synthetic CCoT generation.

#### 449 4.4 GENERALIZATION EVALUATION AND HYPERPARAMETER ANALYSIS

450 To further demonstrate the generalization ability of *SynAdapt*, we evaluate it on **more domains**,  
 451 including scientific question answering (GPQA-Diamond (Rein et al., 2024)) and code generation  
 452 (HumanEval (Chen et al., 2021) and LiveCodeBench (Naman Jain, 2024)). As shown in Table 4,  
 453 *SynAdapt* also exhibits superior performance in both scientific QA and coding tasks. With  $\tau = 0.5$   
 454 for identifying hard questions requiring rethinking, our method achieves performance comparable to  
 455 the raw model while reducing generation length. And with  $\tau = 1.0$ , which means no rethinking of  
 456 any questions, *SynAdapt* still outperforms all other CCoT-based baselines. More results and analyses  
 457 of our method on LiveCodeBench are provided in Appendix F. We also evaluate our method on **more**  
 458 **LLM backbones**, such as DeepSeek-R1-Distill-Llama-8B and DeepSeek-R1-Distill-Qwen-1.5B  
 459 (Guo et al., 2025). As shown in Table 5, *SynAdapt* consistently demonstrates superior performance  
 460 under both  $\tau = 0.5$  and  $\tau = 1.0$  settings.

461 We conduct the **hyperparameter analysis** about CCoT length  $m$  and refining iterations  $k$ . As  
 462 shown in Figures 7 and 8 in the Appendix, our method remains effective and robust across various  
 463 hyperparameter settings. Due to page limitations, more analyses and results are in Appendix G.

Methods	GPQA-Diamond			HumanEval			Methods	R1-Llama-8B			R1-Qwen-1.5B		
	Acc $\uparrow$	Len $\downarrow$	Rel-G $\uparrow$	Pass@1 $\uparrow$	Len $\downarrow$	Rel-G $\uparrow$		Acc $\uparrow$	Len $\downarrow$	Rel-G $\uparrow$	Acc $\uparrow$	Len $\downarrow$	Rel-G $\uparrow$
Raw Model	<b>47.9</b>	7847.1	1.00	<b>75.6</b>	4366.5	1.00	Raw Model	<b>67.2</b>	7998.4	1.00	<b>57.6</b>	9166.2	1.00
<i>SynAdapt</i> ( $\tau=0.5$ )	47.5	<b>6047.0</b>	<b>1.28</b>	73.2	<b>3503.6</b>	<b>1.21</b>	<i>SynAdapt</i> ( $\tau=0.5$ )	66.1	<b>6406.2</b>	<b>1.23</b>	57.3	<b>8836.5</b>	<b>1.03</b>
Coconut	<b>42.9</b>	1406.6	4.99	70.7	750.6	5.44	Coconut	45.5	572.6	9.46	39.6	1767.1	3.57
CompressCoT	41.4	782.9	8.65	71.2	1386.5	2.97	CompressCoT	44.6	1834.3	2.89	38.2	1166.0	5.21
CODI	40.9	676.6	9.89	65.9	<b>602.6</b>	6.32	CODI	38.3	<b>488.2</b>	9.34	40.1	1566.5	4.07
<i>SynAdapt</i> ( $\tau=1.0$ )	42.4	<b>660.2</b>	<b>10.51</b>	<b>72.0</b>	622.4	<b>6.68</b>	<i>SynAdapt</i> ( $\tau=1.0$ )	<b>48.0</b>	582.7	<b>9.80</b>	<b>42.1</b>	<b>690.8</b>	<b>9.70</b>

Table 4: Evaluation of our method across more domains, including GPQA-Diamond (Rein et al., 2024) for scientific question answering and HumanEval (Chen et al., 2021) for code generation. Table 5: Evaluation of our method on DeepSeek-R1-Distill-Llama-8B and DeepSeek-R1-Distill-Qwen-1.5B backbones. We report the average results across all five math benchmarks.

## 478 5 CONCLUSION

479 We propose a novel and efficient reasoning framework, *SynAdapt*, designed to help LLMs learn  
 480 continuous CoT (CCoT). Before fine-tuning, we generate the synthetic CCoT, which serves as a  
 481 more effective alignment target for learning CCoT. Additionally, we train a difficulty classifier  
 482 that identifies hard questions by considering both the question and its corresponding CCoT. By  
 483 dynamically prompting the LLM to re-think hard questions, our method can adapt to both accuracy-  
 484 sensitive and efficiency-sensitive scenarios. Extensive experimental results across various benchmarks,  
 485 domains and LLM backbones consistently demonstrate the effectiveness of *SynAdapt* for efficient  
 reasoning.

486 REPRODUCIBILITY STATEMENT  
487488 We provide the processing details of our training datasets and evaluation benchmarks in Sections C.1  
489 and D.1. The statistic of these dataset are shown in Table 6 and 7. The implementation details of  
490 *SynAdapt* are provided in Section C.3 to facilitate reproducibility. We report the full settings used  
491 during LLM training and evaluation, including all hyperparameters. In addition, we have conducted  
492 experiment to analyze the impact of the hyperparameters in Section G and explain why we choose  
493 these setting.494 To further facilitate reproducibility, we release all source code and the datasets used in our experiments  
495 in the supplementary materials. An anonymous repository containing the code and datasets is  
496 also provided for easy access by reviewers: [https://anonymous.4open.science/r/SynAdapt\\_Review-E677](https://anonymous.4open.science/r/SynAdapt_Review-E677). The repository includes a detailed user guide in the README files, covering installation,  
497 dependencies, and usage instructions.  
498499  
500 ETHICS STATEMENT  
501502 We adhere to the ICLR Code of Ethics in all aspects of this work. Our research utilizes exclusively  
503 publicly available repositories and datasets, ensuring full transparency and reproducibility. To  
504 rigorously validate the effectiveness of our proposed method and minimize the impact of randomness,  
505 we conduct extensive evaluations across a diverse range of domain tasks, large language model (LLM)  
506 backbones, and hyperparameter combinations. We not only evaluate the inference performance of our  
507 method, but also consider its training efficiency as critical factors in our analysis. To ensure a fair and  
508 comprehensive comparison, we rigorously assess training efficiency under consistent experimental  
509 conditions. All experiments are designed and reported in accordance with principles of responsible  
510 research, and we have conscientiously considered potential societal impacts in our work.  
511  
512  
513  
514  
515  
516  
517  
518  
519  
520  
521  
522  
523  
524  
525  
526  
527  
528  
529  
530  
531  
532  
533  
534  
535  
536  
537  
538  
539

540 REFERENCES  
541

542 Pranjal Aggarwal and Sean Welleck. L1: Controlling how long a reasoning model thinks with  
543 reinforcement learning. *arXiv preprint arXiv:2503.04697*, 2025.

