

To Think or Not to Think: The Hidden Cost of Meta-Training with Excessive CoT Examples

Vignesh Kothapalli¹, Ata Fatahibaarzi², Hamed Firooz, Maziar Sanjabi

¹Stanford University, ²LinkedIn AI

Correspondence: vigneshk@cs.stanford.edu

Abstract

Chain-of-thought (CoT) prompting combined with few-shot in-context learning (ICL) has unlocked significant reasoning capabilities in large language models (LLMs). However, ICL with CoT examples is ineffective on novel tasks when the pre-training knowledge is insufficient. We study this problem in a controlled setting using the *CoT-ICL Lab* (Kothapalli et al., 2025) framework, and propose meta-training techniques to learn novel abstract reasoning tasks in-context. Although CoT examples facilitate reasoning, we noticed that their excessive inclusion during meta-training degrades performance when CoT supervision is limited. To mitigate such behavior, we propose CoT-Recipe, a formal approach to modulate the mix of CoT and non-CoT examples in meta-training sequences. We demonstrate that careful modulation via CoT-Recipe can increase the accuracy of transformers on novel tasks by up to 300% even when there are no CoT examples available in-context. We confirm the broader effectiveness of these techniques by applying them to pretrained LLMs (Qwen2.5 series) for symbolic reasoning tasks and observing gains of up to 130% in accuracy. Code is available at: <https://github.com/kvignesh1420/cot-icl-lab>

1 Introduction

Recent advances in large language models (LLMs) have demonstrated remarkable reasoning abilities when prompted to generate step-by-step solutions to problems. A prime example is *chain-of-thought* (CoT) prompting (Wei et al., 2022; Kojima et al., 2022), where appending a prompt with “*Let’s think step by step*” can induce an LLM to generate intermediate “thought” steps (Nye et al., 2022), and enable them to tackle multi-step problems. CoT prompting, specially when combined with *in-context learning* (ICL) (Brown et al., 2020) has yielded impressive gains on arithmetic, common-sense, and symbolic reasoning benchmarks (Wei

et al., 2022; Kojima et al., 2022), and has become a key technique for eliciting reasoning in LLMs.

Despite its successes, CoT in-context learning (CoT-ICL) faces several limitations. First, models often require carefully chosen exemplars (Min et al., 2022b; Zhao et al., 2021) for effective ICL. Next, few-shot CoT prompting uses handcrafted demonstrations of question-answer pairs with rationales, which can be labor-intensive to create for each task (Wei et al., 2022; Kim et al., 2023). Moreover, the benefits of CoT prompting tend to emerge with model scale; while smaller models struggle to produce answers with good reasoning unless fine-tuned to do so (Li et al., 2023a; Huang et al., 2023; Ho et al., 2023; Kim et al., 2023).

These issues are exacerbated when the tasks are entirely novel and the pre-training knowledge of LLMs is insufficient to generate the correct responses. For example, prompting LLMs to answer domain-specific queries whose background knowledge is not included in the pre-training data. In such scenarios, the models have to rely solely on the (possibly limited) task descriptions and the in-context examples to generate a response. Having CoT examples aids in revealing more information about the task, but their availability might be limited due to data curation constraints.

While previous works have explored meta-training approaches (Min et al., 2022a; Chen et al., 2022) with ICL as an objective, the role of CoT exemplars in the data recipes and inference prompts has been largely overlooked. By addressing this gap, our work aims to understand if models can be meta-trained to effectively leverage the (limited) CoT examples at inference for solving novel tasks. In particular, we study this problem in a controlled setting using the *CoT-ICL Lab* framework (Kothapalli et al., 2025) for abstract reasoning with transformers. Although CoT exemplars can aid in learning about the task, we find that their excessive inclusion during meta-training can

