

# 000 SIM-CoT: SUPERVISED IMPLICIT CHAIN-OF- 001 002 THOUGHT 003 004

005 **Anonymous authors**

006 Paper under double-blind review

## 007 008 ABSTRACT 009 010

011 Implicit Chain-of-Thought (CoT) methods offer a token-efficient alternative to  
012 explicit CoT reasoning in Large Language Models (LLMs), but a persistent per-  
013 formance gap has limited their adoption. We identify a core **latent instability**  
014 **issue** when scaling the computational budget of implicit CoT: as the number of  
015 reasoning tokens increases, training often becomes unstable and collapses. Our  
016 analysis shows that this instability arises from latent representations becoming  
017 homogeneous and losing semantic diversity, caused by insufficient step-level super-  
018 vision in current implicit CoT methods. To address this, we propose **SIM-CoT**, a  
019 plug-and-play training module that introduces step-level supervision to stabilize  
020 and enrich the latent reasoning space. SIM-CoT employs an auxiliary decoder  
021 during training to align each implicit token with its corresponding explicit reason-  
022 ing step, ensuring latent states capture distinct and meaningful information. The  
023 auxiliary decoder is removed at inference, preserving the efficiency of implicit  
024 CoT with no added overhead. It also provides interpretability by projecting each  
025 latent token onto an explicit reasoning vocabulary, enabling per-step visualization  
026 and diagnosis. SIM-CoT significantly improves both in-domain accuracy and  
027 out-of-domain stability of implicit CoT methods, boosting Coconut by +8.2% on  
028 GPT-2 and CODI by +3.0% on LLaMA-3.1 8B. It further surpasses the explicit  
029 CoT baseline on GPT-2 by 2.1% with 2.3 $\times$  greater token efficiency, while closing  
030 the performance gap on larger models like LLaMA-3.1 8B.  
031

## 032 1 INTRODUCTION 033

034 “Measure what is measurable, and make measurable what is not so.” — Galileo Galilei  
035

036 The strong reasoning capabilities of Large Language Models (LLMs) (OpenAI, 2024; Google, 2024;  
037 Anthropic, 2024) are often unlocked through explicit Chain-of-Thought (CoT) prompting (Wei et al.,  
038 2022). The explicit CoT approach enables LLMs to solve complex problems in a step-by-step  
039 manner, yielding high performance in domains like mathematics and programming (Guo et al., 2025;  
040 Muennighoff et al., 2025). Despite its advantages, explicit CoT also faces several limitations. For  
041 instance, explicit CoT approaches must verbalize intermediate thoughts from a fixed vocabulary,  
042 thereby precluding the exploration of alternative solution paths (Li et al., 2025; Zhang et al., 2025).  
043 Additionally, the generation of extensive intermediate sequences significantly increases inference  
044 cost and can result in redundant over-thinking steps or unnecessary verbosity (Chen et al., 2024).

045 To address the flexibility and efficiency issues of explicit CoT methods, recent **implicit CoT** ap-  
046 proaches (Hao et al., 2025; Zhang et al., 2025; Li et al., 2025) have been proposed by representing  
047 reasoning in a continuous latent space rather than as a sequence of discrete text tokens. The implicit  
048 CoT methods allow each latent representation to encode richer information than a single explicit  
049 token, often with a significantly smaller number of latents than the length of an explicit reasoning  
050 chain. Early representative implicit work like Coconut (Hao et al., 2025) improves efficiency while  
051 still capturing useful intermediate structure. More recent approaches, such as CODI (Shen et al.,  
052 2025), further apply trajectory-level distillation from explicit reasoning paths to enhance performance.  
053 Despite these advancements, a **performance gap** still exists between existing implicit CoT methods  
and their explicit counterparts. The implicit CoT approaches are *fast, token-efficient but less accurate*,  
which currently limits their broader application.

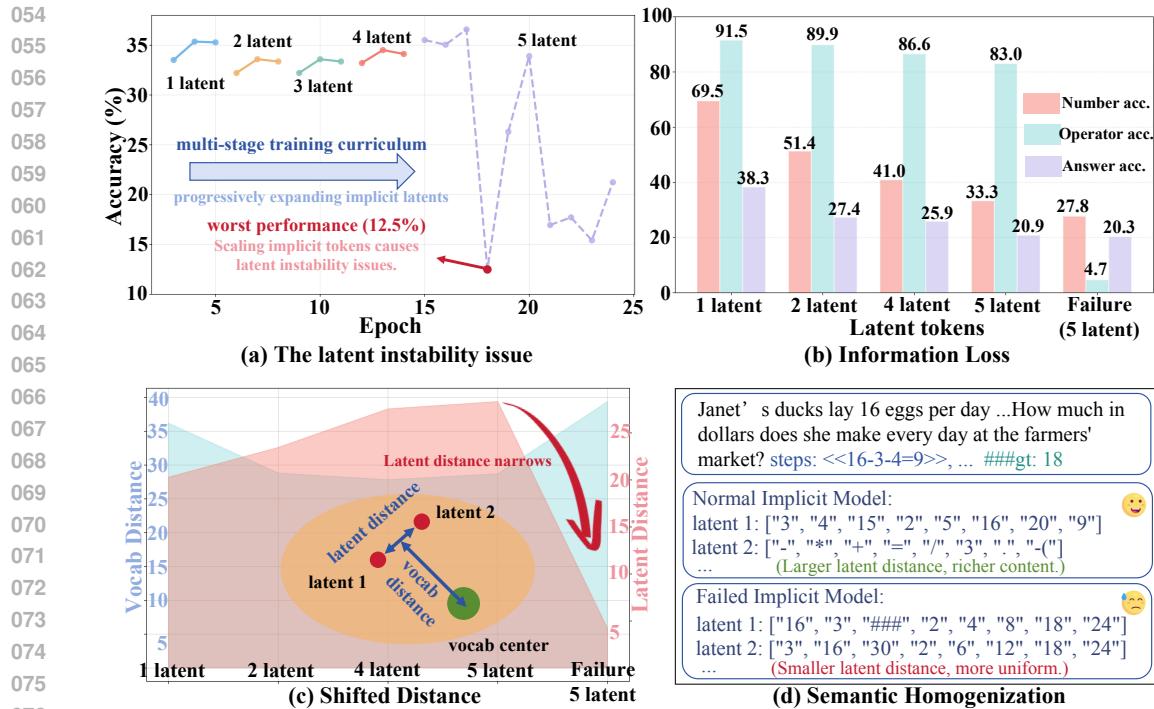


Figure 1: **(a) The latent instability issue:** while using more implicit tokens initially improves accuracy, training becomes unstable and sometimes collapses. **(b) Information Loss:** the implicit tokens of failed models (5 latent tokens) lose crucial information about operators (like  $+$ ,  $-$ ), which makes complex reasoning impossible. **(c) Shifted Distance:** the latent-to-latent distance of failed models shrinks and becomes too similar to each other, while the latent drifts away from the central vocabulary embedding space. **(d) Semantic Homogenization:** failed models produce similar latent representations, resulting in a narrower range of decoded tokens, mostly numbers, as opposed to the more varied content generated by a normal model.

To narrow the performance gap, inspired by the success of explicit CoT that scales computational budget for better performance, we explore a similar strategy for implicit CoT methods by increasing the number of implicit tokens. However, in Fig. 1 **(a)**, we reveal one underlying **latent instability issue** in current implicit CoT approaches. As we extend the number of implicit tokens from the default three (Hao et al., 2025) to five, the training process initially improves accuracy but becomes unstable and sometimes collapses entirely. To interpret the **latent instability issue**, we analyze implicit tokens from models trained on math reasoning data GSM8K-Aug (Deng et al., 2024). We follow previous works (Hao et al., 2025; Deng et al., 2024) to project the implicit tokens through the LM head and examine their top decoded tokens for analysis. As shown in Fig. 1 **(b)**, failed models tend to collapse into homogeneous latent states. While successful reasoning requires capturing both numerical and operator information, the implicit tokens of failed models primarily represent numbers, almost completely losing the critical operator information. Fig. 1 **(c)** further demonstrates that a model’s collapse is accompanied by two changes: a reduction in the inter-latent distance and a drift of the latent states away from the central vocabulary embedding space. The latent representations of failed models become too similar and lose their semantic connection to the tokens they are meant to represent. Fig. 1 **(d)** provides an example of the semantic homogenization. A normal model (top) maintains a large distance between its two latent tokens, allowing them to capture distinct information for numbers and operators. In contrast, a failed model (bottom)’s latent tokens become homogeneous, with both states decoding to similar information, primarily numbers.

Our observation (Fig. 1) reveals the reasons for the latent instability issue: a lack of sufficient step-level supervision for existing implicit methods to maintain the rich and varied internal representations. Without stronger guidance, the latent space collapses, losing its diversity and making it impossible to reliably encode the distinct, step-level reasoning needed for complex reasoning tasks. Motivated by our findings, we propose **Supervised IMplicit-CoT (SIM-CoT)**, a plug-and-play module that introduces step-level supervision for implicit CoT approaches to alleviate the latent instability issue.

108 Instead of supervising only the final answer (Hao et al., 2025) or the trajectory (Shen et al., 2025),  
 109 SIM-CoT uses an auxiliary decoder to align each implicit token with its corresponding explicit  
 110 reasoning step during training. The step-level supervision for implicit tokens stabilizes optimization,  
 111 prevents collapse, and ensures that latent tokens capture meaningful reasoning content. Crucially,  
 112 because the auxiliary decoder is removed during inference, our approach incurs virtually no extra  
 113 computational cost, making it as efficient as standard implicit CoT approaches. Beyond *accuracy*,  
 114 *stability*, and *efficiency*, the auxiliary decoder also affords *interpretability* of implicit reasoning.  
 115 During training, it defines a projection from latent tokens to the explicit reasoning vocabulary,  
 116 enabling us to decode each latent step into a human-interpretable summary for verification or error  
 117 diagnosis.

118 Experiments show that SIM-CoT acts as a plug-and-play module that boosts both accuracy and  
 119 stability. We show that SIM-CoT can be effortlessly combined with various implicit CoT approaches  
 120 such as Coconut (Hao et al., 2025), CODI (Shen et al., 2025), and training-free approaches (Zhang  
 121 et al., 2025) to further enhance reasoning performance. On GPT-2, SIM-CoT surpasses both the  
 122 strong explicit baseline (supervised fine-tuning on explicit CoT data) by 2.1%, and outperforms  
 123 existing implicit methods Coconut and CODI by 8.2% and 4.3%, respectively. The performance  
 124 trend holds as the method scales to larger models such as the LLaMA series. SIM-CoT achieves  
 125 improvements over CODI of 3.4% (LLaMA-3.2 1B), 1.5% (LLaMA-3.2 3B), and 3.0% (LLaMA-3.1  
 126 8B), in addition to a 9.0% gain over Coconut on the LLaMA-3.2 1B model. Furthermore, while  
 127 previous implicit CoT approaches (e.g., Coconut) collapse when scaled to 8 or 16 implicit tokens,  
 128 SIM-CoT remains stable and continues to boost performance.