544 Daman Arora and Andrea Zanette. Training language models to reason efficiently. *arXiv preprint*  
545 *arXiv:2502.04463*, 2025.

546 Amos Azaria and Tom Mitchell. The internal state of an llm knows when it's lying. *arXiv preprint*  
547 *arXiv:2304.13734*, 2023.

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

559 Jeffrey Cheng and Benjamin Van Durme. Compressed chain of thought: Efficient reasoning through  
560 dense representations. *arXiv preprint arXiv:2412.13171*, 2024.

561 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,  
562 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John  
563 Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*,  
564 2021.

565 Yuntian Deng, Yejin Choi, and Stuart Shieber. From explicit cot to implicit cot: Learning to internalize  
566 cot step by step. *arXiv preprint arXiv:2405.14838*, 2024.

567 Yihong Dong, Xue Jiang, Huanyu Liu, Zhi Jin, Bin Gu, Mengfei Yang, and Ge Li. Generalization or  
568 memorization: Data contamination and trustworthy evaluation for large language models. *arXiv*  
569 *preprint arXiv:2402.15938*, 2024.

570 Sachin Goyal, Ziwei Ji, Ankit Singh Rawat, Aditya Krishna Menon, Sanjiv Kumar, and Vaishnavh  
571 Nagarajan. Think before you speak: Training language models with pause tokens. *arXiv preprint*  
572 *arXiv:2310.02226*, 2023.

573 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,  
574 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms  
575 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

576 Tingxu Han, Zhenting Wang, Chunrong Fang, Shiyu Zhao, Shiqing Ma, and Zhenyu Chen. Token-  
577 budget-aware llm reasoning. *arXiv preprint arXiv:2412.18547*, 2024.

578 Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong  
579 Tian. Training large language models to reason in a continuous latent space. *arXiv preprint*  
580 *arXiv:2412.06769*, 2024.

581 Zhiwei He, Tian Liang, Jiahao Xu, Qiuzhi Liu, Xingyu Chen, Yue Wang, Linfeng Song, Dian  
582 Yu, Zhenwen Liang, Wenxuan Wang, et al. Deepmath-103k: A large-scale, challenging, de-  
583 contaminated, and verifiable mathematical dataset for advancing reasoning. *arXiv preprint*  
584 *arXiv:2504.11456*, 2025.

585 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,  
586 Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.

587 Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris  
588 Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al.  
589 Mixtral of experts. *arXiv preprint arXiv:2401.04088*, 2024.

594 Nan Jiang, Ziming Wu, De-Chuan Zhan, Fuming Lai, and Shaobing Lian. Dart: Distilling autoregres-  
 595 sive reasoning to silent thought. *arXiv preprint arXiv:2506.11752*, 2025.  
 596

597 Jean Kaddour, Joshua Harris, Maximilian Mozes, Herbie Bradley, Roberta Raileanu, and Robert  
 598 McHardy. Challenges and applications of large language models. *arXiv preprint arXiv:2307.10169*,  
 599 2023.

600 Yu Kang, Xianghui Sun, Liangyu Chen, and Wei Zou. C3ot: Generating shorter chain-of-thought with-  
 601 out compromising effectiveness. In *Proceedings of the AAAI Conference on Artificial Intelligence*,  
 602 volume 39, pp. 24312–24320, 2025.  
 603

604 Philipp Koehn and Rebecca Knowles. Six challenges for neural machine translation. *arXiv preprint*  
 605 *arXiv:1706.03872*, 2017.

606 Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large  
 607 language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:  
 608 22199–22213, 2022.  
 609

610 Ayeong Lee, Ethan Che, and Tianyi Peng. How well do llms compress their own chain-of-thought? a  
 611 token complexity approach. *arXiv preprint arXiv:2503.01141*, 2025.

612 Manling Li, Shiyu Zhao, Qineng Wang, Kangrui Wang, Yu Zhou, Sanjana Srivastava, Cem Gokmen,  
 613 Tony Lee, Erran Li Li, Ruohan Zhang, et al. Embodied agent interface: Benchmarking llms for  
 614 embodied decision making. *Advances in Neural Information Processing Systems*, 37:100428–  
 615 100534, 2024a.  
 616

617 Zhiyuan Li, Hong Liu, Denny Zhou, and Tengyu Ma. Chain of thought empowers transformers to  
 618 solve inherently serial problems. *arXiv preprint arXiv:2402.12875*, 1, 2024b.

619 Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan  
 620 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. *arXiv preprint*  
 621 *arXiv:2305.20050*, 2023.

622 Zicheng Lin, Tian Liang, Jiahao Xu, Qiuqhi Lin, Xing Wang, Ruilin Luo, Chufan Shi, Siheng Li,  
 623 Yujiu Yang, and Zhaopeng Tu. Critical tokens matter: Token-level contrastive estimation enhances  
 624 llm’s reasoning capability. *arXiv preprint arXiv:2411.19943*, 2024.  
 625

626 Lei Liu, Xiaoyan Yang, Junchi Lei, Yue Shen, Jian Wang, Peng Wei, Zhixuan Chu, Zhan Qin, and  
 627 Kui Ren. A survey on medical large language models: Technology, application, trustworthiness,  
 628 and future directions. *arXiv preprint arXiv:2406.03712*, 2024.  
 629

630 Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. The ai scientist:  
 631 Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292*, 2024.

632 Haotian Luo, Li Shen, Haiying He, Yibo Wang, Shiwei Liu, Wei Li, Naiqiang Tan, Xiaochun Cao,  
 633 and Dacheng Tao. O1-pruner: Length-harmonizing fine-tuning for o1-like reasoning pruning.  
 634 *arXiv preprint arXiv:2501.12570*, 2025.

635 Wenjie Ma, Jingxuan He, Charlie Snell, Tyler Griggs, Sewon Min, and Matei Zaharia. Reasoning  
 636 models can be effective without thinking. *arXiv preprint arXiv:2504.09858*, 2025a.  
 637

638 Xinyin Ma, Guangnian Wan, Runpeng Yu, Gongfan Fang, and Xinchao Wang. Cot-valve: Length-  
 639 compressible chain-of-thought tuning. *arXiv preprint arXiv:2502.09601*, 2025b.  
 640

641 Matéo Mahaut, Laura Aina, Paula Czarnowska, Momchil Hardalov, Thomas Müller, and Lluís  
 642 Márquez. Factual confidence of llms: on reliability and robustness of current estimators. *arXiv*  
 643 *preprint arXiv:2406.13415*, 2024.