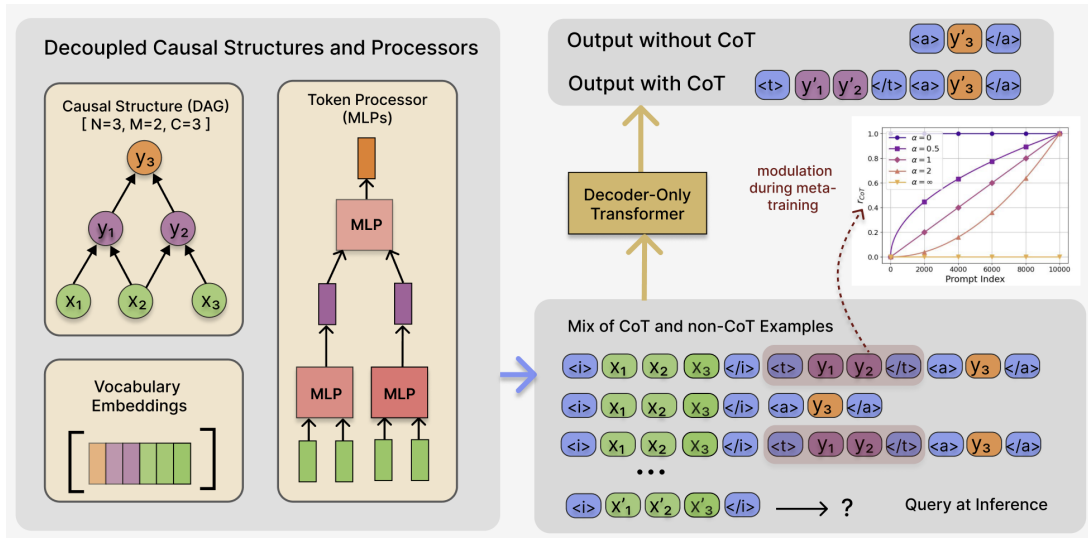


Figure 1: The *CoT-ICL Lab-2.0* framework. (1) We incorporate special tokens (marked in blue) to act as delimiter tokens between input, intermediate/thinking, and answer tokens. (2) Each sequence is constructed using a specific DAG, which determines the number of input and chain tokens per example. (3) The choice of r_{CoT} is varied across sequences as per the CoT-Recipe and modulates the mix of *CoT/standard examples* for meta-training.

be detrimental to the model’s performance when such supervision is limited (during inference). We propose principled data curation recipes to modulate the mix of CoT and non-CoT examples in sequences to address this issue (Figure 1). We also create a novel symbolic reasoning dataset called CIL-LangSym and meta-train LLMs (Qwen-2.5 series) with our data recipes to show that they can reason effectively on these domain-specific queries (1) even in the absence of CoT exemplars and (2) limited task descriptions. In summary, our key contributions are as follows:

1. We introduce *CoT-ICL Lab-2.0*, an extension of *CoT-ICL Lab* by Kothapalli et al. (2025) for meta-training on abstract reasoning tasks. It incorporates special tokens to isolate the ‘input’, ‘thinking’, and ‘answer’ segments of examples, and allows *reasoning control* in trained transformers by enabling dynamic invocation of multi-step reasoning or direct final answers as needed.
2. We introduce CoT-Recipe, a formal approach to modulate the mix of CoT and non-CoT examples in sequences during meta-training. In essence, datasets curated with this approach allow the models to reason even in the absence of CoT exemplars.
3. We leverage the insights from systematic experiments with *CoT-ICL Lab-2.0* to im-

prove the CoT-ICL capabilities of real-world LLMs (Qwen-2.5 series) on symbolic reasoning tasks. Especially when their pre-trained knowledge is insufficient for reasoning.

2 Related Work

Chain-of-Thought Prompting. CoT prompting with (Wei et al., 2022) and without (Kojima et al., 2022) ICL has been an effective strategy to improve model performance in complex reasoning tasks. However, CoT prompting’s effectiveness is strongly dependent on model scale and the quality of exemplars (Li et al., 2023a). Additionally, designing good CoT exemplars for few-shot prompting can be non-trivial (Min et al., 2022b; Zhao et al., 2021) as the exemplars need to be representative and align with the task at hand (Wang et al., 2023; Chu et al., 2024). This highlights the brittleness of the current models in utilizing the CoT exemplar¹. Beyond the basic paradigm of CoT prompting, the ‘ReAct’ framework by Yao et al. (2023) blends CoT prompting with actions that interface with an external environment. In this framework, the model is prompted in a format where it alternates between generating a thought (reflection on the task state) and an action (like a command to a tool or an environment). This is a departure from pure CoT, but it underscores a prompt design principle: by

¹For instance: The Tulu-3 report (Lambert et al., 2024) describes the effectiveness of zero-shot CoT over 1-shot CoT prompts for evaluations on MATH (Hendrycks et al., 2021).

structuring the model’s output into labeled sections (thought vs. action), one can guide it to perform complex interactive tasks. Our design of the *CoT-ICL Lab-2.0* framework by introducing ‘special’ tokens has a similar flavor in that we isolate different parts of the output (reasoning vs. answer), which in turn allows us to modulate the mix of CoT and non-CoT examples in-context.

In-Context Learning and Meta-Training. The ICL capability of LLMs (Brown et al., 2020) to recognize and generalize to unseen patterns during inference has been extensively studied and explored in many practical and theoretical settings (Achiam et al., 2023; Team et al., 2024; Firooz et al., 2025; Dong et al., 2024; Zhou et al., 2024). Beyond the theoretical exploration of this behaviour using single-layer transformers (Oko et al., 2024; Chang et al., 2025) and real-valued inputs (Garg et al., 2022; Bai et al., 2023; Li et al., 2023b), a recently developed framework called *CoT-ICL Lab* (Kothapalli et al., 2025) leveraged a vocabulary of abstract tokens to train transformer models from scratch and shed light on the emergent ICL capabilities. The prompt design employed by Kothapalli et al. (2025) for the abstract reasoning tasks is similar to the meta-training approaches (Min et al., 2022a; Chen et al., 2022; Min et al., 2022b) used in the broader literature with NLP prompts. However, as with the design of any few-shot prompt (Kim et al., 2023; Biderman et al., 2024), the sequences are usually limited to either including/excluding the ‘chain’ tokens (i.e, the thinking tokens) in *all* the in-context examples. Thus, the current understanding of the proportion of CoT vs non-CoT examples needed to meta-train these models and facilitate ICL is rather limited. Our work sheds light on this underexplored aspect and demonstrates the effectiveness of such modulation when the tasks (at inference) have limited CoT examples.