129 In summary, our contributions are as follows: **1)** We provide a systematic analysis of the latent insta-  
 130 bility issue of implicit CoT approaches, showing that instability and collapse arise from insufficient  
 131 supervision. **2)** We introduce SIM-CoT, which applies step-level supervision to the model’s implicit  
 132 tokens. SIM-CoT not only integrates seamlessly with existing implicit CoT approaches and boosts  
 133 performance with minimal inference overhead, but also affords interpretability of implicit reasoning  
 134 by projecting each latent token onto an explicit reasoning vocabulary, enabling per-step visualization  
 135 of semantic roles and diagnosis. **3)** Through extensive experiments, we demonstrate that SIM-CoT  
 136 not only improves accuracy in the in-domain dataset, but also generalizes effectively to out-of-domain  
 137 datasets. The performance gains are consistent across a range of LLMs, including GPT-2 and recent  
 138 LLaMA 3 models (1B, 3B, and 8B).

## 2 ANALYSIS OF IMPLICIT COT: THE LATENT INSTABILITY ISSUE

142 We first present an analysis (Fig. 1) of the limitations in implicit latent CoT approaches. We follow  
 143 Coconut (Hao et al., 2025) and analyze implicit latents by projecting them through the LM head and  
 144 examining the top-8 decoded tokens to understand the semantic and geometric properties.

145 **Latent Instability Issue.** Fig. 1 **(a)** shows the training process of Coconut when the number of  
 146 implicit latent tokens is progressively increased. Initially, as the number of latents increases from one  
 147 to four, the model’s accuracy generally improves, suggesting that using more latents can enhance  
 148 performance. However, a significant drop in accuracy occurs when the number of latents is scaled to  
 149 five, with performance collapsing to its worst point of 12.5%. The latent instability issue indicates  
 150 that the implicit reasoning approach is sensitive to the choice of the number of latent tokens, as shown  
 151 by the sharp drop and subsequent fluctuations in accuracy after adding the fifth latent.

152 **Information Loss.** Fig. 1 **(b)** presents an analysis of how different levels of accuracy are affected by  
 153 the number of latent tokens, using accuracy metrics at three levels: number, operator, and answer. The  
 154 bar chart reveals a clear trend: as the number of latent tokens increases from 1 to 5, there is a general  
 155 decline in performance across all three metrics, especially for the operator accuracy. The strong  
 156 correlation between increased latent tokens and declining performance, particularly the sharp fall  
 157 during failure, suggests that implicit latents do not consistently capture the necessary compositional  
 158 reasoning process without more explicit, fine-grained supervision.

159 **Shifted Distance.** Fig. 1 **(c)** examines the geometric properties of the latent representations during  
 160 training. Two metrics are analyzed: the Latent Distance (red), which measures the average distance  
 161 between pairs of latent vectors, and the Vocab Distance (blue), which measures the average distance  
 from each latent vector to the center of the vocabulary embedding space. When the latent CoT model

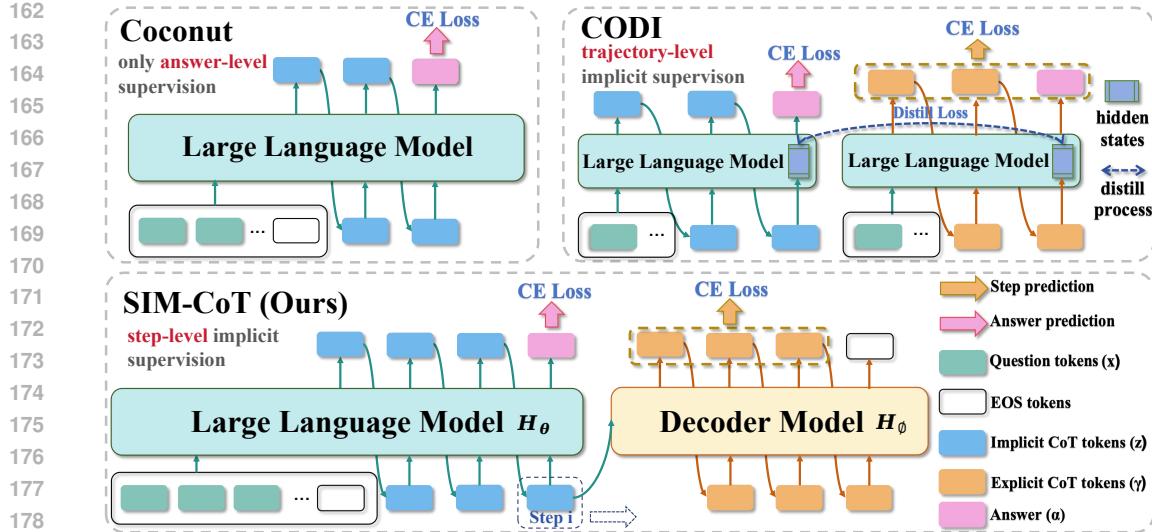


Figure 2: The framework comparison between **Coconut** (upper left), **CODI** (upper right), and our **SIM-CoT** (bottom). Unlike **Coconut** and **CODI**, which apply coarse-grained supervision on answers or trajectories, our **SIM-CoT** employs a decoder to **align implicit latents with step-level reasoning**, enhancing performance while maintaining inference efficiency.

collapses, the latent distance decreases sharply, indicating that the latent vectors are collapsing and becoming nearly identical, losing their distinctiveness. Simultaneously, the vocab distance increases, showing that these collapsing latents are drifting away from the main lexical embedding space and are no longer grounded in the fundamental token representations used by the model.

**Semantic Homogenization.** Fig. 1 (d) provides a qualitative analysis of the content of the latent tokens in a normal case versus a failed model. In the normal implicit model (middle), the decoded tokens from the latents are diverse and meaningful. In the failed implicit model (bottom), the semantic content of the latents becomes highly homogeneous. Latent 1 and Latent 2 contain mainly numbers, lacking operators or symbolic information needed for calculation. This shows that successful training produces latents with step-wise reasoning, while without explicit supervision, the latent space collapses into uniform numerical forms.

**Summary.** Our analysis across Fig. 1 (a-d) highlights a crucial trade-off between diversity and stability. When the model collapses, it loses both its diversity (as the latents become too similar) and its stability (as the latents move away from the token space), leading to catastrophic information loss and a complete failure of the reasoning process, as shown by the sharp drop in overall accuracy. These combined findings show that without proper guidance, the latent space degenerates, losing its ability to represent distinct reasoning steps. These challenges motivate our proposed method, which introduces **step-level implicit supervision** to stabilize the training process and enrich unique semantic content of each latent, all while maintaining efficiency during inference.

### 3 METHODOLOGY

**Overview.** As shown in Fig. 2, early implicit reasoning studies differ mainly in supervision granularity: **Coconut** (top left) uses answer-level supervision, while **CODI** (top right) introduces trajectory-level signals via distillation. Both remain coarse and do not tell the model which latent should encode which step. We propose **SIM-CoT**, which provides **step-level implicit supervision**: During an **implicit phase**, the LLM runs for a fixed number  $K$  of reasoning steps; at each step  $k$  it takes the **last hidden state** as the implicit latent  $z_k$  and appends it to the sequence as the next “token” vector. After  $K$  steps, the model switches back to **explicit** decoding over the vocabulary to generate the final answer. A decoder is used only in training to align each  $z_k$  with the textual content of the  $k$ -th reasoning step; at inference, the decoder is removed, so the runtime is essentially that of direct answer generation plus  $K$  forward positions, which is far shorter than explicit CoT token lengths.

216 3.1 NOTATION  
217218 Let  $\mathcal{V}$  be the vocabulary and  $E \in \mathbb{R}^{|\mathcal{V}| \times d}$  the token embedding matrix. A question is  $x =$   
219  $(x_1, \dots, x_T) \in \mathcal{V}^T$  with embedded prefix

220 
$$U^{(0)} = (e(x_1), \dots, e(x_T)), \quad e(\cdot) \in \mathbb{R}^d.$$
  
221

222 We run an autoregressive LLM  $F_\theta$  on any prefix  $U = (u_1, \dots, u_m)$  of  $d$ -dimensional vectors (tokens  
223 or latents). Denote the last-layer hidden state at the final position by

224 
$$H_\theta(U) \in \mathbb{R}^d.$$
  
225

226 For supervision, the  $k$ -th textual step is  $s_k = (y_{k,1}, \dots, y_{k,L_k}) \in \mathcal{V}^{L_k}$ , and the answer is  $a =$   
227  $(a_1, \dots, a_{L_a}) \in \mathcal{V}^{L_a}$ . The auxiliary decoder has parameters  $\phi$ ; the LLM has parameters  $\theta$ .228 3.2 IMPLICIT PHASE: LATENT CONSTRUCTION BY LAST HIDDEN STATES  
229230 We fix the number of implicit reasoning steps  $K$  in advance. For each step  $k = 1, \dots, K$ ,  
231

232 
$$z_k = H_\theta(U^{(k-1)}) \in \mathbb{R}^d, \quad U^{(k)} = U^{(k-1)} \oplus z_k, \quad (1)$$
  
233

234 where  $\oplus$  denotes concatenation along the time axis. The implicit chain-of-thought is therefore represented  
235 as a continuous sequence of hidden states  $z_{1:K} = (z_1, \dots, z_K)$ , which are autoregressively  
236 generated and appended to the context before the model switches to explicit decoding.237 3.3 EXPLICIT PHASE: ANSWER DECODING OVER THE VOCABULARY  
238239 After constructing the implicit latents  $z_{1:K}$ , the model switches to explicit decoding to generate the  
240 final answer. Let  $W_o \in \mathbb{R}^{|\mathcal{V}| \times d}$  be the output projection (LM head). With teacher forcing on the  
241 partial answer  $a_{<t}$ , the generation is  
242

243 
$$h_{T+K+t} = H_\theta(U^{(K)} \oplus e(a_{<t})), \quad (2)$$
  
244

245 
$$p_\theta(a_t | x, z_{1:K}, a_{<t}) = \text{softmax}(W_o h_{T+K+t})_{a_t}, \quad (3)$$
  
246

247 
$$p_\theta(a | x, z_{1:K}) = \prod_{t=1}^{L_a} p_\theta(a_t | x, z_{1:K}, a_{<t}). \quad (4)$$
  
248

249 3.4 TRAINING-TIME DECODER AND STEP-LEVEL SUPERVISION  
250251 During training, a decoder  $p_\phi$  (architecturally identical to the LLM) takes only the  $k$ -th implicit latent  
252  $z_k$  as conditioning signal and autoregressively generates the  $k$ -th textual step  $s_k = (y_{k,1}, \dots, y_{k,L_k})$ .  
253 This provides **step-level** supervision that directly grounds  $z_k$  to its corresponding reasoning content:

254 
$$p_\phi(s_{1:K} | z_{1:K}) = \prod_{k=1}^K p_\phi(s_k | z_k) = \prod_{k=1}^K \prod_{t=1}^{L_k} p_\phi(y_{k,t} | z_k, y_{k,<t}). \quad (5)$$
  
255

256 *Parameterization.* For step  $k$ , the decoder is conditioned on the implicit latent  $z_k$  obtained from the  
257 LLM. Since  $z_k$  does not correspond to any token in the vocabulary, it is not included in the loss  
258 calculation. Instead,  $z_k$  is injected as an additional prefix vector that initializes the decoder’s hidden  
259 state for step generation. Concretely, the decoder input sequence is

260 
$$U_k^{\text{dec}} = [z_k; e(y_{k,1}), \dots, e(y_{k,L_k})],$$
  
261

262 where  $e(\cdot)$  denotes the embedding function of the LLM shared between both models. During training  
263 with teacher forcing, the decoder predicts each token  $y_{k,t}$  autoregressively:

264 
$$p_\phi(y_{k,t} | z_k, y_{k,<t}) = \text{softmax}(W^{\text{dec}} h_{k,t}^{\text{dec}})_{y_{k,t}},$$
  
265

266 where  $h_{k,t}^{\text{dec}}$  is the decoder hidden state at position  $t$  and  $W^{\text{dec}}$  is the LM head of the decoder.