644 mathai. Aime2025 dataset. <https://huggingface.co/datasets/math-ai/aime25>,  
 645 2024.

646 Maxwell-Jia. Aime2024 dataset. [https://huggingface.co/datasets/Maxwell-Jia/AIME\\_2024](https://huggingface.co/datasets/Maxwell-Jia/AIME_2024), 2024.

648 Tergel Munkhbat, Namgyu Ho, Seo Hyun Kim, Yongjin Yang, Yujin Kim, and Se-Young Yun.  
 649 Self-training elicits concise reasoning in large language models. *arXiv preprint arXiv:2502.20122*,  
 650 2025.

651 Roman F Nalewajski. Elements of information theory. In *Perspectives in Electronic Structure Theory*,  
 652 pp. 371–395. Springer, 2011.

653 Alex Gu Wen-Ding Li Fanjia Yan Tianjun Zhang Sida Wang Armando Solar-Lezama Koushik Sen  
 654 Ion Stoica Naman Jain, King Han. Livecodebench: Holistic and contamination free evaluation of  
 655 large language models for code. *arXiv preprint*, 2024.

656 Isaac Ong, Amjad Almahairi, Vincent Wu, Wei-Lin Chiang, Tianhao Wu, Joseph E Gonzalez,  
 657 M Waleed Kadous, and Ion Stoica. Routellm: Learning to route llms with preference data. *arXiv  
 658 preprint arXiv:2406.18665*, 2024.

659 OpenAI. Learning to reason with llms, 2025.

660 Zhuoshi Pan, Qianhui Wu, Huiqiang Jiang, Menglin Xia, Xufang Luo, Jue Zhang, Qingwei Lin,  
 661 Victor Rühle, Yuqing Yang, Chin-Yew Lin, et al. Llmlingua-2: Data distillation for efficient and  
 662 faithful task-agnostic prompt compression. *arXiv preprint arXiv:2403.12968*, 2024.

663 Jacob Pfau, William Merrill, and Samuel R Bowman. Let’s think dot by dot: Hidden computation in  
 664 transformer language models. *arXiv preprint arXiv:2404.15758*, 2024.

665 Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. Deepspeed: System opti-  
 666 mizations enable training deep learning models with over 100 billion parameters. In *Proceedings  
 667 of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pp.  
 668 3505–3506, 2020.

669 David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani,  
 670 Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a benchmark. In  
 671 *First Conference on Language Modeling*, 2024.

672 Matthew Renze and Erhan Guven. The benefits of a concise chain of thought on problem-solving in  
 673 large language models. In *2024 2nd International Conference on Foundation and Large Language  
 674 Models (FLLM)*, pp. 476–483. IEEE, 2024.

675 Nikunj Saunshi, Nishanth Dikkala, Zhiyuan Li, Sanjiv Kumar, and Sashank J Reddi. Reasoning with  
 676 latent thoughts: On the power of looped transformers. *arXiv preprint arXiv:2502.17416*, 2025.

677 Yi Shen, Jian Zhang, Jieyun Huang, Shuming Shi, Wenjing Zhang, Jiangze Yan, Ning Wang, Kai  
 678 Wang, Zhaoxiang Liu, and Shiguo Lian. Dast: Difficulty-adaptive slow-thinking for large reasoning  
 679 models. *arXiv preprint arXiv:2503.04472*, 2025a.

680 Zhenyi Shen, Hanqi Yan, Linhai Zhang, Zhanghao Hu, Yali Du, and Yulan He. Codi: Compressing  
 681 chain-of-thought into continuous space via self-distillation. *arXiv preprint arXiv:2502.21074*,  
 682 2025b.

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

686 Heming Xia, Chak Tou Leong, Wenjie Wang, Yongqi Li, and Wenjie Li. Tokenskip: Controllable  
 687 chain-of-thought compression in llms. *arXiv preprint arXiv:2502.12067*, 2025.

688 Silei Xu, Wenhao Xie, Lingxiao Zhao, and Pengcheng He. Chain of draft: Thinking faster by writing  
 689 less. *arXiv preprint arXiv:2502.18600*, 2025a.

690 Yige Xu, Xu Guo, Zhiwei Zeng, and Chunyan Miao. Softcot: Soft chain-of-thought for efficient  
 691 reasoning with llms. *arXiv preprint arXiv:2502.12134*, 2025b.

692 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang  
 693 Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*,  
 694 2025.

695

702 Edward Yeo, Yuxuan Tong, Morry Niu, Graham Neubig, and Xiang Yue. Demystifying long  
703 chain-of-thought reasoning in llms. *arXiv preprint arXiv:2502.03373*, 2025.  
704

705 Ping Yu, Jing Xu, Jason Weston, and Ilia Kulikov. Distilling system 2 into system 1. *arXiv preprint*  
706 *arXiv:2407.06023*, 2024.

707 Qifan Yu, Zhenyu He, Sijie Li, Xun Zhou, Jun Zhang, Jingjing Xu, and Di He. Enhancing auto-  
708 regressive chain-of-thought through loop-aligned reasoning. *arXiv preprint arXiv:2502.08482*,  
709 2025.

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

713 Hanlin Zhu, Shibo Hao, Zhiting Hu, Jiantao Jiao, Stuart Russell, and Yuandong Tian. Reasoning  
714 by superposition: A theoretical perspective on chain of continuous thought. *arXiv preprint*  
715 *arXiv:2505.12514*, 2025.

716 zwhe99. Amc23 dataset. <https://huggingface.co/datasets/zwhe99/amc23>, 2024.