3 Preliminaries and Setup

Notation. Let $[K] = \{1, \dots, K\}$. We consider a vocabulary \mathcal{V} of ‘abstract’ tokens that is associated with a data embedding matrix $\mathbf{E}_{\text{data}} \in \mathbb{R}^{|\mathcal{V}| \times d}$. Here d denotes the data embedding dimension and the entries of \mathbf{E}_{data} are sampled i.i.d from $\mathcal{N}(0, 1)$. Let \mathcal{G}, \mathcal{H} denote the causal structure and token processing function classes respectively. Formally, \mathcal{G} represents functions that select M tokens from an arbitrary number of input tokens. \mathcal{H} represents functions that process the data embeddings (i.e,

rows of \mathbf{E}_{data}) of M tokens and output a single token. In this setup, we are interested in learning a class of functions $f \in \mathcal{F}$ that are compositional $f = f_C \circ f_{C-1} \cdots \circ f_1$ and can be formulated as $f_c = h_c \circ g_c$, where $g_c \in \mathcal{G}, h_c \in \mathcal{H}$. Given N input tokens $\mathbf{x} = (x_1, \dots, x_N)$, f recursively generates C chain tokens $\mathbf{y} = (y_1, \dots, y_C)$ as:

$$y_c = h_c(g_c(\mathbf{x}, y_1, \dots, y_{c-1})), \forall c \in [C]. \quad (1)$$

Here $\mathbf{y}_{:C-1} = (y_1, \dots, y_{C-1})$ denote the *intermediate/thinking tokens*, and y_C represents the final *answer token*.

Special tokens. We reserve a subset of tokens ($\mathcal{V}_{\text{special}} \in \mathcal{V}$) to serve as delimiters. The rest of the vocab is denoted by $\mathcal{V}_{\text{normal}}$. These special tokens comprise of $t_{\text{bos}}, t_{\text{eos}}, t_{\text{pad}}, t_{\text{inp_start}}, t_{\text{inp_end}}, t_{\text{think_start}}, t_{\text{think_end}}, t_{\text{ans_start}}$, and $t_{\text{ans_end}}$. The role of these tokens are similar to the ones used in NLP. For example: $t_{\text{think_start}}, t_{\text{think_end}}$ here correspond to the $\langle \text{think} \rangle, \langle / \text{think} \rangle$ tokens in NLP (see Table 3).

In-context example design. We design a *standard example* $\mathbf{e}(f)$ using tokens from $\mathcal{V}_{\text{special}}$, the N input tokens $\mathbf{x} \in \mathcal{V}_{\text{normal}}^N$ and the *answer token* $y_C \in \mathcal{V}_{\text{normal}}$ from (1) as:

$$\begin{aligned} \mathbf{x}' &= (t_{\text{inp_start}}, \mathbf{x}, t_{\text{inp_end}}), \\ \mathbf{a}' &= (t_{\text{ans_start}}, y_C, t_{\text{ans_end}}), \\ \mathbf{e}(f) &= (\mathbf{x}', \mathbf{a}', t_{\text{eos}}). \end{aligned} \quad (2)$$

Similarly, we include tokens $t_{\text{think_start}}, t_{\text{think_end}}$ along with the $C - 1$ *intermediate tokens* $\mathbf{y}_{:C-1} \in \mathcal{V}_{\text{normal}}^{C-1}$ in a *standard example* to obtain a *CoT example* $\mathbf{e}_{\text{CoT}}(f)$ as follows:

$$\begin{aligned} \mathbf{x}' &= (t_{\text{inp_start}}, \mathbf{x}, t_{\text{inp_end}}), \\ \mathbf{t}' &= (t_{\text{think_start}}, \mathbf{y}_{:C-1}, t_{\text{think_end}}), \\ \mathbf{a}' &= (t_{\text{ans_start}}, y_C, t_{\text{ans_end}}), \\ \mathbf{e}_{\text{CoT}}(f) &= (\mathbf{x}', \mathbf{t}', \mathbf{a}', t_{\text{eos}}). \end{aligned} \quad (3)$$

Tokenized sequence design. Instead of designing a sequence with only *standard* or *CoT examples*, we employ a CoT probability parameter r_{CoT} and diversify the sequences with a mix of both types (Figure 1). Formally, a sequence $\mathbf{p}^K(f, r_{\text{CoT}})$ with K examples is given by:

$$\begin{aligned} \mathbf{p}^K(f, r_{\text{CoT}}) &= \left(t_{\text{bos}}, \left(\mathbf{z}^{(i)} \right)_{i=1}^K \right), \\ \text{where } \mathbf{z}^{(i)} &= \begin{cases} \mathbf{e}_{\text{CoT}}^{(i)}(f) & \text{if } r_{\text{CoT}} \geq u^{(i)}, \\ \mathbf{e}^{(i)}(f) & \text{otherwise.} \end{cases} \end{aligned} \quad (4)$$

Here $u^{(i)} \sim \mathcal{U}(0, 1)$ denotes a scalar sampled from the uniform distribution $\mathcal{U}(0, 1)$ for deciding the i^{th} example design. Note that choosing $r_{\text{CoT}} = 1$ gives us the special case of all *CoT examples* in-context, whereas $r_{\text{CoT}} = 0$ gives a sequence with all *standard examples*.