270 The training loss for step  $k$  is then  
 271

$$272 \quad \mathcal{L}_{\text{step},k} = - \sum_{t=1}^{L_k} \log p_\phi(y_{k,t} \mid z_k, y_{k,<t}),$$

274 which supervises only the textual step tokens. The decoder is used exclusively for this supervision  
 275 during training and is discarded at inference.  
 276

277 **3.5 OBJECTIVES**  
 278

279 Training involves two complementary cross-entropy losses: one for supervising the textual steps  
 280 through the decoder, and one for supervising the final answer through the base LLM.  
 281

282 **Step-level supervision.** For each implicit latent  $z_k$ , the decoder  $p_\phi$  generates the corresponding  
 283 reasoning step  $s_k = (y_{k,1}, \dots, y_{k,L_k})$ . Since  $z_k$  is not a vocabulary token, the loss is computed only  
 284 over the textual step tokens:

$$285 \quad \mathcal{L}_{\text{step}} = - \sum_{k=1}^K \sum_{t=1}^{L_k} \log p_\phi(y_{k,t} \mid z_k, y_{k,<t}). \quad (6)$$

287 This loss grounds each latent  $z_k$  to a specific reasoning step, ensuring that the latent sequence carries  
 288 fine-grained semantics.  
 289

290 **Answer supervision.** After  $K$  implicit steps, the LLM  $F_\theta$  switches back to explicit decoding to  
 291 generate the final answer  $a = (a_1, \dots, a_{L_a})$ . We optimize the standard language modeling loss:  
 292

$$293 \quad \mathcal{L}_{\text{ans-lm}} = - \sum_{t=1}^{L_a} \log p_\theta(a_t \mid x, z_{1:K}, a_{<t}). \quad (7)$$

295 **Total objective.** The overall loss is a weighted sum:  
 296

$$297 \quad \mathcal{L} = \lambda_{\text{step}} \mathcal{L}_{\text{step}} + \lambda_{\text{lm}} \mathcal{L}_{\text{ans-lm}}. \quad (8)$$

298 Gradients from  $\mathcal{L}_{\text{step}}$  propagate through the decoder into the latent representations  $z_{1:K}$  and further  
 299 into the LLM (via Eq. equation 1), shaping the hidden states to encode step-level reasoning.  
 300 Meanwhile,  $\mathcal{L}_{\text{ans-lm}}$  trains the base model to produce the final answer directly, so the decoder can be  
 301 discarded at inference time without affecting efficiency. Implementation details, inference procedures,  
 302 and diagnostic analyses are provided in Appendix D.  
 303

304 **4 EXPERIMENT**

305 **4.1 EXPERIMENTAL SETUP**  
 306

307 **Training Data.** We follow previous works (Deng et al., 2024; Hao et al., 2025) to use the  
 308 **GSM8k-Aug** dataset Deng et al. (2024) for training implicit CoT models. The GSM8k-Aug expands  
 309 the original GSM8k training set (Cobbe et al., 2021) to 385k examples by using GPT-4  
 310 for data generation. To facilitate implicit CoT training, the GSM8k-Aug removes the reasoning  
 311 chain of natural language, preserving only a sequence of structured mathematical expressions.  
 312 Each expression is logically linked to the previous step, as illustrated by the example:  
 313 <<12 \* 3=36>><<9 \* 2=18>><<17 \* 2=34>><<36+18+34=88>>.

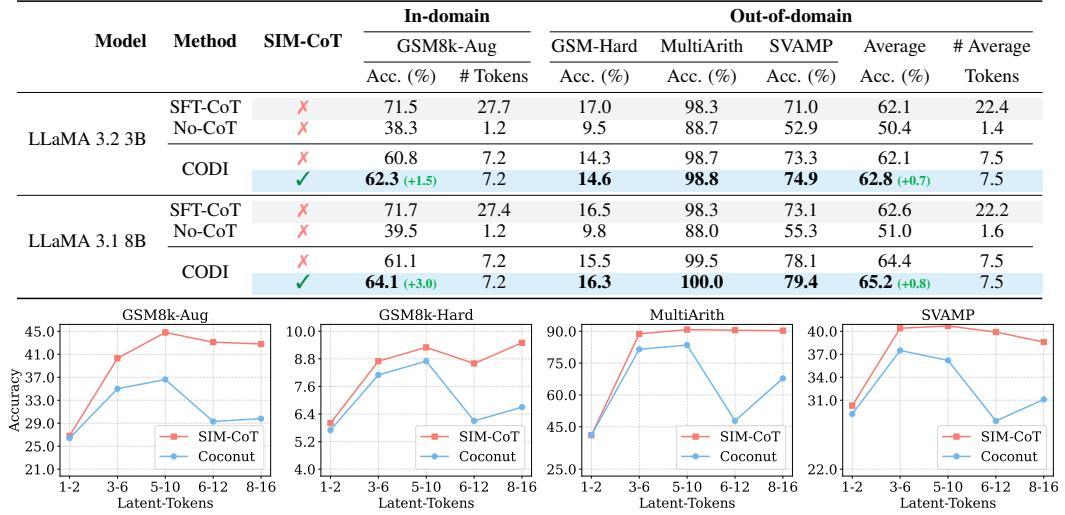
314 **Evaluation Benchmarks.** We report results on the **GSM8k-Aug** test set (Cobbe et al., 2021), which  
 315 serves as our in-domain (ID) evaluation benchmark. To further evaluate mathematical reasoning  
 316 under a distribution shift, we also evaluate models on three out-of-domain (OOD) benchmarks: (1)  
 317 **SVAMP** (Patel et al., 2021), a dataset of grade-school arithmetic word problems that introduces  
 318 simple variations to assess robustness; (2) **GSM-Hard** (Gao et al., 2022), a modified version of the  
 319 GSM8k test split where numbers are replaced with larger magnitudes to increase problem difficulty;  
 320 and (3) **MultiArith** (Roy & Roth, 2015), a subset of MAWPS (Koncel-Kedziorski et al., 2016)  
 321 consisting of multi-step arithmetic word problems. Please refer to the Appendix D for more details.  
 322

323 **Implementation Details.** We follow the training setup of previous works (Hao et al., 2025; Shen  
 324 et al., 2025), and adopt consistent hyperparameter choices for GPT-2, LLaMA 1B/3B/8B. Detailed  
 325 configurations, such as learning rates, curriculum strategies, are provided in Appendix C.  
 326

324  
 325  
 326  
 327  
 328 Table 1: **Main results on GPT-2.** We report accuracy (%) on *in-domain* (GSM8k-Aug) and *out-of-domain* (GSM-Hard, MultiArith, SVAMP) benchmarks. Our SIM-CoT is shown to provide accuracy  
 329 gains on top of existing methods such as Coconut (Hao et al., 2025) and CODI (Shen et al., 2025).

330   331   332   333   334   335   336   Method	337   338   339   SIM-CoT	340   341   342   343   344   345   346   347   348   In-domain		349   350   351   352   353   354   355   356   357   358   359   360   361   362   363   364   365   366   367   368   369   370   371   372   373   374   375   376   377   378   379   380   381   382   383   384   385   386   387   388   389   390   391   392   393   394   395   396   397   398   399   400   401   402   403   404   405   406   407   408   409   410   411   412   413   414   415   416   417   418   419   420   421   422   423   424   425   426   427   428   429   430   431   432   433   434   435   436   437   438   439   440   441   442   443   444   445   446   447   448   449   450   451   452   453   454   455   456   457   458   459   460   461   462   463   464   465   466   467   468   469   470   471   472   473   474   475   476   477   478   479   480   481   482   483   484   485   486   487   488   489   490   491   492   493   494   495   496   497   498   499   500   501   502   503   504   505   506   507   508   509   510   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   541   542   543   544   545   546   547   548   549   550   551   552   553   554   555   556   557   558   559   560   561   562   563   564   565   566   567   568   569   570   571   572   573   574   575   576   577   578   579   580   581   582   583   584   585   586   587   588   589   590   591   592   593   594   595   596   597   598   599   600   601   602   603   604   605   606   607   608   609   610   611   612   613   614   615   616   617   618   619   620   621   622   623   624   625   626   627   628   629   630   631   632   633   634   635   636   637   638   639   640   641   642   643   644   645   646   647   648   649   650   651   652   653   654   655   656   657   658   659   660   661   662   663   664   665   666   667   668   669   670   671   672   673   674   675   676   677   678   679   680   681   682   683   684   685   686   687   688   689   690   691   692   693   694   695   696   697   698   699   700   701   702   703   704   705   706   707   708   709   710   711   712   713   714   715   716   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   757   758   759   760   761   762   763   764   765   766   767   768   769   770   771   772   773   774   775   776   777   778   779   780   781   782   783   784   785   786   787   788   789   790   791   792   793   794   795   796   797   798   799   800   801   802   803   804   805   806   807   808   809   810   811   812   813   814   815   816   817   818   819   820   821   822   823   824   825   826   827   828   829   830   831   832   833   834   835   836   837   838   839   840   841   842   843   844   845   846   847   848   849   850   851   852   853   854   855   856   857   858   859   860   861   862   863   864   865   866   867   868   869   870   871   872   873   874   875   876   877   878   879   880   881   882   883   884   885   886   887   888   889   890   891   892   893   894   895   896   897   898   899   900   901   902   903   904   905   906   907   908   909   910   911   912   913   914   915   916   917   918   919   920   921   922   923   924   925   926   927   928   929   930   931   932   933   934   935   936   937   938   939   940   941   942   943   944   945   946   947   948   949   950   951   952   953   954   955   956   957   958   959   960   961   962   963   964   965   966   967   968   969   970   971   972   973   974   975   976   977   978   979   980   981   982   983   984   985   986   987   988   989   990   991   992   993   994   995   996   997   998   999   1000   1001   1002   1003   1004   1005   1006   1007   1008   1009   1010   1011   1012   1013   1014   1015   1016   1017   1018   1019   1020   1021   1022   1023   1024   1025   1026   1027   1028   1029   1030   1031   1032   1033   1034   1035   1036   1037   1038   1039   1040   1041   1042   1043   1044   1045   1046   1047   1048   1049   1050   1051   1052   1053   1054   1055   1056   1057   1058   1059   1060   1061   1062   1063   1064   1065   1066   1067   1068   1069   1070   1071   1072   1073   1074   1075   1076   1077   1078   1079   1080   1081   1082   1083   1084   1085   1086   1087   1088   1089   1090   1091   1092   1093   1094   1095   1096   1097   1098   1099   1100   1101   1102   1103   1104   1105   1106   1107   1108   1109   1110   1111   1112   1113   1114   1115   1116   1117   1118   1119   1120   1121   1122   1123   1124   1125   1126   1127   1128   1129   1130   1131   1132   1133   1134   1135   1136   1137   1138   1139   1140   1141   1142   1143   1144   1145   1146   1147   1148   1149   1150   1151   1152   1153   1154   1155   1156   1157   1158   1159   1160   1161   1162   1163   1164   1165   1166   1167   1168   1169   1170   1171   1172   1173   1174   1175   1176   1177   1178   1179   1180   1181   1182   1183   1184   1185   1186   1187   1188   1189   1190   1191   1192   1193   1194   1195   1196   1197   1198   1199   1200   1201   1202   1203   1204   1205   1206   1207   1208   1209   1210   1211   1212   1213   1214   1215   1216   1217   1218   1219   1220   1221   1222   1223   1224   1225   1226   1227   1228   1229   1230   1231   1232   1233   1234   1235   1236   1237   1238   1239   1240   1241   1242   1243   1244   1245   1246   1247   1248   1249   1250   1251   1252   1253   1254   1255   1256   1257   1258   1259   1260   1261   1262   1263   1264   1265   1266   1267   1268   1269   1270   1271   1272   1273   1274   1275   1276   1277   1278   1279   1280   1281   1282   1283   1284   1285   1286   1287   1288   1289   1290   1291   1292   1293   1294   1295   1296   1297   1298   1299   1300   1301   1302   1303   1304   1305   1306   1307   1308   1309   1310   1311   1312   1313   1314   1315   1316   1317   1318   1319   1320   1321   1322   1323   1324   1325   1326   1327   1328   1329   1330   1331   1332   1333   1334   1335   1336   1337   1338   1339   1340   1341   1342   1343   1344   1345   1346   1347   1348   1349   1350   1351   1352   1353   1354   1355   1356   1357   1358   1359   1360   1361   1362   1363   1364   1365   1366   1367   1368   1369   1370   1371   1372   1373   1374   1375   1376   1377   1378   1379   1380   1381   1382   1383   1384   1385   1386   1387   1388   1389   1390   1391   1392   1393   1394   1395   1396   1397   1398   1399   1400   1401   1402   1403   1404   1405   1406   1407   1408   1409   1410   1411   1412   1413   1414   1415   1416   1417   1418   1419   1420   1421   1422   1423   1424   1425   1426   1427   1428   1429   1430   1431   1432   1433   1434   1435   1436   1437   1438   1439   1440   1441   1442   1443   1444   1445   1446   1447   1448   1449   1450   1451   1452   1453   1454   1455   1456   1457   1458   1459   1460   1461   1462   1463   1464   1465   1466   1467   1468   1469   1470   1471   1472   1473   1474   1475   1476   1477   1478   1479   1480   1481   1482   1483   1484   1485   1486   1487   1488   1489   1490   1491   1492   1493   1494   1495   1496   1497   1498   1499   1500   1501   1502   1503   1504   1505   1506   1507   1508   1509   1510   1511   1512   1513   1514   1515   1516   1517   1518   1519   1520   1521   1522   1523   1524   1525   1526   1527   1528   1529   1530   1531   1532   1533   1534   1535   1536   1537   1538   1539   1540   1541   1542   1543   1544   1545   1546   1547   1548   1549   1550   1551   1552   1553   1554   1555   1556   1557   1558   1559   1560   1561   1562   1563   1564   1565   1566   1567   1568   1569   1570   1571   1572   1573   1574   1575   1576   1577   1578   1579   1580   1581   1582   1583   1584   1585   1586   1587   1588   1589   1590   1591   1592   1593   1594   1595   1596   1597   1598   1599   1600   1601   1602   1603   1604   1605   1606   1607   1608   1609   1610   1611   1612   1613   1614   1615   1616   1617   1618   1619   1620   1621   1622   1623   1624   1625   1626   1627   1628   1629   1630   1631   1632   1633   1634   1635   1636   1637   1638   1639   1640   1641   1642   1643   1644   1645   1646   1647   1648   1649   1650   1651   1652   1653   1654   1655   1656   1657   1658   1659   1660   1661   1662   1663   1664   1665   1666   1667   1668   1669   1670   1671   1672   1673   1674   1675   1676   1677   1678   1679   1680   1681   1682   1683   1684   1685   1686   1687   1688   1689   1690   1691   1692   1693   1694   1695   1696   1697   1698   1699   1700   1701   1702   1703   1704   1705   1706   1707   1708   1709   1710   1711   1712   1713   1714   1715   1716   1717   1718   1719   1720   1721   1722   1723   1724   1725   1726   1727   1728   1729   1730   1731   1732   1733   1734   1735   1736   1737   1738   1739   1740   1741   1742   1743   1744   1745   1746   1747   1748   1749   1750   1751   1752   1753   1754   1755   1756   1757   1758   1759   1760   1761   1762   1763   1764   1765   1766   1767   1768   1769   1770   1771   1772   1773   1774   1775   1776   1777   1778   1779   1780   1781   1782   1783   1784   1785   1786   1787   1788   1789   1790   1791   1792   1793   1794   1795   1796   1797   1798   1799   1800   1801				