717

718

719

720

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756 THE USE OF LARGE LANGUAGE MODELS  
757758 In this paper, we strictly adhere to the usage policies of large language models (LLMs). LLMs were  
759 employed solely to assist with language polishing and to improve the readability of the manuscript.  
760 All generated content was carefully reviewed and verified by the authors before inclusion in the paper  
761 to ensure accuracy and integrity. No LLM outputs were used in a manner that could compromise  
762 reproducibility, scientific validity, or compliance with ethical standards.  
763764 A APPENDIX OVERVIEW  
765766 The appendix is organized into two main parts: Appendices B–D provide detailed related works  
767 and more experimental setup of our *SynAdapt*. Appendices E–G present additional case studies and  
768 experimental results, further demonstrating the effectiveness of our *SynAdapt*.  
769770 B DETAILS OF RELATED WORK  
771772 In this section, we provide a detailed overview of related works on LLM efficient reasoning, which  
773 can be broadly categorized into four main types: **SFT-based methods**, **RL-based methods**, **Prompt-  
774 based methods**, and **CCoT-based methods**.  
775776 For **SFT-based methods**, Yu et al. (2024) proposes to collect CoT and answer data from reasoning  
777 LLMs and directly discard the CoT part. And then they fine-tune the LLM using only the answers  
778 to help the model reduce reasoning length. Ma et al. (2025b) fine-tunes the LLM simultaneously  
779 on data with CoT and data without CoT, using specific instructions to distinguish between the two.  
780 During inference, they use the instructions to prevent the model from outputting CoT. Munkhbat et al.  
781 (2025) applies best-of-n sampling to LLM, selecting the shortest CoT, and fine-tune the model on  
782 these short CoTs to reduce reasoning length. Xia et al. (2025) assesses the semantic importance of  
783 tokens in the initial CoT, retaining only the most important tokens for fine-tuning the LLM. Kang  
784 et al. (2025) dynamically samples simplified CoTs from the model after each fine-tuning epoch for  
785 the next round of fine-tuning.786 All of the above methods either discard the CoT or use a simplified version for fine-tuning the LLM  
787 to reduce reasoning length. While these approaches effectively shorten the reasoning length, they  
788 overlook important details in the original CoT, leading to significant performance degradation during  
789 further fine-tuning.790 For **RL-based methods**, Arora & Zanette (2025) introduces a length-based reward, where shorter  
791 correct answers receive higher rewards, and uses policy gradient (PG) methods to fine-tune the LLM  
792 to reduce reasoning length. Luo et al. (2025) enhances this reward by comparing the generated  
793 answer length to a reference answer and applies PPO optimization for LLM fine-tuning. Yeo et al.  
794 (2025) further introduces a cosine-based reward and applies a penalty for exceeding the length limit.  
795 Aggarwal & Welleck (2025) uses length-constrained prompts to sampling data during RL fine-tuning.  
796 Shen et al. (2025a) employs SimPO to fine-tune the LLM using a length-preference dataset.  
797798 Although these RL-based methods can reduce reasoning length to some extent while maintaining  
799 LLM performance, RL fine-tuning requires significant resources. For example, they need to repeatedly  
800 sample new data for updating the action of LLM. Moreover, the reduction in length is limited and  
801 cannot be applied to those efficiency-sensitive scenarios. For instance, in real-life medical QA  
802 scenarios, efficiency is critical. Diagnosis advice must be concise, enabling doctors and patients to  
803 quickly access key details and conclusions, especially in emergencies. Previous studies (Kaddour  
804 et al., 2023; Liu et al., 2024) have highlighted that overly long responses can lead to errors, such as  
805 confusing similar drug names or omitting critical contraindications.806 For **Prompt-based methods**, Renze & Guven (2024) proposes to prompt the LLM to perform CoT  
807 reasoning while explicitly instructing it to be concise. Xu et al. (2025a) focuses on adding instructions  
808 in the prompt to condense each reasoning step and limit verbosity. Lee et al. (2025) explores various  
809 prompt types to reduce reasoning length, such as prompting to output only numbers or only use bullet  
810 points. Han et al. (2024) estimates a token budget for each question, allocating more tokens for harder  
811 questions, and instructing the LLM to stay within this budget during reasoning for efficiency.

810 Most prompt-based methods reduce reasoning length by adding additional length constraint instruc-  
 811 tions in the prompt. While this approach is low-cost, its impact on reducing length is limited. LLMs  
 812 still tend to generate redundant reasoning CoTs, especially when faced with hard questions.  
 813

814 For **CCoT-based methods**, Hao et al. (2024) was the first to propose to fine-tune the LLM to reason  
 815 continuously and utilize the last hidden state as the continuous CoT (CCoT) to replace traditional  
 816 discrete CoT (DCoT), which often contain redundant tokens. They introduce curriculum learning to  
 817 gradually replace DCoT with CCoT during fine-tuning, without explicit alignment with the original  
 818 DCoT. Xu et al. (2025b) is similar to Coconut, but it incorporates an additional assistant LLM  
 819 with a projection module to generate the CCoT. Although it provides slight improvements, it also  
 820 incurs additional resource costs. Shen et al. (2025b) employs self-distillation to learn CCoT by  
 821 simultaneously fine-tuning on both DCoT and CCoT and explicitly aligns the last token hidden state  
 822 between the two. Cheng & Van Durme (2024) measures token importance in advance and aligns the  
 823 CCoT only with the hidden states of those important tokens in the DCoT.  
 824

825 Current CCoT-based methods can successfully compress reasoning steps into a latent space, replacing  
 826 the original DCoT with a more efficient CCoT and significantly reducing generation length. However,  
 827 they often suffer from unsatisfactory performance degradation. This is mainly because they either do  
 828 not apply explicit alignment between DCoT and CCoT or only use partial DCoT (e.g., the last token  
 829 or a subset of important tokens) to supervise CCoT learning. These weak supervisory signals fail to  
 830 help LLM to learn a well CCoT representation, leading to significant performance drops. Therefore,  
 831 designing stronger supervisory signals for CCoT learning is crucial for real-world applications.  
 832

## 833 C DETAILS OF ACCURACY-EFFICIENCY TRADE-OFF EVALUATION

### 834 C.1 DATASET DETAILS

835 For the training set, we use the DeepMath-103K dataset He et al. (2025), which contains numerous  
 836 math problems with three distinct reasoning paths from DeepSeek-R1 Guo et al. (2025), covering  
 837 various math topics and difficulty levels. For each question, we randomly select one reasoning path  
 838 as the discrete CoT and exclude samples with reasoning paths exceeding 12,000 tokens. Moreover,  
 839 as pointed out by Dong et al. (2024), the public datasets, containing numerous samples, suffer from a  
 840 'data contamination' issue, where some samples may be similar to evaluation benchmark. Directly  
 841 training on this data may cause the model to memorize these samples, leading to unnaturally high  
 842 performance. Additionally, including too many training samples introduces excessive training costs,  
 843 which contradicts our goal of high efficiency. Therefore, we only sample a portion of the original  
 844 DeepMath-103K dataset for training. Specifically, we randomly sample 10% of the training samples  
 845 for each difficulty level to create the final **DeepMath** dataset, ensuring the distribution of question  
 846 difficulty remains consistent. The total size of the DeepMath dataset is 9,660.  
 847