3.1 Model Training and Evaluation

Models. We follow the setup of Kothapalli et al. (2025) and train 3 custom transformer models based on the Llama-3 (Dubey et al., 2024) architecture with depth $l \in \{4, 8, 12\}$ (Table 1). For notational consistency, we denote a model with L layers as TF- L throughout the paper. Appendix C presents details about the experiment settings.

Training objective. We employ the Cross-Entropy (CE) based next-token prediction loss with masking for training the TF models (Kothapalli et al., 2025). The CE loss is computed only on the 4 tokens: $t_{\text{ans_start}}, t_{\text{ans_end}}, y_C^{(i)}$ and t_{eos} per *standard example*. Similarly, the loss is computed only on the $C + 5$ tokens per *CoT example*, i.e., $t_{\text{think_start}}, t_{\text{think_end}}, t_{\text{ans_start}}$, all the *chain tokens* $y^{(i)}, t_{\text{ans_end}}$, and t_{eos} .

Evaluation prompts. Considering a test function $\tilde{f} \in \mathcal{F}$ and the query input tokens $\tilde{\mathbf{x}} = (\tilde{x}_1, \dots, \tilde{x}_N) \in \mathcal{V}_{\text{normal}}^N$, the evaluation prompt $\tilde{\mathbf{p}}^K(\tilde{f}, r_{\text{CoT}})$ is defined as:

$$\tilde{\mathbf{x}}' = (t_{\text{inp_start}}, \tilde{\mathbf{x}}, t_{\text{inp_end}}), \quad (5)$$

$$\tilde{\mathbf{p}}^K(\tilde{f}, r_{\text{CoT}}) := \left(\mathbf{p}^{K-1}(\tilde{f}, r_{\text{CoT}}), \tilde{\mathbf{x}}' \right).$$

Forcing strategies. Incorporating special tokens in the prompt design allows us to measure the performance of the trained TF models using the following 3 approaches: (1) ‘Force Think’, (2) ‘Force Answer’ and (3) ‘No Forcing’. Consider C_{eos} as the number of tokens generated by the model, and TF(\cdot) as greedy auto-regressive single token generation. The recursive formulation $\hat{y}_o, \forall o \leq C_{\text{eos}}$ with the ‘Force Think’ strategy is as follows:

$$\text{TF} \left(\underbrace{\tilde{\mathbf{p}}^K(\tilde{f}, r_{\text{CoT}})}_{\text{evaluation seq}}, \underbrace{t_{\text{think_start}}}_{\text{force token}}, \underbrace{\hat{y}_1, \dots, \hat{y}_{o-1}}_{\text{previous step outputs}} \right). \quad (6)$$

By replacing $t_{\text{think_start}}$ with $t_{\text{ans_start}}$ in (6), we condition the model to directly provide the final answer, whereas ‘No Forcing’ completely removes the *force token* suffix and allows the model to choose between thinking and no thinking modes.

Measuring accuracy. We denote the ground truth chain tokens for $\tilde{\mathbf{x}}$ in the evaluation prompt as $\tilde{\mathbf{y}} = (\tilde{y}_1, \dots, \tilde{y}_C) \in \mathcal{V}_{\text{normal}}^C$, which are generated by the formulation given in (1) using $\tilde{f} \in \mathcal{F}$. Considering the model’s output sequence as: $\hat{\mathbf{y}} = (\hat{y}_1, \hat{y}_2, \dots, \hat{y}_{C_{\text{eos}}})$, we search for the following pattern: $(t_{\text{ans_start}}, \hat{y}_k, t_{\text{ans_end}})$ in $\hat{\mathbf{y}}$ and treat \hat{y}_k as the predicted answer \hat{y}_{pred} . If the pattern does not exist then we set $\mathbb{I}_{\hat{y}_{\text{pred}}=\tilde{y}_C} = 0$ since the model failed to follow the output format². Given \tilde{T} evaluation sequences, we measure the overall accuracy as $\frac{1}{\tilde{T}} \sum_{t=1}^{\tilde{T}} \mathbb{I}_{\hat{y}_{\text{pred}}=\tilde{y}_C}$.

3.2 Choices of \mathcal{G}, \mathcal{H} in CoT-ICL Lab

The *CoT-ICL Lab* framework is designed with the above formalization and employs the following function classes \mathcal{G}, \mathcal{H} to generate chain tokens (1).

DAG representation of causal dependencies. \mathcal{G} is a class of topologically sorted DAGs whose structure is determined by the choice of (N, M, C) . See Figure 1, which represents a DAG sampled from $\mathcal{G}(N = 3, M = 2, C = 3)$.

Process data embeddings with MLPs. \mathcal{H} is a class of MLPs of depth l with the activation ϕ ; denoted as $\mathcal{H}(l, \phi)$ ³. We maintain a TokenProcessorCache of finite MLPs with random weights and sample from it accordingly. For a given value of C , we sample C MLPs from this cache (one for each chain token) and use them to generate the chain tokens of all K examples within that sequence.

Remark. As the DAG and MLPs are unique to every sequence, one can intuitively think of meta-training as teaching the model to figure out the underlying DAG and approximate the MLP transformations solely from the in-context examples.