378 Table 3: Main results on larger LLaMA models (3B and 8B). We report accuracy (%) on **in-domain**  
 379 (GSM8k-Aug) and **out-of-domain** (GSM-Hard, MultiArith, SVAMP) benchmarks.



390 Figure 3: Ablation study on different numbers of implicit latents. The x-axis denotes the number of  
 391 implicit latents and implicit tokens (joined with “-”), while the y-axis denotes accuracy. The blue line  
 392 corresponds to our method SIM-CoT, and the orange line corresponds to the baseline Coconut.

401 average improvement of +4.3 points when using Coconut as the backbone. From the third column of  
 402 Table 2, our method further improves upon the current SOTA implicit reasoning method CODI by  
 403 +1.0 point. Moreover, when scaling model size from GPT-2 to LLaMA-1B, SIM-CoT enlarges the  
 404 performance gap against iCoT, Coconut, and other baselines.

405 We attribute the robustness of SIM-CoT to its step-level implicit supervision. Unlike SFT-CoT,  
 406 which forces the model to mimic deterministic natural language annotations, and unlike CODI,  
 407 which applies trajectory-level alignment to a coarse-grained reasoning path, our method introduces  
 408 a moderate form of supervision. This design ensures the plausibility of each reasoning step while  
 409 preserving the diversity of reasoning trajectories, thereby improving generalization to unseen inputs.

410 **Inference Efficiency.** In terms of inference speed, our method maintains the same efficiency as  
 411 other implicit reasoning approaches on both GPT-2 and LLaMA-1B. On GPT-2, SIM-CoT not only  
 412 surpasses SFT-CoT on both in-domain and out-of-domain benchmarks, but also achieves a  $2.3 \times$   
 413 and  $2.2 \times$  speedup on Coconut, respectively. On LLaMA-1B, SIM-CoT remains comparable to  
 414 SFT-CoT in accuracy while delivering  $1.9 \times$  and  $1.7 \times$  speedups on in-domain and out-of-domain  
 415 benchmarks, respectively. These results demonstrate the effectiveness of our approach in retaining or  
 416 even enhancing the performance of explicit CoT while substantially reducing inference cost.

### 417 4.3 ABLATION STUDIES

418 **Ablation on the Number of Implicit Tokens.** We study the effect of varying the number of  
 419 implicit latents on GPT-2, comparing SIM-CoT with Coconut trained on GSM8k-Aug and evaluated  
 420 on GSM8k-Aug, GSM-Hard, MultiArith, and SVAMP (Fig. 3). Following Coconut, each latent  
 421 corresponds to two tokens. As shown in Fig. 5, most problems involve two to six steps with a  
 422 small proportion of harder cases, so we set the maximum number of implicit latents to 8. For each  
 423 configuration, we report the best performance, and results show that SIM-CoT provides more stable  
 424 training and achieves consistent gains over Coconut, indicating that step-level implicit supervision  
 425 scales effectively with larger latent capacity.

426 **Ablation on Scaling to Larger Backbones.** To examine robustness and scalability, we extend experiments  
 427 to larger LLaMA backbones, including LLaMA 3.2 3B and LLaMA 3.1 8B. Table 3 reports  
 428 results on GSM8k-Aug (in-domain) and GSM-Hard, MultiArith, and SVAMP (out-of-domain).

429 Overall, SIM-CoT scales effectively to larger backbones, consistently surpassing or matching explicit  
 430 CoT on out-of-domain tasks while reducing reliance on trajectory-level supervision.

432 Table 4: Comparison of (a) LLaMA 1B with different decoders and (b) latent token distance analysis.  
 433 In (a), we evaluate the effect of using larger decoders with a 1B model on both in-domain (GSM8k-  
 434 Aug) and out-of-domain benchmarks (GSM-Hard, MultiArith, SVAMP). In (b), we report average  
 435 pairwise distances among latent tokens (Dist.) and their distances to the vocabulary center (Dist. to  
 436 VC) under different settings, including failed cases and the effect after applying SIM-CoT.

437 (a) LLaMA 1B with different decoders.

438

Model	In-domain		Out-of-domain	
	GSM8k-Aug	GSM-Hard	MultiArith	SVAMP
Baseline	52.7	11.9	95.0	60.6
+ 1B Decoder	<b>56.1</b>	<b>12.7</b>	<b>96.2</b>	<b>61.5</b>
+ 3B Decoder	50.4	11.6	95.6	59.8
+ 8B Decoder	50.0	11.7	94.2	56.8

439 (b) Latent token distance analysis.

440

Setting	Dist.	Dist. to VC
1 latent	20.30	36.20
2 latent	23.46	28.82
4 latent	27.56	27.83
5 latent	28.34	28.34
Fail 5 latent	4.21	39.39
After SIM-CoT	32.81	29.80

441 Table 5: Ablation study of soft thinking on LLaMA 3.2 1B. We report accuracy (%) on the in-domain  
 442 dataset (GSM8k-Aug) and out-of-domain datasets (GSM-Hard, MultiArith, and SVAMP). Adding  
 443 soft thinking consistently improves both Coconut and SIM-CoT across all benchmarks, showing its  
 444 effectiveness in enhancing implicit reasoning.

445

Method	GSM8k-Aug	GSM-Hard	MultiArith	SVAMP
Coconut	36.6	8.1	83.5	36.2
+ Soft Thinking	36.7	8.3	85.2	36.0
SIM-CoT	44.8	9.3	90.8	40.7
+ Soft Thinking	<b>45.0</b>	<b>9.4</b>	<b>91.5</b>	<b>40.8</b>

446 On **LLaMA 3.2 3B**, SIM-CoT improves over CODI by +1.5 points on GSM8k-Aug and +1.6  
 447 points on SVAMP, while maintaining comparable performance on GSM-Hard and MultiArith. This  
 448 demonstrates that step-level implicit supervision strengthens strong implicit reasoning baselines even  
 449 at larger scales.

450 On **LLaMA 3.1 8B**, SIM-CoT yields gains of +3.0 points on GSM8k-Aug, +1.3 on SVAMP, and  
 451 +0.8 on MultiArith relative to CODI, while maintaining stable accuracy on GSM-Hard. Compared  
 452 with SFT-CoT, it achieves higher accuracy on MultiArith (100.0 vs. 98.3) and SVAMP (79.4 vs.  
 453 73.1), while remaining similar on GSM-Hard.

454 Together, these results confirm that SIM-CoT scales effectively to larger backbones, providing  
 455 consistent gains across both in-domain and out-of-domain benchmarks with reduced reliance on  
 456 trajectory-level supervision.

457 **Ablation on Different Decoder Sizes.** We investigate how decoder size affects performance by  
 458 replacing the decoder of the LLaMA 1B backbone with larger versions from the same vocabulary  
 459 family and evaluating on GSM8k-Aug, GSM-Hard, MultiArith, and SVAMP. As shown in Table 4(a),  
 460 integrating a 1B-scale decoder leads to consistent improvements across all benchmarks. However,  
 461 simply scaling the decoder to larger variants (3B or 8B) does not yield additional benefits and instead  
 462 slightly reduces accuracy.