848 For the test set, we consider several widely adopted math-related benchmarks: **AIME25** mathai.  
 849 (2024), **AIME24** Maxwell-Jia. (2024), **AMC23** zwhe99. (2024), **MATH500** Lightman et al. (2023),  
 850 and **GSM8K** Cobbe et al. (2021). The difficulty of these benchmarks gradually decreases, covering a  
 851 wide range from complex math competitions to simple grade school math. The details of both the  
 852 train and test dataset sizes are shown in Table 6.  
 853

854 Train Dataset	Test Dataset					
	855 DeepMath	AIME25	AIME24	AMC23	MATH500	GSM8K
9660	856	30	30	40	500	1319

857 Table 6: The size of our used train dataset and five math-related evaluation benchmarks, covering  
 858 various difficulty levels.  
 859

### 860 C.2 BASELINES DETAILS

861 Here, we provide more details about all the compared efficient reasoning baselines. We consider  
 862 not only CCoT-based baselines but also other SFT-based and prompt-based methods. We exclude  
 863

864 RL-based methods, as these require substantial resources to apply RL learning to LLMs, making  
 865 them inefficient and impractical for real-world applications.  
 866

867 We further mainly categorize these baselines into two scenarios based on their focus. Baselines for the  
 868 **accuracy-sensitive scenario** primarily aim to maintain performance while shortening the generation  
 869 length. Here are the details of these baselines:

870 **CoT-FT** belongs to SFT-based methods. We directly uses the CoT and answers from the training  
 871 set, to supervise fine-tune (SFT) the LLM. This method aims to maintain accuracy while slightly  
 872 reducing the generation length.

873 **TokenSkip** (Xia et al., 2025) belongs to SFT-based methods. As proposed by TokenSkip, different  
 874 tokens in the CoT have varying semantic importance, and tokens with low semantic value can be  
 875 skipped during SFT of the LLM. Specifically, we use LLMLingua-2 (Pan et al., 2024) to assess the  
 876 importance of each token and obtain a compressed CoT. We set the compression ratio to 0.7 because  
 877 too low ratio will make the CoT inconsistent for fine-tuning while too low ratio only provides a slight  
 878 reduction in generation length. We utilize the compressed CoT along with corresponding answer to  
 879 fine-tune LLM to reduce generation length while maintaining performance.

880 **NoThinking** (Ma et al., 2025a) is a prompt-based method. NoThinking proposes to directly prompt  
 881 the LLM to avoid generating a CoT, which effectively reduces the generation length with fine-tuning  
 882 process. Specifically, we append the instruction “*Okay, I think I have finished thinking.</think>*” to  
 883 the initial prompt, instructing the LLM to skip reasoning and directly output the answer without CoT.

884 **CoD** (Xu et al., 2025a) is another prompt-based method. Different from NoThinking directly prompts  
 885 LLM to skip reasoning and do not output CoT, Chain-of-Draft (CoD) preserves the reasoning process  
 886 but condenses each reasoning step by inserting the “*only keep a minimum draft for each thinking step,  
 887 with 5 words at most.*” instruction.

888 **TokenBudget** (Han et al., 2024) is also a prompt-based method. Following TokenBudget, we prompt  
 889 the LLM in advance to estimate the difficulty of each question and determine the essential token  
 890 budget. During inference, we incorporate the token budget into the initial prompt by adding the  
 891 instruction, “*Let’s think step by step and use fewer than [[Token Budget]] tokens*”, guiding the LLM  
 892 to reduce unnecessary generation.

893 In contrast, baselines for the **efficiency-sensitive scenario** prioritize improving efficiency, even at the  
 894 cost of performance. Here are the details of these baselines:

895 **NoCoT-FT** (Yu et al., 2024) is an SFT-based method. However, unlike previous SFT-based methods,  
 896 NoCoT-FT distills the ability from the reasoning model to the model that does not output any CoT,  
 897 by fine-tuning solely on the answer part from the reasoning model. Specifically, we discard the CoT  
 898 part in our training set and fine-tune the LLM only with the answer.

900 **SelfTraining** (Munkhbat et al., 2025) is another SFT-based method. As proposed by SelfTraining,  
 901 we apply best-of-n sampling to the LLM to generate multiple answers for each question, then select  
 902 the shortest correct answer to fine-tune the LLM and reduce generation length. During sampling, we  
 903 also provide demonstrations as few-shots to instruct the LLM to generate the answer directly without  
 904 CoT. The sampled answers are then used to fine-tune the LLM to skip the CoT.

905 **Coconut** (Hao et al., 2024) is one of CCoT-based methods. According to Coconut, we apply  
 906 curriculum learning to help the LLM gradually learn Continous CoT (CCoT). Specifically, we fine-  
 907 tune the LLM for 3 epochs, gradually reducing the initial DCoT tokens to none as the epochs progress,  
 908 and replacing them with CCoT. Finally, we can internalize the DCoT into the CCoT.

909 **CompressCoT** (Cheng & Van Durme, 2024) belongs to CCoT-based methods. Following Compress-  
 910 CoT, we first identify important tokens in the discrete CoT using LLMLingua-2 (Pan et al., 2024) and  
 911 compute the mid-layer hidden states of these tokens as the target. We then fine-tune the LLM with the  
 912 LoRA module to generate the CCoT similar to target. Simultaneously, we fine-tune another LoRA  
 913 module to predict the correct answer based on the CCoT. During inference, we first use the prior  
 914 LoRA module to generate the CCoT and then use the other LoRA module to generate the answer  
 915 based on it.

916 **CODI** (Shen et al., 2025b) is another CCoT-based methods. As proposed by CODI, we fine-tune the  
 917 LLM with two tasks: the teacher task, which generates the discrete CoT tokens and the final correct

918 answer, and the student task, which generates the CCoT and the correct answer. We then explicitly  
 919 align the last token hidden states from the DCoT and CCoT to achieve self-distillation from DCoT to  
 920 CCoT.

921

### 922 C.3 IMPLEMENTATION DETAILS OF OUR *SynAdapt*

923

924 We adopt the DeepSeek-R1-Distill-Qwen-7B (Guo et al., 2025) as the LLM backbone and we also  
 925 evaluate the our method on other backbones in Section 4.4. For the **Synthetic CCoT generation**, we  
 926 fix the LLM backbone and make the randomly initialized synthetic CCoT to be trainable. The length  
 927 of synthetic CCoT is set as  $m = 512$  and the we optimize it using the learning rate at 1e-3 for 32  
 928 steps. During optimization, we use a batch size of 16 to ensure high efficiency.