4 The CoT-ICL Lab-2.0

Following the formalization in the above section and the design choices for \mathcal{G}, \mathcal{H} as per *CoT-ICL Lab*, we (1) incorporate special tokens, (2) diversify the sequences in a dataset by randomly choosing N, M, C per sequence from a list of choices, and (3) modulating the mix of *CoT/standard examples* based on CoT-Recipes. The algorithms to generate

²This is similar to measuring the instruction following capability of LLMs where the output should be formatted in a certain fashion based on the in-context examples.

³We choose $l = 1, \phi = \text{LeakyReLU}$ based on the token distribution analysis of sequences by Kothapalli et al. (2025).

the entire dataset based on these techniques are formalized in Appendix A.

4.1 Ext 1: Special Tokens

We extend the in-context examples with special tokens as formulated by (2) for a *standard example* and (3) for a *CoT example*. Without delimiters, a *CoT example* is formulated as $e_{CoT}(f) = (\mathbf{x}, \mathbf{y})$, and a *standard example* as $e(f) = (\mathbf{x}, y_C)$, where the loss is computed only on \mathbf{y} or y_C respectively. Such a design has the following limitation that if one were to mix *CoT* and *standard examples* in a single sequence, the model cannot differentiate between the first intermediate token of a *CoT example* from the answer token of a *standard example*.

4.2 Ext 2: Diversification with N, M, C

To enhance the diversity of the dataset comprising T sequences, we introduce variability through randomized sampling of the parameters N, M, C . Specifically, we define discrete sets of available choices: \mathbf{N}, \mathbf{M} , and \mathbf{C} for each sequence, and sample one value from each set: $N \sim \mathbf{N}, M \sim \mathbf{M}, C \sim \mathbf{C}$. These values are used to construct a DAG using $\mathcal{G}(N, M, C)$ for all the K examples in the sequence. Such a sequence design is not possible in the older *CoT-ICL Lab* design as the same (N, M, C) are used for all T sequences.

4.3 Ext 3: Diversification with CoT-Recipe

By leveraging the richer sequence design of (4), we define CoT-Recipe, a parameterized approach to systematically assign $r_{CoT}^{(j)}$ based on the sequence index $j \in [0, T-1]$. This formulation allows us to control the expected proportion of *CoT* versus *standard examples* in context. Formally, a CoT-Recipe is a partial power-law function with offset as:

$$\text{CoT-Recipe}(\alpha, a, b)(u) = a \cdot u^\alpha + b, \quad (7)$$

where $\alpha \in \mathbb{R}_{\geq 0}$ governs the shape (e.g., linear, sub-linear), while $a, b \in \mathbb{R}$ scale and shift the curve respectively (see the illustration in Figure 1 and Appendix B.2 for calculations on the expected token counts). Using this formalization, we set $r_{CoT}^{(j)}$ for a sequence/prompt with index $j \in [0, T-1]$ as $r_{CoT}^{(j)} = \text{CoT-Recipe}(\alpha, a, b)(j/T)$, which gives:

$$r_{CoT}^{(j)} = a \cdot (j/T)^\alpha + b. \quad (8)$$

Design Rationale. Although $r_{CoT}^{(j)}$ is deterministically computed based on sequence index $j \in$

$[0, T-1]$, we note that the sequences are randomly shuffled prior to training. The formulation in (7) allows us to have a clear mental model for designing the CoT-Recipe partial function, while explicitly avoiding a curriculum-like effect during training.

5 To Think or Not To Think?

In this section, we systematically meta-train the models with varying CoT-Recipe parameter α , and evaluate them on datasets with varying fractions of *CoT examples*. In essence, we aim to understand which choices of α can lead to effective meta-training and allow the models to solve novel tasks even with limited *CoT examples*.

Meta-training Setup. We choose $|\mathcal{V}| = 1024, d = 10, \mathbf{N} = \mathbf{M} = \mathbf{C} = \{4\}, K = 40$ and vary $\alpha \in \{0, 0.5, 1, 2, \infty\}$ with $a = 1, b = 0$, for creating the training ($T = 64 \times 10^5$) datasets.

Evaluation Setup. We create a separate evaluation dataset $\tilde{\mathcal{D}}$ with $|\tilde{\mathcal{D}}| = \tilde{T} = 10^4$ sequences using $\tilde{\mathbf{N}} = \tilde{\mathbf{M}} = \tilde{\mathbf{C}} = \{4\}, \tilde{K} = 40$ and $r_{CoT}^{(j)} = 1, \forall j \in [0, \tilde{T}-1]$. Since $r_{CoT}^{(j)} = 1$, all the evaluation sequences contain $K-1$ *CoT examples* along with the query input tokens. Next, we transform $\tilde{\mathcal{D}}$ by randomly choosing K' *CoT examples* (i.i.d) per sequence and dropping the intermediate tokens along with the special tokens $t_{\text{think_start}}$ and $t_{\text{think_end}}$. Thus converting them into K' *standard examples*.

Limiting CoT supervision by increasing K' . By varying $K' \in \{0, 10, 20, 30, 39\}$ and applying different forcing strategies, we evaluate the TF model’s performance when CoT supervision is limited, i.e as K' increases, the *standard examples* outnumber the *CoT examples*.