463 These results suggest that moderate decoder scaling can enhance reasoning ability, but excessively  
 464 large decoders may introduce optimization challenges or misalignment with the 1B backbone,  
 465 ultimately limiting generalization. A plausible explanation is that the 1B encoder and 1B decoder  
 466 originate from the same model family and thus share a more compatible representation space,  
 467 facilitating stable learning. In contrast, larger decoders (3B or 8B) may require implicit projection to  
 468 align with the 1B backbone, which can introduce representational mismatches and hinder training  
 469 stability.

470 **Ablation on Soft Thinking.** We also study the effect of integrating soft thinking (Zhang et al., 2025;  
 471 Wu et al., 2025) with both Coconut and SIM-CoT. For clarity, the detailed experimental setup, results,  
 472 and analyses are provided in Appendix A.

473 **Interpretability of Implicit Reasoning.** Implicit reasoning models generate continuous latent  
 474 thoughts that do not correspond to vocabulary tokens and thus cannot be directly decoded into

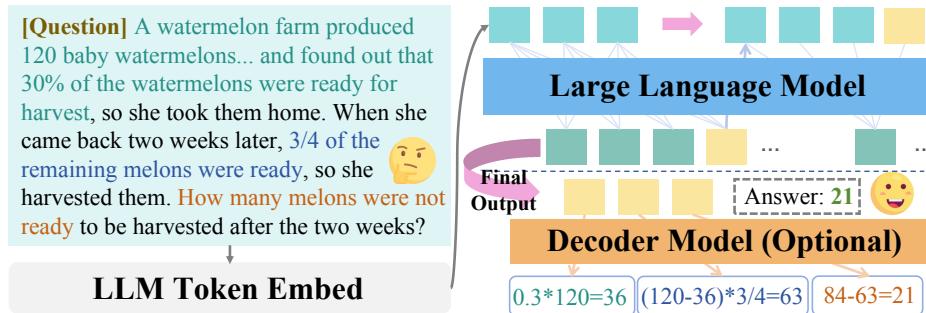


Figure 4: SIM-CoT case study on GSM8k. The generated implicit continuous tokens are subsequently interpreted by our decoder, which visualizes the solution intermediate steps leading to the final output.

human-readable text. To address this, we reuse the training decoder to project each latent step into interpretable textual space, enabling per-step visualization of latent semantics (Fig. 4). For completeness, detailed descriptions and additional visualization examples are provided in Appendix G.

To better understand the latent space, we analyze two geometric measures: the average pairwise distance among latent tokens (Dist.) and their distance to the vocabulary center (Dist. to VC), summarized in Table 4(b) and Fig. 6. As the number of latent tokens increases from 1 to 5, Dist. rises from 20.30 to 28.34, indicating improved separability. In the failed 5-latent case, however, Dist. collapses to 4.21, showing that tokens converge to nearly identical points; SIM-CoT avoids this collapse and increases Dist. to 32.81. For Dist. to VC, values decrease from 36.20 (1 latent) to around 28 as more tokens are introduced, reflecting better alignment with the vocabulary manifold. The failed 5-latent case instead spikes to 39.39, indicating drift, whereas SIM-CoT stabilizes this measure at 29.80. Fig. 6 qualitatively confirms these patterns: normal tokens remain separated and grounded, failed tokens collapse and drift outward, and SIM-CoT restores a structured configuration. Overall, these analyses show that SIM-CoT improves stability while maintaining an interpretable latent space.

## 5 RELATED WORK

A large body of work has studied explicit chain-of-thought (CoT) prompting, including self-consistency (Wei et al., 2022; Wang et al., 2023), least-to-most prompting (Zhou et al., 2023), reflection-based reasoning (Shinn et al., 2023; Madaan et al., 2023), and the integration of external tools (Yao et al., 2023). Other work investigates step-level supervision to structure explicit reasoning (Zheng et al., 2023; Wei et al., 2025). While effective, explicit CoT increases inference cost with longer sequences and often produces redundant steps, limiting efficiency and reasoning diversity (Li et al., 2025; Zhang et al., 2025; Xu et al., 2025).

Implicit CoT aims to reduce output length while retaining multi-step reasoning. Prior work explores knowledge internalization (Deng et al., 2024), architectural modification (Saunshi et al., 2025; Chen et al., 2025; Cheng & Van Durme, 2024; Su et al., 2025; Mohtashami et al., 2023; Geiping et al., 2025), training-free latent construction (Zhang et al., 2025; Wu et al., 2025), and auto-regressive latent reasoning (Xu et al., 2025; Tan et al., 2025). Coconut applies answer-level supervision (Hao et al., 2025), and CODI uses trajectory-level distillation (Shen et al., 2025). Our work introduces step-level supervision, which distributes signals across latent steps and improves stability. See extended discussion in Appendix B.

## 6 CONCLUSION

We introduce SIM-CoT, a training-based implicit reasoning method with step-level supervision on latent tokens. On GPT-2, SIM-CoT outperforms the strong explicit baseline SFT-CoT, while also surpassing implicit baselines such as Coconut and CODI. When scaling to larger LLaMA backbones, the performance achieves consistent gains over existing implicit reasoning methods and maintains fast inference efficiency. Ablation studies further show that it improves training stability with more latent tokens and can benefit from integration with training-free techniques such as soft thinking. Distance analysis confirms that SIM-CoT produces latent representations that are diverse yet stable.

540 REFERENCES  
541

542 Anthropic. Claude 3.5 sonnet, 2024. URL <https://www.anthropic.com/news/claude-3-5-sonnet>.

544 Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuwei Liu, Mengfei Zhou, Zhuosheng Zhang, et al. Do not think that much for  $2+3=?$  on the overthinking of 01-like llms. [arXiv preprint arXiv:2412.21187](https://arxiv.org/abs/2412.21187), 2024.

548 Yilong Chen, Junyuan Shang, Zhenyu Zhang, Yanxi Xie, Jiawei Sheng, Tingwen Liu, Shuohuan Wang, Yu Sun, Hua Wu, and Haifeng Wang. Inner thinking transformer: Leveraging dynamic depth scaling to foster adaptive internal thinking. [arXiv preprint arXiv:2502.13842](https://arxiv.org/abs/2502.13842), 2025.

551 Jeffrey Cheng and Benjamin Van Durme. Compressed chain of thought: Efficient reasoning through 552 dense representations. [arXiv preprint arXiv:2412.13171](https://arxiv.org/abs/2412.13171), 2024.

554 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, 555 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve 556 math word problems. [arXiv preprint arXiv:2110.14168](https://arxiv.org/abs/2110.14168), 2021.

557 Yuntian Deng, Yeqin Choi, and Stuart Shieber. From explicit cot to implicit cot: Learning to internalize 558 cot step by step. [ArXiv](https://arxiv.org/abs/2405.14838), abs/2405.14838, 2024.

560 Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and 561 Graham Neubig. Pal: Program-aided language models. [arXiv preprint arXiv:2211.10435](https://arxiv.org/abs/2211.10435), 2022.

562 Jonas Geiping, Sean McLeish, Neel Jain, John Kirchenbauer, Siddharth Singh, Brian R Bartoldson, 563 Bhavya Kailkhura, Abhinav Bhatele, and Tom Goldstein. Scaling up test-time compute with latent 564 reasoning: A recurrent depth approach. [arXiv preprint arXiv:2502.05171](https://arxiv.org/abs/2502.05171), 2025.

566 Google. Our next-generation model: Gemini 1.5, 2024. URL <https://blog.google/technology/ai/google-gemini-next-generation-model-february-2024>.

569 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, 570 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms 571 via reinforcement learning. [arXiv preprint arXiv:2501.12948](https://arxiv.org/abs/2501.12948), 2025.

572 Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhitong Hu, Jason Weston, and Yuandong Tian. 573 Training large language models to reason in a continuous latent space. In [COLM](https://colm.csail.mit.edu), 2025.

574 Rik Koncel-Kedziorski, Subhro Roy, Aida Amini, Nate Kushman, and Hannaneh Hajishirzi. MAWPS: 575 A math word problem repository. In [NAACL](https://naacl.org/naacl2016), 2016.

577 Jindong Li, Yali Fu, Li Fan, Jiahong Liu, Yao Shu, Chengwei Qin, Menglin Yang, Irwin King, and 578 Rex Ying. Implicit reasoning in large language models: A comprehensive survey. [arXiv preprint 579 arXiv:2509.02350](https://arxiv.org/abs/2509.02350), 2025.

580 Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, 581 Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-Refine: Iterative refinement with 582 self-feedback. In [NeurIPS](https://neurips.cc), 2023.

584 Meta. Llama 3.2: Revolutionizing edge ai and vision with open, cus- 585 tomizable models, 2024. URL <https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices>.

587 Amirkeivan Mohtashami, Matteo Pagliardini, and Martin Jaggi. Cotformer: More tokens with 588 attention make up for less depth. In [WANT@ NeurIPS](https://want.csail.mit.edu) 2023, 2023.

590 Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke 591 Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time 592 scaling. [arXiv preprint arXiv:2501.19393](https://arxiv.org/abs/2501.19393), 2025.

593 OpenAI. Hello gpt-4o, 2024. URL <https://openai.com/index/hello-gpt-4o>.

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

596

597 Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language  
 598 models are unsupervised multitask learners. [OpenAI Technical Report](#), 2019.

599 Subhro Roy and Dan Roth. Solving general arithmetic word problems. In [EMNLP](#), 2015.

600

601 Nikunj Saunshi, Nishanth Dikkala, Zhiyuan Li, Sanjiv Kumar, and Sashank J Reddi. Reasoning with  
 602 latent thoughts: On the power of looped transformers. [arXiv preprint arXiv:2502.17416](#), 2025.

603

604 Zhenyi Shen, Hanqi Yan, Linhai Zhang, Zhanghao Hu, Yali Du, and Yulan He. Codi: Compressing  
 605 chain-of-thought into continuous space via self-distillation. [arXiv preprint arxiv:2502.21074](#),  
 606 2025.

607

608 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion:  
 609 Language agents with verbal reinforcement learning. In [NeurIPS](#), 2023.

610

611 DiJia Su, Hanlin Zhu, Yingchen Xu, Jiantao Jiao, Yuandong Tian, and Qinqing Zheng. Token  
 612 assorted: Mixing latent and text tokens for improved language model reasoning. [arXiv preprint  
 613 arXiv:2502.03275](#), 2025.

614

615 Wenhui Tan, Jiaze Li, Jianzhong Ju, Zhenbo Luo, Jian Luan, and Ruihua Song. Think silently,  
 616 think fast: Dynamic latent compression of llm reasoning chains. [arXiv preprint arXiv:2505.16552](#),  
 617 2025.

618

619 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, and Denny Zhou. Self-  
 620 consistency improves chain of thought reasoning in language models. In [ICLR](#), 2023.

621

622 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V  
 623 Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In  
 624 [NIPS](#), 2022.

625

626 Ting-Ruen Wei, Haowei Liu, Xuyang Wu, and Yi Fang. A survey on feedback-based multi-step  
 627 reasoning for large language models on mathematics. [arXiv preprint arXiv:2502.14333](#), 2025.

628

629 Junhong Wu, Jinliang Lu, Zixuan Ren, Ganqiang Hu, Zhi Wu, Dai Dai, and Hua Wu. Llms have a  
 630 heart of stone: Demystifying the soft thinking ability of large reasoning models. [arXiv preprint  
 631 arXiv:2508.03440](#), 2025.

632

633 Yige Xu, Xu Guo, Zhiwei Zeng, and Chunyan Miao. SoftCoT: Soft chain-of-thought for efficient  
 634 reasoning with llms. In [ACL](#), 2025.