929 For **Synthetic CCoT Enhanced Fine-tuning**, we use LoRA (Hu et al., 2022) to fine-tune the LLM  
 930 for learning CCoT. The lora rank is set to be 8 and the alpha value at 32. We use the DeepSpeed  
 931 (Rasley et al., 2020) framework to fine-tune the LLM. We fine-tine LLM for 3 epochs with a batch  
 932 size of 1 and a gradient accumulation step of 16. We employ the AdamW optimizer with a learning  
 933 rate set to 4e-5. The refinement steps of the draft is  $k = 4$  and the length of CCoT is also  $m = 512$ .  
 934 We also analyze these hyperparameters in Section 4.4.

935 For **Adaptive Reasoning via CCoT**, we firstly generate the CCoT with the length of 512 and use  
 936 the difficulty classifier to judge the difficulty score  $\tau$  based on question and CCoT. The score ranges  
 937 from 0 to 1, with scores below the threshold  $\tau$  considered as simple, and those above as hard. For  
 938 the efficiency-sensitive scenario, we set  $\tau = 1.0$ , treating all questions as simple. For the accuracy-  
 939 sensitive scenario, we set  $\tau = 0.5$  to classify some questions as hard. We also try more  $\tau$  values,  
 940 as shown in Figure 3(a). During answer generation, we use greedy decoding and set the maximum  
 941 generation length to 32,768 tokens. The generation prompt and the prompt for re-thinking hard  
 942 questions are provided in Appendix H. All our training and evaluation experiments are conducted on  
 943 the H20 GPU.

944

## 945 D DETAILS OF DIFFICULTY CLASSIFIER EVALUATION

946

947 

### D.1 DATASET DETAILS

948 Here, we will introduce the details of the two datasets used to evaluate the hard question identification  
 949 performance. For the **MATH500** dataset (Lightman et al., 2023), we use the original difficulty labels,  
 950 which range from 1 to 5, with higher values indicating more difficult questions. Questions with a  
 951 difficulty level of 5 are considered hard, while the others are easy. The detailed statistics are shown in  
 952 Table 7.

953 For the **MixD** dataset, we combine questions from AIME25, AIME24, and AMC23 to form the  
 954 hard question set. Questions from the GSM8K dataset are considered easy. We random select  
 955 20% questions from GSM8K randomly selected to form the easy question set to avoid severe data  
 956 imbalance problem. We then mix both the hard and easy questions to create our MixD dataset. The  
 957 detailed statistics are shown in Table 7.

Dataset	Total Cnt	Number of Hard	Number of Easy
<b>MATH500</b>	500	134	366
<b>MixD</b>	363	100	263

958  
 959 Table 7: The statistics for our used test dataset used to evaluate hard question identification perfor-  
 960 mance.

961  
 962 

### D.2 BASELINES DETAILS

963 In this section, we provide a more detailed introduction to the baselines for hard question identification  
 964 as follows:

965  
 966 **Seq\_PPL** (Mahaut et al., 2024) uses sequence probability (PPL) to reflect the confidence of the LLM.  
 967 We compute the PPL of the LLM on each question, which is equivalent to the sequence probability

972		Question: Find the sum of all integer bases $b > 9$ for which $17_b$ is a divisor of $97_b$ .
973		Correct Answer: 70
974		<b>[Coconut (Hao et al. 2024):</b>
975		Response: To find the sum of all integer bases $b > 9$ for which $17_b$ is a divisor of $97_b$ , we start by converting the numbers from base $b$ to base 10. The number $17_b$ in base 10 is: $1 \cdot b + 7 \cdot b^0 = 9 \cdot b + 7 \cdot b^0 + 7$ . We need $b+7$ to divide $9b+7$ . This condition can be expressed as: $9b+7 \equiv 0 \pmod{b+7}$ for some integer $k$ . Now, rearranging and solving for $b$ , we get: $9b+7 \equiv kb+7 \pmod{b+7}$ or $(9-k)b \equiv 0 \pmod{b+7}$ . For $b$ to be an integer, the denominator $b+7$ must divide the numerator $7(k-1)$ . Therefore, we need to consider the divisors of 7, which are 1 and 7. We will test each of these values to find valid bases $b$ .
976		Case 1: $9-k=1$ . In this case, $k=8$ , and we can substitute into the equation for $b$ : $b=7(8-1)/1=7-7=1=49$ . Let's check if $49+7=56$ divides $9 \cdot 49 + 7 = 448$ : $448 \div 56 = 8$ . Since this is a valid division, $b=49$ is a valid base.
977		Case 2: $9-k=-1$ . In this case, $k=10$ , and we substitute into the equation for $b$ : $b=7(10-1)/-1=7 \cdot 9/-1=63$ . A negative base is invalid, so this case does not provide a valid solution.
978		Case 3: $9-k=7$ . In this case, $k=2$ , and we substitute into the equation for $b$ : $b=7(2-1)/7=1/7=1$ . Base 1 is invalid, so this case does not provide a valid solution.
979		Case 4: $9-k=-7$ . In this case, $k=16$ , and we substitute into the equation for $b$ : $b=7(16-1)/-7=7 \cdot 15/-7=-15$ . Again, a negative base is invalid, so this case does not provide a valid solution.
980		Conclusion: After considering all possible cases, the only valid base $b$ greater than 9 is $b=49$ . Therefore, the sum of all such bases is: 49.
981		Disadvantages: <b>Redundant and Incorrect</b>
982		<b>[CompressCoT (Cheng and Van Durme 2024):</b>
983		Response: Convert to base 10: $17_b$ in base 10 is $1 \cdot b + 7 \cdot b^0 = 9 \cdot b + 7 \cdot b^0 + 7$ . We need $b+7$ to divide $9b+7$ . This translates to the equation: $9b+7 \equiv 0 \pmod{b+7}$ . Then $9b+7 \equiv kb+7k$ , implies $(9-k)b \equiv 0 \pmod{b+7}$ . Thus, $b \equiv k \pmod{b+7}$ . We find that $k$ must be a divisor of 7, leading to possible values of $k$ as 1, 7, -1, -7. Checking these values: $k=1$ gives $b=7$ (invalid since $b>9$ ); $k=7$ gives $b=14$ ; $k=-1$ gives $b=-1$ (invalid); $k=-7$ gives $b=14$ (invalid). The only valid solution is $b=14$ .
984		Disadvantages: <b>Concise but Incorrect</b>
985		<b>[CODI (Shen et al. 2025b):</b>
986		Response: To find the sum of all integer bases such that divides, let's break this down carefully and take it step by step. We begin by converting the numbers from base $b$ to base 10. In base $b$ , $17_b$ is equal to $b+7$ , and $97_b$ is equal to $9b+7$ . This gives us the base-10 equivalents for both numbers. Now, the problem asks us to find when $17_b$ divides $97_b$ , which translates to the condition: $b+7 \mid 9b+7$ . At this stage, it's useful to simplify the divisibility condition. To do so, subtract $9(b+7)$ from $9b+7$ . Thus, the condition simplifies to $b+7 \mid 56$ . We can now focus on finding the divisors of 56 (as we are interested in positive divisors). The divisors of 56 are 1, 2, 4, 7, 8, 14, 28, 56. Next, we must satisfy the condition that $b+7 > 9$ , since the problem specifies that $b > 9$ . Therefore, we need to consider only the divisors of 56 that are greater than 9. These are 28 and 56, which correspond to $b=21$ and $b=49$ . Now, we simply add these two values of $21+49=70$ .
987		Let's take a moment to recheck and think about any potential alternative approaches. It seems that these two values of (21 and 49) are the only ones that satisfy the condition. Therefore, we can confidently say that the final sum is 70. To conclude, our careful analysis, combined with a little rethinking and checking of conditions, confirms that the sum of all integer bases where divides is indeed 70.
988		Disadvantages: <b>Redundant but Correct</b>
989		<b>[SynAdapt]:</b>
990		Response: We convert to base 10: $17_b = b+7$ and $97_b = 9b+7$ . The divisibility condition $b+7 \mid 9b+7$ reduces to $b+7 \mid 56$ by subtracting $9(b+7)$ . Considering $b > 9$ , the valid divisors are 28 and 56, yielding $b=21$ and $49$ . Their sum is 70.
991		Disadvantages: <b>Concise and Correct</b>