5.1 Over-Reliance on CoT/standard Examples

Lack of modulation during meta-training leads to over-reliance on CoT/standard examples. As the TF models trained using $\alpha = 0$ have never encountered *standard examples* in a training sequence, they tend to rely on the *CoT examples* for generating the outputs. Observe from Figure 2 that such a recipe leads to a gradual reduction of accuracy across all model sizes and forcing strategies as K' increases. In particular, the (largest) TF-12 model evaluated with the ‘Force Think’ strategy exhibits a reduction in accuracy from ≈ 0.93 (when $K' = 0$) to ≈ 0.18 (when $K' = 39$). On the contrary, models trained using $\alpha = \infty$ are

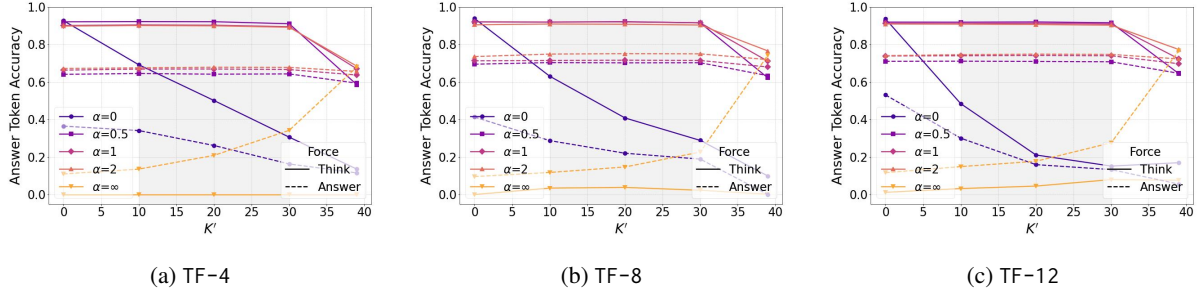


Figure 2: accuracy of models trained with varying α , $\mathbf{N} = \mathbf{M} = \mathbf{C} = \{4\}$ and evaluated on datasets with $\tilde{\mathbf{N}} = \tilde{\mathbf{M}} = \tilde{\mathbf{C}} = \{4\}$. Here K' indicates the number of *standard examples* in test prompts.

incapable of producing the intermediate chain tokens. Thus, the ‘Force Think’ strategy leads to ≈ 0 accuracy across all models. However, the ‘Force Answer’ strategy results in a gradual increase in accuracy as K' increases and the fraction of *CoT examples* reduces. Thus indicating an over-reliance on *standard examples*. Such behavior is not desirable since we want the model to generalize the thinking process.

Careful selection of α facilitates thinking even without *CoT examples*. Training datasets created with $\alpha = \{0.5, 1, 2\}$ modulate the proportion of *CoT examples* in the sequences and thus prevent the model’s over-reliance as seen with $\alpha = \{0, \infty\}$. Especially, observe from Figure 2 that the ‘Force Think’ strategy can be applied to all models for maintaining a high accuracy even as the proportion of *standard examples* increases (i.e., K' increases). Thus, avoiding the gradual collapse as observed with $\alpha = 0$. Surprisingly, notice that the models trained using $\alpha = 2$ can leverage the ‘Force Think’ strategy even when there are no *CoT examples* in-context (i.e., $K' = 39$), and perform on-par with the counterparts that were trained using $\alpha = \infty$ and evaluated with the ‘Force Answer’ strategy. This observation signifies that with careful selection of the recipes, one can train models to be good at thinking even when there are no *CoT examples* available in-context. In particular, the accuracy improvements over $\alpha = 0$ tend to be larger as K' increases and can be greater than 300% with TF-12 for $K' = 39$ i.e., a increase from ≈ 0.18 with $\alpha = 0$ to ≈ 0.78 with $\alpha = 2$.

6 Task Diversity and Length Generalization

In the previous section, we have seen that the CoT-Recipe formalization allows us to modulate the mix of CoT examples and facilitate reasoning

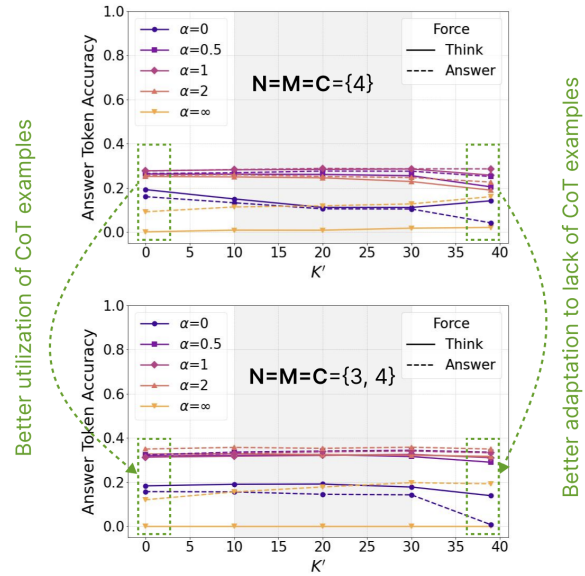


Figure 3: Input length generalization of TF-12 models when tested with $\tilde{\mathbf{N}} = \{5\}, \tilde{\mathbf{M}} = \tilde{\mathbf{C}} = \{4\}$.

control in the models. As a next step, we show that increasing the diversity of training data via $\mathbf{N}, \mathbf{M}, \mathbf{C}$ can aid the models to generalize to out-of-distribution (OOD) settings with longer inputs.