635

636 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao.  
 637 ReAct: Synergizing reasoning and acting in language models. In [ICLR](#), 2023.

638

639 Zhen Zhang, Xuehai He, Weixiang Yan, Ao Shen, Chenyang Zhao, Shuohang Wang, Yelong Shen,  
 640 and Xin Eric Wang. Soft thinking: Unlocking the reasoning potential of llms in continuous concept  
 641 space. [arXiv preprint arXiv:2505.15778](#), 2025.

642

643 Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min,  
 644 Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen,  
 645 Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and  
 646 Ji-Rong Wen. A survey of large language models. [arXiv preprint arXiv:2303.18223](#), 2023. URL  
 647 <http://arxiv.org/abs/2303.18223>.

648

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

651

652 Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans,  
 653 Ed Chi, and Quoc V. Le. Least-to-most prompting enables complex reasoning in large language  
 654 models. In [NeurIPS](#), 2023.

655

648

649

650

651

652

653

654

655

656

657

658

659

660

661

662

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

## APPENDIX

## USAGE OF LARGE LANGUAGE MODELS

In this paper, we used LLMs only for minor language polishing and formatting, without generating ideas, analyses, or experimental results.

## OUTLINE

In this appendix, we provide additional analyses and supporting materials to complement the main text. In Sec. A, we present experiments on combining soft thinking with SIM-CoT, including setup, results, and a detailed formulation with pseudocode. In Sec. B, we provide an overview of related work in explicit and implicit chain-of-thought reasoning. In Secs. C and D, we describe our implementation, hyperparameter configurations, **boundary conditions for implicit token alignment**, **training overhead** and training/inference procedures, including benchmark and dataset details. In Sec. E, we introduce the SIM-CoT training procedure and provide pseudocode for step-level supervision. In Sec. F, we offer geometric diagnostics of the latent space, analyzing inter-latent distances and distance to the vocabulary center. In Sec. G, we discuss interpretability analysis, including latent visualization and summary findings. Finally, we provide declarations on LLM usage and additional case studies on GSM8k to further illustrate the reasoning process and visualization choices.

## A ADDITIONAL ANALYSIS ON SOFT THINKING

Soft thinking (Zhang et al., 2025; Wu et al., 2025) is a training-free method for implicit reasoning in which the latent space is represented as a weighted average over the vocabulary embedding space. In contrast, SIM-CoT learns latent representations directly from data during training. To our knowledge, no prior work has evaluated a combination of these two approaches; our experiments provide the first such evaluation.

## A.1 SETUP

We apply the proposed soft thinking mechanism on top of both Coconut and SIM-CoT, while adopting GPT-2 as the backbone model. To assess the effectiveness of this approach, we perform evaluations on a diverse set of mathematical reasoning benchmarks. The in-domain evaluation is carried out on GSM8k-Aug, which provides augmented training and testing samples closely aligned with the original GSM8k distribution. To further examine generalization beyond the training domain, we include three out-of-domain benchmarks: GSM-Hard, which contains more challenging arithmetic problems with subtle variations in reasoning steps; MultiArith, which evaluates performance on multi-step arithmetic operations requiring careful sequencing of addition, subtraction, multiplication, and division; and SVAMP, which focuses on variations of elementary word problems designed to test robustness to superficial changes in problem statements.

## A.2 RESULTS

Table 5 (b) reports the results. Adding soft thinking improves accuracy in most cases. For Coconut, improvements are observed on GSM-Hard (+0.2) and MultiArith (+1.7), with a slight decrease on SVAMP (-0.2). For SIM-CoT, soft thinking consistently enhances performance: GSM8k-Aug (+0.2), GSM-Hard (+0.1), MultiArith (+0.7), and SVAMP (+0.1).

## A.3 FORMULATION

Let  $z \in \mathbb{R}^d$  denote a continuous latent token, and  $E \in \mathbb{R}^{|\mathcal{V}| \times d}$  be the embedding matrix of the vocabulary  $\mathcal{V}$ . Our goal is to enrich the representational capacity of  $z$  by incorporating soft thinking, which allows the latent space to draw information not only from its continuous representation but also from the semantic structure of the vocabulary. The process can be described in three steps.

702 **Step 1. Vocabulary distribution.** The continuous latent token  $z$  is first mapped into a probability  
 703 distribution over the vocabulary space:

704 
$$p = \text{softmax}(Wz),$$

706 where  $W \in \mathbb{R}^{|\mathcal{V}| \times d}$  is the output projection matrix and  $p \in \mathbb{R}^{|\mathcal{V}|}$  is the resulting distribution. This  
 707 step can be viewed as interpreting the latent token in terms of vocabulary-level semantics, where each  
 708 token in  $\mathcal{V}$  is assigned a likelihood according to its relevance to  $z$ .

710 **Step 2. Soft-thinking embedding.** Using the distribution  $p$ , we compute a weighted mixture of  
 711 vocabulary embeddings:

712 
$$z_{\text{soft}} = E^T p = \sum_{v \in \mathcal{V}} p_v E_v,$$

714 where  $E_v$  is the embedding vector corresponding to token  $v$ . This operation can be seen as constructing  
 715 a "soft token" that captures multiple semantic hypotheses simultaneously, instead of committing  
 716 to a single discrete vocabulary token. As a result,  $z_{\text{soft}}$  provides richer and smoother information than  
 717 a hard token lookup.

718 **Step 3. Combination.** Finally, we combine the original continuous latent  $z$  with the soft-thinking  
 719 embedding  $z_{\text{soft}}$ :

721 
$$z' = \alpha z + \beta z_{\text{soft}},$$

722 where  $\alpha = \text{continuous\_weight}$  and  $\beta = \text{soft\_weight}$  are hyperparameters that balance the  
 723 contribution of the continuous and soft-thinking components. This formulation allows  $z'$  to retain  
 724 the model's learned continuous representations while also grounding them in the vocabulary space.  
 725 Intuitively, the continuous part encourages compact reasoning within the latent space, whereas the  
 726 soft-thinking component brings in semantic priors from the vocabulary, leading to more stable and  
 727 interpretable reasoning.

728 The pseudocode implementation of the above process is presented as follows.

---

730 **Algorithm 1** Soft Thinking with Continuous Tokens

---

731 **Require:** Continuous latent  $z$ , embedding matrix  $E$ , weights  $\alpha, \beta$   
 732 1: **if**  $\beta > 0$  **then**  
 733 2:   Compute logits:  $l \leftarrow Wz$   
 734 3:   Convert to probabilities:  $p \leftarrow \text{softmax}(l)$   
 735 4:   Form soft embedding:  $z_{\text{soft}} \leftarrow E^T p$   
 736 5:   Update latent:  $z' \leftarrow \alpha z + \beta z_{\text{soft}}$   
 737 6: **else**  
 738 7:    $z' \leftarrow z$   
 739 8: **end if**  
 740 9: **return**  $z'$

---

741  
 742   A.4 ANALYSIS

744 The results demonstrate that soft thinking complements training-based implicit reasoning. The hybrid  
 745 latent  $z'$  integrates semantics learned through training and distributional information from vocabulary  
 746 mixing, which enables the model to explore diverse intermediate states rather than committing  
 747 to a single deterministic path. This leads to improvements in both in-domain and out-of-domain  
 748 benchmarks. Our findings suggest that combining training-free construction with training-based  
 749 supervision provides gains beyond either approach in isolation.

750  
 751   B RELATED WORK

752 **Explicit chain-of-thought reasoning.** Chain-of-thought (CoT) prompting enables large language  
 753 models (LLMs) to generate intermediate reasoning steps before producing the final answer (Wei et al.,  
 754 2022). This approach has been widely studied and extended in many directions. Self-consistency  
 755 samples multiple reasoning paths and selects the majority answer to improve reliability (Wang et al.,

2023). Least-to-most prompting decomposes a complex question into simpler sub-problems and solves them in order (Zhou et al., 2023). Reflection-based reasoning allows the model to revise or verify its own intermediate steps, leading to better correctness (Shinn et al., 2023; Madaan et al., 2023). Other works focus on using external tools or symbolic solvers together with explicit reasoning, which further improves accuracy in mathematics and program synthesis (Yao et al., 2023). Methods such as progressive-hint prompting (Zheng et al., 2023) and step-level feedback (Wei et al., 2025) study how supervision can be incorporated into explicit reasoning to make reasoning more structured. Despite these advances, explicit CoT has clear drawbacks. Because it generates long token sequences, inference cost grows rapidly with reasoning length, and many intermediate steps are redundant or irrelevant to the final answer. Moreover, since explicit reasoning is restricted to tokens from a fixed vocabulary, it often commits to a single trajectory and shows limited reasoning diversity (Li et al., 2025; Zhang et al., 2025; Xu et al., 2025).

**Implicit chain-of-thought reasoning.** Implicit CoT performs multi-step computation in a continuous latent space instead of emitting long textual traces, reducing decoded length while keeping internal structure. Prior work follows four practical routes. First, **knowledge internalization** trains models to carry out reasoning internally by progressively removing explicit traces or by using dedicated control embeddings; examples include iCoT-SI (Deng et al., 2024), which removes steps during training to internalize reasoning. Second, **architectural modification** controls compute by reusing or skipping layers, or by adding light recurrence, so models can refine hidden states without lengthening outputs (Saunshi et al., 2025; Chen et al., 2025; Cheng & Van Durme, 2024; Su et al., 2025; Mohtashami et al., 2023; Geiping et al., 2025). Third, **training-free** methods construct continuous latents directly from the model’s probability distribution over the vocabulary; Soft Thinking mixes embeddings by probability to form “concept” tokens that explore alternative paths without updating weights, which improves efficiency and diversity but does not bind each latent to step-level semantics (Zhang et al., 2025; Wu et al., 2025).

The fourth route, **auto-regressive latent reasoning**, updates and concatenates latent states in place of some token-level decoding and is the most relevant to our work (Xu et al., 2025; Tan et al., 2025). Coconut applies **answer-level** supervision—training on the final answer while leaving intermediate latents weakly constrained (Hao et al., 2025). CODI adds **trajectory-level** distillation by aligning an implicit trajectory with an explicit CoT trace, narrowing the gap to explicit CoT but giving only coarse guidance to intermediate steps (Shen et al., 2025). However, the implicit token length in CODI is fixed during training, which limits its flexibility and makes it less suitable for scaling to variable or longer reasoning chains. Our framework remains in the auto-regressive setting but changes the supervision: during training, each latent is aligned with its corresponding textual step (**step-level** supervision), distributing learning signals across the full latent chain to improve stability and semantic fidelity of intermediate states; at inference, the decoder is discarded, ensuring that the decoding cost remains identical to that of standard implicit CoT methods (e.g., Coconut).

## C IMPLEMENTATION AND TRAINING DETAILS

We provide the full hyperparameter settings, training procedures, and additional analysis used in our experiments. Unless otherwise specified, we use the AdamW optimizer with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , and weight decay of 0.1. Batch size is set to 128 for GPT-2 and LLaMA 1B, and 64 for LLaMA 3B and 8B. Early stopping is applied with a patience of 3 epochs. We now describe the training setups for Coconut, CODI, and our SIM-CoT.

### C.1 COCONUT TRAINING SETUP

Following Hao et al. (2025), GPT-2 and LLaMA 1B (Radford et al., 2019; Meta, 2024) are trained with a fixed learning rate of  $1 \times 10^{-4}$ . One implicit latent corresponds to two implicit tokens. A curriculum is applied: every three epochs, one explicit reasoning step is replaced by an implicit latent until the maximum number of latent steps is reached. After this expansion, training continues for 15 additional epochs.