Figure 4: An representative example of the generated output from our *SynAdapt* and other CCoT-based baselines is provided. We highlight the crucial wrong steps that lead to incorrect answers in red, and the correct reasoning steps in green. Redundant parts in the answer are marked with a gray background. We also provide short analyses explaining the disadvantages or advantages of the generated responses.

of the question. We treat those questions with high PPL are considered hard, while those with low PPL are categorized as simple.

**PromptLLM** (Han et al., 2024) prompts the LLM to assess the difficulty of a question and predict the essential token budget required for solving it. We also prompt the LLM to predict the token budget and restrict the range to 128-32,768 tokens. Questions that require a high token budget are considered hard, while those with a low token budget are classified as simple.

**RouteLLM** (Ong et al., 2024) trains a hard question classifier using a BERT backbone. The classifier assigns high scores to hard questions and routes them to stronger LLMs, such as GPT-4 OpenAI (2025), while easier questions are processed by weaker LLMs, like Mixtral-8x7B Jiang et al. (2024). Therefore, we directly use their released model weights <sup>2</sup> and classify those questions with high scores as hard.

**Probe\_Q** (Azaria & Mitchell, 2023) trains a classifier based on the LLM’s hidden state to assess truthfulness. Similarly, we provide the LLM with the question and train a classifier to evaluate difficulty based on the last token’s hidden state from LLM. This approach is similar to ours, but it does not leverage information from the CCoT for assessing question difficulty.

## E CASE STUDIES

### E.1 RESPONSE EXAMPLE FROM VARIOUS BASELINES

We provide a representative example to demonstrate the effectiveness of our *SynAdapt* by comparing its generated response with those from other CCoT-based baselines, including **Coconut**, **CompressCoT**, and **CODI**.

As shown in Figure 4, the response from **Coconut** contains numerous redundant parts, which primarily serve communication or linguistic purposes, rather than contributing to the reasoning process needed to derive the correct answer. Moreover, the answer generated is incorrect, highlighting that indirect

<sup>2</sup>[https://huggingface.co/routellm/bert\\_gpt4\\_augmented](https://huggingface.co/routellm/bert_gpt4_augmented)

1026 training without explicit alignment with DCoT fails to effectively learn CCoT. **CompressCoT**  
 1027 successfully generates a concise response without redundancy but still outputs the wrong answer.  
 1028 This is because it aligns only with a subset of isolated, incoherent DCoT tokens, which fail to capture  
 1029 the full reasoning process, resulting in performance degradation. For **CODI**, the generated response  
 1030 provides the correct answer but retains redundant parts. This occurs because it applies alignment only  
 1031 at the final position, limiting its ability to learn CCoT and produce concise output.

1032 In contrast, our method generates both a concise and correct answer. This is due to our use of  
 1033 synthetic CCoT as the fine alignment target and applying full alignment during CCoT fine-tuning.  
 1034 These results strongly demonstrate the effectiveness of our method for efficient reasoning.  
 1035

## 1036 E.2 CCoT FOR HARD QUESTION EXAMPLE

1038 We provide an illustrative example  
 1039 demonstrating that solely relying on  
 1040 CCoT is insufficient to solve hard ques-  
 1041 tions. As shown in Figure 6, when the  
 1042 LLM relies only on CCoT, it generates  
 1043 a concise but incorrect answer. It may  
 1044 be because CCoT restricts the LLM’s  
 1045 ability to verify reasoning steps, confin-  
 1046 ing it to the incorrect answer. However,  
 1047 when prompted to re-think the question,  
 1048 the LLM can rectify the previous mis-  
 1049 take and derive the right answer. This  
 1050 effectively demonstrates that compress-  
 1051 ing DCoT into CCoT inevitably results  
 1052 in information loss, limiting the model’s  
 1053 reflective ability and leading to incorrect  
 1054 answers.

## 1055 E.3 INDISTINGUISHABLE HARD QUESTION EXAMPLE

1056 We provide an illustrative example  
 1057 showing how some hard questions are  
 1058 similar to simple ones, making them dif-  
 1059 ficult to distinguish. As seen in Figure  
 1060 6, both the easy and hard questions are  
 1061 very similar, both focusing on the quater-  
 1062 nions topic and are short in length. If  
 1063 we only assess difficulty based on the  
 1064 question itself, both would be categor-  
 1065 ized as easy, leading to performance  
 1066 degradation.

1067 However, when considering the CoT  
 1068 process, there exist significant differ-  
 1069 ences. For the easy question, the CoT  
 1070 is short and easily leads to the correct  
 1071 answer. In contrast, the hard question  
 1072 involves more reasoning steps and a  
 1073 longer CoT. By incorporating both the  
 1074 CoT and the question, we can accu-  
 1075 rately identify these indistinguishable  
 1076 hard questions. This highlights the value of reasoning information in identifying hard questions. This  
 1077 is also why our difficulty classifier is build up on both the question and CCoT, which can effectively  
 1078 utilize the information in CCoT.