Setup. We train additional TF models on new diverse datasets created using $\mathbf{N} = \mathbf{M} = \mathbf{C} = \{3, 4\}, K = 40$ (with rest of the parameters for CoT-Recipe, \mathcal{G}, \mathcal{H} being the same as Section 5). For evaluation, we consider $\tilde{\mathbf{N}} = \{5\}, \tilde{\mathbf{M}} = \tilde{\mathbf{C}} = \{4\}, \tilde{K} = 40$ and employ the same process as Section 5 to create the evaluation datasets.

Diversity with $\mathbf{N}, \mathbf{M}, \mathbf{C}$ improves length generalization. Figure 3 illustrates that the TF-12 models leverage the diversity of $\mathbf{N} = \mathbf{M} = \mathbf{C} = \{3, 4\}$ and attain a peak evaluation accuracy of ≈ 0.38 , when compared to ≈ 0.28 with models trained using $\mathbf{N} = \mathbf{M} = \mathbf{C} = \{4\}$. Furthermore, models trained with $\alpha = \{0.5, 1, 2\}$ consistently outperform the $\alpha = \{0, \infty\}$ cases across all sizes

(see Figure 9 for TF-4 and Figure 10 for TF-8). More importantly, the ‘Force Think’ strategy with $\alpha = 0$ models on $K' = 0$ is not as effective as $\alpha = \{0.5, 1, 2\}$, unlike the in-domain generalization setting (see Figure 2).

Forcing strategies and OOD tasks. Surprisingly, Figure 3 also shows that for any K' , the peak accuracy with ‘Force Answer’ strategy across all α is comparable and sometimes even higher than the ‘Force Think’ strategy. This is a failure mode where the model is unable to generate the required number of thinking steps (5 in this case) even with the ‘Force Think’ strategy and fails to arrive at the right answers. In terms of recipes, $\alpha = 2$ tends to be the best choice for TF-12. While we observe that the gap in accuracy for both the strategies is quite narrow across model sizes and K' in Figure 3, Figure 9, and Figure 10, the gap was observed to be relatively wider for the in-domain generalization setting (see Figure 2). Thus, it sheds light on the limitations of enforcing thinking behavior and the brittleness of OOD generalization.

7 Symbolic Reasoning with LLMs

Our analysis above highlighted the importance of CoT-Recipe for reasoning with abstract tokens. To verify if these insights can be transferred to pretrained LLMs, we leverage the design patterns of *CoT-ICL Lab-2.0* to create a fully interpretable symbolic reasoning dataset called CIL-LangSym. In particular, we aim to understand (1) if CoT-Recipe parameters obtained in the previous sections can still be effective when each intermediate step and answer spans multiple tokens and (2) how forcing strategies affect length generalization when the model is pre-trained on a much larger and diverse natural language data.

7.1 Data Generation

Random words with ASCII lowercase. We adhere to the setup in Section 3 and generate N input words per in-context example, each comprising of W ASCII lowercase characters (a-z). In essence, we transition from N abstract tokens in *CoT-ICL Lab-2.0* to N words per in-context example.

Causal structure via DAGs. Similar to *CoT-ICL Lab-2.0*, we consider the function class \mathcal{G} of topologically sorted DAGs to implant the causal dependencies between words.

String processing function. Unlike the token processing function class \mathcal{H} in *CoT-ICL Lab-2.0* that relied on the data embeddings \mathbf{E}_{data} and MLPs to generate the abstract chain tokens, we employ a string processing function s based on string slicing and character offset operations for CIL-LangSym. Formally, the function s takes the M filtered words (from \mathcal{G}) as inputs and outputs a single *intermediate word* (Figure 4). See Appendix E for the algorithmic description of these design aspects and Figure 45 for an example prompt.

Remark. Notice that the task description does not reveal the underlying string transformation, and the model is required to figure out the task solely from the in-context *CoT/standard examples*.

7.2 Experiments

Setup. The CIL-LangSym dataset is created with the following parameters: $\mathbf{N} = \tilde{\mathbf{N}} = \{4\}$, $\mathbf{M} = \tilde{\mathbf{M}} = \{2\}$, $\mathbf{C} = \tilde{\mathbf{C}} = \{3\}$, $K = \tilde{K} = 40$ for training ($T = 1000$), and evaluation ($\tilde{T} = 10,000$). We consider the Qwen-2.5 series models to highlight the difficulty of the tasks and to study the role of CoT-Recipes and forcing strategies. We allocate a budget of 1000 and 100 tokens for the ‘Force Think’ and ‘Force Answer’ strategies respectively.