810 C.2 CODI TRAINING SETUP  
811

812 For larger backbones such as LLaMA 3B and LLaMA 8B, we adopt task-specific hyperparameter  
813 settings to ensure stable training. In particular, we use a learning rate of  $3 \times 10^{-4}$  for LLaMA 3B  
814 and train for 8 epochs, while for LLaMA 8B the learning rate is reduced to  $1 \times 10^{-4}$  with 6 training  
815 epochs. These choices are motivated by the increased sensitivity of larger models to optimization  
816 dynamics, where smaller learning rates and fewer epochs help to prevent overfitting and instability.

817 When reproducing CODI on GPT-2 and LLaMA 1B, we strictly follow the configurations reported  
818 by Shen et al. (2025). Specifically, we use a learning rate of  $3 \times 10^{-3}$  with 40 epochs for GPT-2, and  
819 a learning rate of  $8 \times 10^{-4}$  with 10 epochs for LLaMA 1B. Adopting these settings ensures that our  
820 results are directly comparable to prior work and isolates the effect of our proposed method, rather  
821 than confounding it with differences in optimization schedules.

822 C.3 SUMMARY OF HYPERPARAMETERS  
823824 Table 6: Training hyperparameters across different models.  
825

826 Model	827 Method	828 LR	829 Epochs
828 GPT-2	829 Coconut	$1 \times 10^{-4}$	15 + curriculum
829 LLaMA 1B	830 Coconut	$1 \times 10^{-4}$	15 + curriculum
830 GPT-2	831 CODI	$3 \times 10^{-3}$	40
831 LLaMA 1B	832 CODI	$8 \times 10^{-4}$	10
832 LLaMA 3B	833 CODI	$3 \times 10^{-4}$	8
833 LLaMA 8B		$1 \times 10^{-4}$	6

834 C.4 TRAINING-TIME OVERHEAD  
835

836 We analyze the computational overhead introduced during training by the auxiliary decoder. Since  
837 the auxiliary decoder has the same architecture and number of parameters as the original decoder, and  
838 it participates in an additional forward pass during training, the overall parameter count and memory  
839 usage are approximately doubled compared with the implicit baselines Coconut and CODI.

840 To quantify the training-time overhead, Table 7 reports the wall-clock training hours of SIM-CoT and  
841 the corresponding implicit models under identical hardware settings (H800). Across different model  
842 scales, SIM-CoT introduces only a moderate increase in training time, ranging from approximately 2  
843 hours for smaller backbones to around 16 hours for larger ones.

844 Table 7: Training hours of SIM-CoT compared with implicit baselines under identical hardware  
845 (H800).  
846

847 Model	848 Training Hours
848 Coconut GPT	$\sim 180\text{h}$
849 SIM-CoT (Coconut GPT)	$\sim 192\text{h}$
850 CODI GPT	$\sim 16\text{h}$
851 SIM-CoT (CODI GPT)	$\sim 18.2\text{h}$
852 CODI 1B	$\sim 16.5\text{h}$
853 SIM-CoT (CODI 1B)	$\sim 18.5\text{h}$
854 CODI 3B	$\sim 34\text{h}$
855 SIM-CoT (CODI 3B)	$\sim 42\text{h}$
856 CODI 8B	$\sim 71\text{h}$
857 SIM-CoT (CODI 8B)	$\sim 87\text{h}$

858 These results show that the additional computational cost introduced by SIM-CoT remains modest  
859 relative to the improvements in stability and accuracy it provides.

864

865

866

## C.5 BOUNDARY CONDITIONS IN IMPLICIT TOKEN-STEP ALIGNMENT

867

868

869

870

871

Since the model produces a fixed number of implicit tokens ( $K$ ) independent of the length of the ground-truth reasoning chain, we formalize the boundary conditions that govern how these tokens are aligned with textual reasoning steps. Let the annotated reasoning contain  $N$  steps. Under this formulation, two boundary cases emerge:

872

873

874

**Case 1:  $N > K$  (Longer reasoning chains).** When the reasoning chain contains more steps than latent tokens, a direct one-to-one alignment is impossible. We adopt a *many-to-one* strategy: the first  $K - 1$  latent tokens are aligned individually to the first  $K - 1$  reasoning steps, while the final token  $z_K$  receives supervision from the concatenation of the remaining steps (steps  $K$  through  $N$ ). This boundary condition ensures that information from longer chains is preserved rather than truncated.

875

876

877

878

879

880

**Case 2:  $N < K$  (Shorter reasoning chains).** When the reasoning chain is shorter than the number of latent tokens, the first  $N$  latent tokens align one-to-one with the available reasoning steps. The remaining latent tokens ( $z_{N+1}$  to  $z_K$ ) are aligned with the final answer. This encourages the model to utilize its surplus latent capacity to refine the target solution.

881

882

883

884

These boundary rules guarantee that each implicit token is assigned a semantically meaningful supervision signal, regardless of how the reasoning length compares with the fixed latent budget  $K$ . They also provide stability during training by preventing both supervision sparsity (when  $N < K$ ) and information loss (when  $N > K$ ).

885

886

## D TRAINING AND INFERENCE DETAILS

887

888

**Curriculum for  $K$ .** We use a curriculum schedule to gradually increase the number of implicit steps. Each latent corresponds to two implicit tokens. Let  $K_{\max}$  denote the maximum number of latents. Starting from  $K^{(0)} = 0$ , the number of implicit steps after epoch  $e$  is

$$K^{(e)} = \min\left(K_{\max}, \left\lfloor \frac{e}{\Delta e} \right\rfloor\right),$$

where  $\Delta e$  is the update interval in epochs. Once  $K^{(e)}$  reaches  $K_{\max}$ , it remains fixed for the remainder of training.

889

890

**Inference and Efficiency.** At inference time, the auxiliary decoder is removed and only the base model is executed:

891

892

$$U^{(0)} = [e(x_1), \dots, e(x_T)], \quad \text{for } k = 1, \dots, K : z_k = H_{\theta}(U^{(k-1)}), \quad U^{(k)} = U^{(k-1)} \oplus z_k,$$

893

894

and after  $K$  implicit steps the model switches back to explicit decoding to generate the final answer sequence  $a$  as in Eq. equation 4. The total decoding length is  $T + K + L_a$ , where  $T$  is the input length,  $K$  the number of implicit steps, and  $L_a$  the answer length. In practice, the cost is comparable to other implicit reasoning methods because  $K$  is moderate. In tasks where explicit CoT requires long trajectories ( $L_{\text{CoT}} \gg K$ ), the implicit formulation reduces decoding positions, providing efficiency gains without loss of reasoning accuracy.

895

896

897

898

899

**Benchmark Detail. SVAMP** (Patel et al., 2021) (Simple Variations on Arithmetic Math Word Problems) is a benchmark dataset designed to test the robustness of math word problem solvers to superficial changes. It contains 1,000 elementary-level arithmetic word problems (grade 4 and below), each involving a single unknown and solvable by an arithmetic expression with no more than two operators. The problems are transformations of existing datasets (such as MAWPS and ASDiv-A) with controlled variations in wording, structure, and number values to reduce artifacts and superficial cues. SVAMP’s average number of reasoning steps required is around 1.2, similar to the base datasets, but model performance drops significantly when tested on SVAMP, showing that many models rely on heuristic patterns rather than deep understanding.

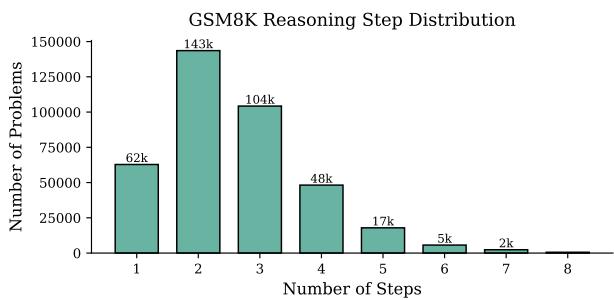


Figure 5: Distribution of reasoning steps in the GSM8K-Aug training dataset. Most problems involve two to four steps, with a long-tail of harder cases. For visualization, step counts with fewer than 200 problems are omitted, though all examples are used in training.

**GSM-Hard** (Gao et al., 2022) is a more challenging variant of the GSM8K dataset, intended to test models’ ability to cope with harder numerical values. It retains the same problem statements as the original GSM8K, but replaces many of the numbers with larger and less common numerical values, making superficial arithmetic computation and reasoning harder. The dataset has about 1,319 examples (matching the GSM8K test set size).

**MultiArith** (Roy & Roth, 2015) is a dataset of multi-step arithmetic word problems designed to challenge systems to correctly sequence multiple operations. It contains 600 problems collected from educational sources, each solvable by one equation involving two or more of the four basic operations (addition, subtraction, multiplication, division). The problems require reasoning over multiple sentences to extract and combine numeric quantities, understand the implied operations, and compute the result. MultiArith has been widely used as a benchmark for evaluating arithmetic reasoning generalization, particularly for models that go beyond single-operation problems. Analyses show that many models struggle on these examples compared to simpler datasets, highlighting the importance of handling compositional and sequential numerical reasoning.

**Training Data.** **GSM8K-Aug** (Deng et al., 2024) is the only training corpus we use. It is an augmented dataset derived from GSM8K (Cobbe et al., 2021), expanding the original 8.5k training problems to roughly 385k examples through paraphrasing, numerical resampling, and synthetic generation with GPT-4. The distribution of reasoning steps in GSM8K-Aug is illustrated in Fig. 5, where the majority of problems require two to four steps, while a long-tail of six or more steps persists. This balance of common and complex instances makes GSM8K-Aug particularly suitable for training models that need to generalize across reasoning difficulty levels.

## E SIM-COT TRAINING IMPLEMENTATION

We provide pseudocode for the SIM-CoT training process, which illustrates how continuous latent embeddings are aligned with explicit supervision at the step level. In particular, each reasoning step in the explicit chain is mapped to a corresponding latent representation, and the training objective enforces consistency between the predicted latent tokens and the ground-truth step annotations. This design ensures that the model learns to represent intermediate reasoning steps in a compact latent space while still retaining interpretability through explicit alignment. By supervising at the step level rather than only at the final answer or trajectory level, SIM-CoT enables finer control over the reasoning process and reduces instability that often arises when scaling to longer chains or larger numbers of implicit tokens.

## F GEOMETRIC DIAGNOSTICS OF THE LATENT SPACE

We analyze the geometry of latent representations with two metrics.

*Inter-latent distance.*

$$\text{Dist}(z_{1:K}) = \frac{2}{K(K-1)} \sum_{1 \leq i < j \leq K} \|z_i - z_j\|_2. \quad (9)$$

A larger value indicates better separation, reducing the risk of collapse.