Figure 5: An illustrative example of solving hard question relying solely on CCoT or rethinking via DCoT. We highlight the crucial wrong steps that lead to incorrect answers in red, and the correct reasoning steps in green.




Figure 6: An illustrative example of an easy question and the similar hard question, with their corresponding CoT processes. We also present the identification results using only the question or both the question and CoT. The key differences from the CoT of the hard question, compared to the easy question, are highlighted in red color.

## 1080 F MORE EVALUATION OF *SynAdapt*'S GENERALIZATION

1083 To further assess the generalization ability of our  
 1084 *SynAdapt*, we extend the evaluation beyond the mathematical domain to more various domains, including  
 1085 GPQA-Diamond (Rein et al., 2024) for scientific question answering, HumanEval (Chen et al.,  
 1086 2021) and LiveCodeBench (Naman Jain, 2024) for  
 1087 code generation. The results on GPQA-Diamond and  
 1088 HumanEval are shown in Table 4, while those on  
 1089 LiveCodeBench are presented in Table 8.

1091 As shown in Table 8, *SynAdapt* remains effective for  
 1092 the **LiveCodeBench** benchmark. Under the accuracy-  
 1093 sensitive setting with  $\tau = 0.5$ , our method achieves  
 1094 comparable performance to the raw model while sub-  
 1095 substantially reducing generation length, resulting in a Rel-G score of 1.14. Under the efficiency-sensitive  
 1096 setting with  $\tau = 1.0$ , it outperforms the other existing CCoT-based methods and achieves the best  
 1097 Rel-G score of 7.55. These results provide additional evidence that our method generalizes well to  
 1098 those coding tasks, rather than being limited to the mathematical domain.

## 1100 G HYPERPARAMETER ANALYSIS

1103 **The Length of CCoT  $m$ .** We analyze the hy-  
 1104 perparameter  $m$  in our method, which controls  
 1105 the length of the CCoT. As shown in Figure  
 1106 7(a), increasing  $m$  leads to higher accuracy as  
 1107 well as longer generation length. This is mainly  
 1108 because a longer CCoT contains more reason-  
 1109 ing steps, which benefits problem solving and  
 1110 improves accuracy. At the same time, longer  
 1111 CCoT also boosts the likelihood of model to  
 1112 generate redundant content, such as repeated  
 1113 verification steps, which simultaneously in-  
 1114 creases the generation length.

1115 To further measure the trade off performance  
 1116 between high accuracy and low generation length, we compute the Rel-G score to capture the actual  
 1117 performance gain, as shown in Equation 9. As illustrated in Figure 7(b), initially increasing the CCoT  
 1118 length  $m$  improves accuracy and leads to a higher Rel-G score. However, further increasing  $m$  causes  
 1119 the model to generate excessive redundant content, resulting in a decline in the Rel-G score. Overall,  
 1120 setting  $m = 512$  yields the best Rel-G score, indicating the optimal balance between accuracy and  
 1121 efficiency.

1123 **The Refining Iterations of CCoT  $k$ .** We also  
 1124 analyze the hyperparameter  $k$  in our method,  
 1125 which controls the refining iterations of CCoT.  
 1126 As shown in Figure 8(a), in the initial stage, in-  
 1127 creasing  $k$  allows the CCoTs to progressively  
 1128 refine potentially incorrect reasoning steps, en-  
 1129 abling the LLM to produce more accurate and  
 1130 concise answers. As a result, accuracy increases  
 1131 while the generation length decreases. However,  
 1132 when  $k$  is further increased beyond 4, redundant  
 1133 refinement steps may confuse the LLM,  
 leading to longer generation lengths and a slight  
 decrease in accuracy.

Methods	LiveCodeBench		
	Pass@1 $\uparrow$	Len $\downarrow$	Rel-G $\uparrow$
Raw Model	<b>46.4</b>	8642.8	1.00
<i>SynAdapt</i> ( $\tau=0.5$ )	41.1	<b>6689.8</b>	<b>1.14</b>
Coconut	26.6	900.3	5.50
CompressCoT	26.4	1323.4	3.72
CODI	25.4	689.0	6.87
<i>SynAdapt</i> ( $\tau=1.0$ )	<b>26.7</b>	<b>658.5</b>	<b>7.55</b>

Table 8: Evaluation results of our method and those CCoT-based methods on LiveCodeBench Naman Jain (2024) for code generation task.




Figure 7: The performance of our methods when using different CCoT Length  $m$ . We report the average results across five math benchmarks.




Figure 8: The performance of our methods when using different refining iterations  $k$  for CCoT generation. We report the average results across five math benchmarks.

1134 To further measure the accuracy-efficiency trade-off, we compute the Rel-G score, following Equation  
 1135 9. As shown in Figure 8(b), initially increasing  $k$  improves accuracy and reduces generation length,  
 1136 resulting in a higher Rel-G score. However, when  $k$  exceeds 4, the increase in generation length and  
 1137 slight decrease in accuracy lead to a lower Rel-G score. Therefore, setting  $k = 4$  achieves the best  
 1138 Rel-G score, indicating the optimal trade-off between accuracy and efficiency.  
 1139

## 1140 H USED PROMPT TEMPLATES

1141  
 1142 In this section, we present the prompts used in our method. For easy questions, we directly prompt  
 1143 the LLM to generate an answer based on the CCoT, as shown in Figure 9. For hard questions, we  
 1144 prompt the LLM to re-think and generate discrete CoT, condensing each reasoning step, as illustrated  
 1145 in Figure 10.  
 1146

```

1147 < | begin_of_sentence | >
1148 Please reason step by step, and put your final answer within \boxed{}.
1149
1150 < | User | >
1151 [[INSERT USER QUESTION HERE]]
1152
1153 < | Assistant | >
1154 <think> [[INSERT CCoT HERE]] </think>
1155

```

1156 Figure 9: The prompt used for directly generating answers based on the CCoT.  
 1157

```

1158 < | begin_of_sentence | >
1159 Think step by step, but only keep minimum draft for each thinking step,
1160 with 5 words at most.
1161 Return the answer at the end of the response within \boxed{}.
1162
1163 < | User | >
1164 [[INSERT USER QUESTION HERE]]
1165
1166 < | Assistant | >
1167 <think>
1168

```

1169 Figure 10: The prompt used for re-thinking hard questions via discrete CoT process while condensing  
 1170 each CoT step.  
 1171  
 1172  
 1173  
 1174  
 1175  
 1176  
 1177  
 1178  
 1179  
 1180  
 1181  
 1182  
 1183  
 1184  
 1185  
 1186  
 1187