---

972 **Algorithm 2** SIM-CoT Training Procedure

---

973 **Require:** Batch size  $b$ , number of thoughts  $C$ , continuous embeddings  $Z$ , tokenized inputs  $X$ ,  
 974 embedding matrix  $E$

975 1: **for** each thought  $t = 1, \dots, C$  **do**

976 2:   **for** each sample  $i = 1, \dots, b$  **do**

977 3:     Extract continuous embeddings  $z_{i,t}$  from  $Z$

978 4:     Obtain token embeddings  $e_{i,t}$  from  $E(X_{i,t})$

979 5:     Concatenate embeddings:  $h_{i,t} \leftarrow [z_{i,t}; e_{i,t}]$

980 6:     Build attention mask  $m_{i,t}$  up to EOS

981 7:     Assign position ids  $p_{i,t}$

982 8:     Prepare labels  $y_{i,t}$  with masked tokens set to  $-100$

983 9:   **end for**

984 10: **end for**

985 11: Pad and stack  $\{h, m, p, y\}$  to maximum sequence length

986 12: Prepare 4D attention mask:  $\hat{M} \leftarrow \text{PrepareMask}(M)$

987 13: Forward pass:  $\hat{O} \leftarrow \text{ExplainableLLM}(H, \hat{M}, P)$

988 14: Extract logits:  $L \leftarrow \hat{O}.\text{logits}$

989 15: Shift logits and labels:  $L' \leftarrow L[:, :-1]$ ,  $Y' \leftarrow Y[:, 1:]$

990 16: Compute cross-entropy loss:  $\ell = \text{CrossEntropy}(L', Y')$

991 17: Normalize  $\ell$  over valid positions

992 **Ensure:** Final training loss  $\ell$

---

993  
 994 *Distance to vocabulary center.* Let  $\mu = \frac{1}{|\mathcal{V}|} \sum_{v \in \mathcal{V}} E_v$  denote the mean embedding. Then,

995  
 996 
$$\text{DistVC}(z_{1:K}) = \frac{1}{K} \sum_{k=1}^K \|z_k - \mu\|_2. \quad (10)$$

997  
 998

999 Moderate values indicate that latents remain close enough to the lexical manifold for stability, while  
 1000 avoiding collapse toward the center. These diagnostics are not used in training but serve as indicators  
 1001 of diversity and stability in the learned latent space.

1002 **G ADDITIONAL DETAILS FOR INTERPRETABILITY ANALYSIS**

1003 **G.1 MAKING IMPLICIT REASONING VISIBLE**

1004  
 1005 Continuous thoughts produced by implicit reasoning models are represented as latent embeddings  
 1006 that do not correspond to discrete vocabulary tokens, and therefore cannot be directly decoded by a  
 1007 tokenizer. This makes it difficult to interpret how the model internally organizes multi-step reasoning.  
 1008 To address this, we reuse the decoder that was employed for step-level supervision during training,  
 1009 and apply it at inference time to map each latent embedding into a human-readable token sequence.

1010  
 1011 As illustrated in Figure 4, the process begins with a natural language problem (e.g., a math word  
 1012 problem) that is embedded and passed into the large language model. The model generates a sequence  
 1013 of implicit latent tokens, which capture intermediate reasoning steps in continuous space. These latent  
 1014 tokens are then fed into the optional decoder, which translates them into interpretable expressions.  
 1015 Each latent corresponds to one reasoning step, and the autoregressive generation order encodes the  
 1016 dependency structure across steps.

1017  
 1018 For example, in the GSM8k case study shown in the figure, the first latent is decoded as  $0.3 \times 120 = 36$ ,  
 1019 representing the number of watermelons harvested initially. The second latent builds upon this result  
 1020 to compute  $120 - 36 = 84$ , the remaining melons. The third latent then calculates  $\frac{3}{4} \times 84 = 63$ , and  
 1021 the final latent derives the answer  $84 - 63 = 21$ . For clarity and to save space, the figure merges  
 1022 the second and third steps into a single box, but the actual implicit reasoning unfolds across four  
 1023 distinct latent steps. This sequence of decoded latents mirrors the logic of explicit chain-of-thought  
 1024 reasoning, while being produced implicitly within the latent space.

1025 By projecting implicit tokens into interpretable space, we gain direct visibility into how the model  
 1026 structures multi-step reasoning. This not only enables analysis of the correctness and consistency of

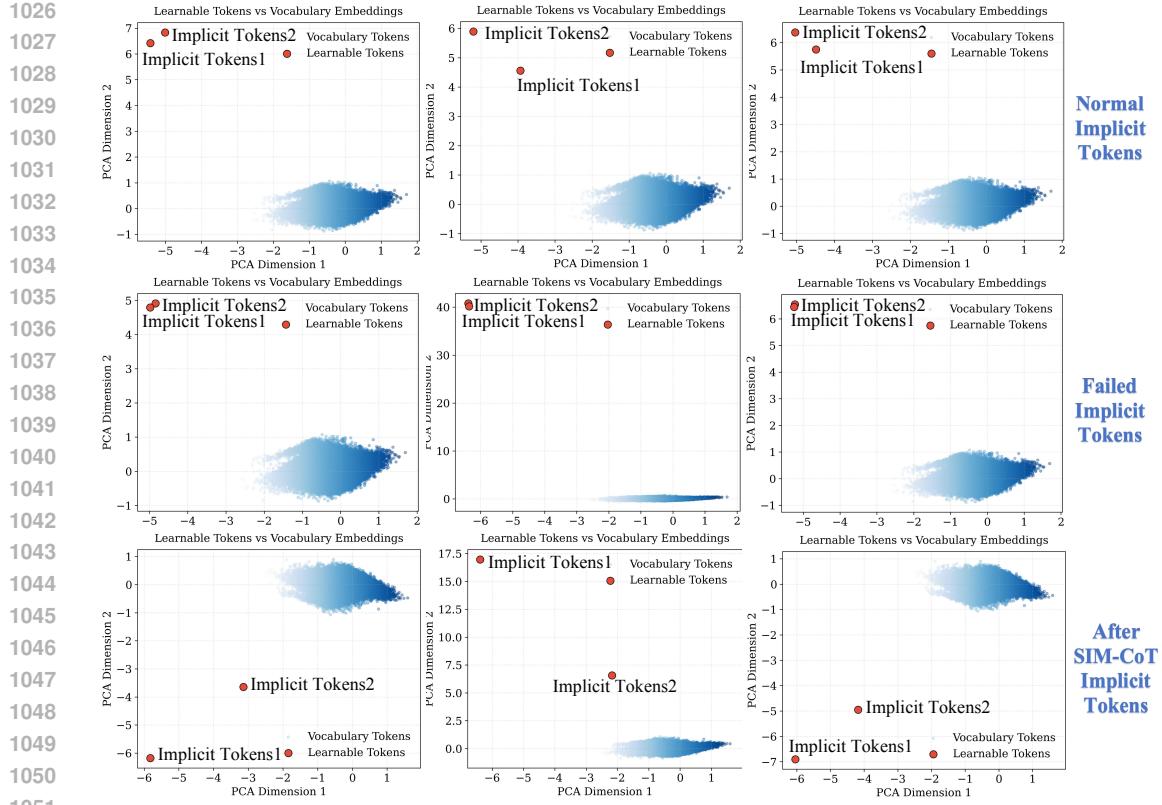


Figure 6: Visualization of distances among implicit tokens and their distances to the vocabulary center. The first row shows normal implicit tokens with well-separated representations, the second row illustrates failed implicit tokens where distances collapse and drift away from the vocabulary center, and the third row presents implicit tokens after applying SIM-CoT, which restores both separation and stability in the latent space.

intermediate steps, but also highlights the dependencies across steps that underlie the final prediction. The visualization confirms that SIM-CoT can encode semantically meaningful and logically ordered reasoning steps in its latent space, bridging the gap between implicit and explicit reasoning.

## G.2 SUMMARY

Overall, the results demonstrate that SIM-CoT establishes a balance between **diversity** and **stability** in the latent space. Larger inter-latent distances mitigate representation collapse, while moderate distances to the vocabulary center prevent excessive drift. This equilibrium supports stable implicit reasoning and provides robustness when scaling to more latent tokens.

## H ADDITIONAL CASE STUDIES ON GSM8K

In practice, implicit reasoning continues to produce latent tokens even after the correct answer has been reached. These trailing latents no longer introduce new steps but simply repeat the final prediction. For clarity, we omit such redundant tokens in the visualizations. As a result, only the latents that correspond to meaningful intermediate steps are displayed in Figure 7, while those mapping directly to the final answer are hidden. This choice improves readability without changing the underlying reasoning process.

Notably, the decoded reasoning steps consistently match the semantic structure of explicit chain-of-thought annotations, while being generated implicitly within the latent space. The final predictions align with the ground-truth answers, demonstrating that SIM-CoT is capable of encoding interpretable and step-ordered reasoning without requiring explicit supervision at inference time.

1080  
1081  
1082  
1083  
1084  
1085

[Question] Toulouse has twice as many sheep as Charleston. Charleston has 4 times as many sheep as Seattle. How many sheep do Toulouse, Charleston, and Seattle have together if Seattle has 20 sheep? 🤔     ### Answer: 260

$$4*20=80 \text{ latent1} \rightarrow 2*80=160 \text{ latent2} \rightarrow 160+80+20=260 \text{ latent3} \text{ 😊}$$

1086  
1087  
1088  
1089  
1090  
1091  
1092  
1093  
1094  
1095  
1096

[Question] Claire makes a 3 egg omelet every morning for breakfast. How many dozens of eggs will she eat in 4 weeks? 🤔     ### Answer: 7

$$3*7=21 \text{ latent1} \rightarrow 21*4=84 \text{ latent2} \rightarrow 84/12=7 \text{ latent3} \text{ 😊}$$

1097  
1098  
1099  
1100  
1101  
1102  
1103

[Question] Billy sells DVDs. He has 8 customers on Tuesday. His first 3 🤔 customers buy one DVD each. His next 2 customers buy 2 DVDs each. His last 3 customers don't buy any DVDs. How many DVDs did Billy sell on Tuesday?

### Answer: 7

$$3*1=3 \text{ latent1} \rightarrow 2*2=4 \text{ latent2} \rightarrow 3+4=7 \text{ latent3} \text{ 😊}$$

1104  
1105  
1106  
1107  
1108  
1109

[Question] Richard lives in an apartment building with 15 floors. Each floor contains 8 units, and  $\frac{3}{4}$  of the building is occupied. What's the total number of unoccupied units in the building? 🤔     ### Answer: 30

$$15*8=120 \text{ latent1} \rightarrow 120*\frac{3}{4}=90 \text{ latent2} \rightarrow 120-90=30 \text{ latent3} \text{ 😊}$$

1110  
1111  
1112  
1113  
1114  
1115  
1116

[Question] Poppy is solving a 1000-piece jigsaw puzzle. She places a quarter of the pieces on the board, then her mom places a third of the remaining pieces. How many jigsaw pieces are left to be placed? 🤔     ### Answer: 500

$$1000/4=250 \text{ latent1} \rightarrow 1000-250=750 \text{ latent2} \rightarrow 750/3=250 \text{ latent3} \rightarrow 750-250=500 \text{ latent4} \text{ 😊}$$

1117  
1118  
1119  
1120  
1121  
1122  
1123  
1124

[Question] The marching band is ordering new uniforms. Each uniform comes with a hat that costs \$25, a jacket that costs three times as much as the hat, and pants that cost the average of the costs of the hat and jacket. How much does each uniform cost total? 🤔     ### Answer: 150

$$25*3=75 \text{ latent1} \rightarrow 25+75=100 \text{ latent2} \rightarrow 100/2=50 \text{ latent3} \rightarrow 100+50=150 \text{ latent4} \text{ 😊}$$

1125  
1126  
1127  
1128  
1129  
1130  
1131  
1132  
1133

Figure 7: Additional SIM-CoT case studies on GSM8k. Each example illustrates how implicit latent tokens correspond to intermediate reasoning steps. Arrows indicate the dependency relations across steps, while colored spans in the question highlight the textual evidence that supports each step. The decoded sequence of latent steps produces the correct final answer, which matches the ground-truth label.