Insufficiency of pre-trained knowledge. We evaluate 8 Qwen-2.5 series models (Yang et al., 2024): 0.5B-Instruct, 1.5B, 1.5B-Instruct, Math-1.5B-Instruct, 7B, 7B-Instruct along with the DeepSeek distilled reasoning models: DeepSeek-R1-Distill-Qwen-1.5B, DeepSeek-R1-Distill-Qwen-7B on evaluation datasets for $K' = \{0, 10, 20, 30, 39\}$ and temperature=0.6. The evaluation datasets are prepared using the same approach as Section 5. By relying solely on CoT-ICL prompting, we noticed that all the 8 models scored an accuracy of 0 across all K' and both the forcing strategies. This baseline establishes that the models cannot leverage their pre-trained knowledge for this symbolic reasoning task and ensures that the performance improvements are largely attributed to meta-training based on the CoT-Recipe⁴.

CoT-Recipe facilitates reasoning control even in pretrained LLMs. We use CoT-Recipe with $\alpha = \{0, 0.5, 1, 2, \infty\}$ for SFT (meta-training) of 0.5B-Instruct, 1.5B-Instruct,

⁴Results were also consistent across other temperature values ranging from 0 to 1 with 0.1 increments.

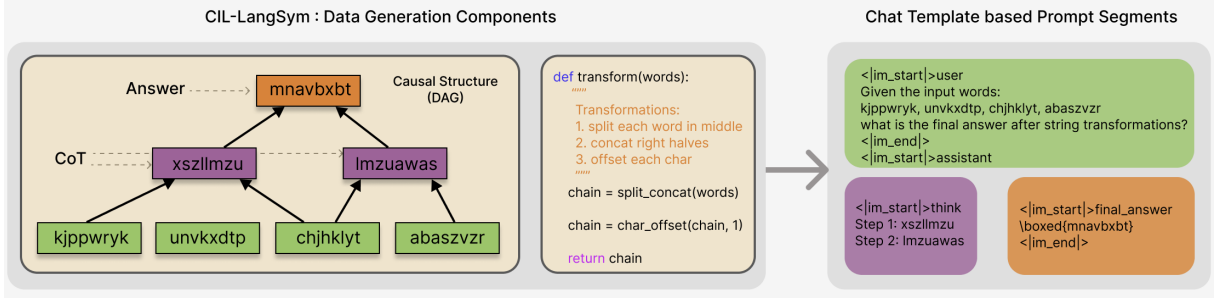


Figure 4: Chat template of a *CoT/standard example* in CIL-LangSym based on the Qwen-2.5-1.5B-Instruct tokenizer. Given $N = 4, M = 2, C = 3$ and word length $W = 8$, the DAG determines the ground truth causal dependencies, and the transform function illustrates the string processing of the M parent words. We apply the chat template to differentiate the question, thinking, and final answer segments of the examples and also ensure that the task description does not reveal the underlying string transformation in natural language.




Figure 5: accuracy of models trained with varying α , $\mathbf{N} = \{4\}, \mathbf{M} = \{2\}, \mathbf{C} = \{3\}$ and evaluated on datasets with $\tilde{\mathbf{N}} = \{4\}, \tilde{\mathbf{M}} = \{2\}, \tilde{\mathbf{C}} = \{3\}$.

and 7B-Instruct models to focus solely on the model size. Given that we use only 1000 prompts, Figure 5 shows a clear benefit of model size as the peak accuracy attained by 0.5B-Instruct is only 0.004 whereas 1.5B-Instruct, 7B-Instruct can reach up-to 0.15, and 0.67 respectively. Given a large enough model such as 7B-Instruct, notice from Figure 5c that as long as there are CoT examples available in-context, the CoT-Recipes with $\alpha \neq \infty$ facilitate the model to leverage the ‘Force Think’ strategy. Similar to the observations with *CoT-ICL Lab-2.0*, these models showcase an over-reliance on the *CoT examples* and exhibit a reversal in trend of accuracy as α increases from 0 to ∞ for $K' = 39$. In these scenarios with limited CoT supervision, if one were to employ the ‘Force Think’ strategy, then choosing $\alpha = 2$ can result in accuracy improvements of over 130% when $K' = 39$ (consistent with the observations from *CoT-ICL Lab-2.0* in Section 5). We provide guidance on choosing α in Section 8, an example prompt in Figure 45 and model outputs in Table 2.

Remark. Since these LLMs are pre-trained on a large corpus on natural language data, notice from Figure 5c that the accuracy of the $\alpha = 0$ model

with ‘Force Answer’ strategy exhibits an increasing trend rather than a deteriorating one as observed with *CoT-ICL Lab-2.0*. This behavior is unique to the pre-trained models and understanding the role of model size and pre-training datasets can be an interesting avenue for future research.

Correct thinking steps do not necessarily imply a correct final answer. We break down the accuracy of the Qwen-2.5-7B-Instruct model trained with $\alpha = 0$ by analyzing the correctness of the intermediate reasoning steps for the evaluation prompts. Since $\tilde{\mathbf{C}} = \{3\}$, there will be 2 intermediate steps. Figure 6 considers the scenario with ‘Force Think’ strategy and $K' = 0$ (i.e., all examples have CoT) to highlight that: out of the 6752 correctly predicted prompts, around 80% of them have both the intermediate steps to be correct. Whereas, out of the 3428 prompts with wrong final answers, $\approx 30\%$ of them have both the steps to be correct and $\approx 60\%$ of them have at least one step to be correct. We also present breakdowns for other models and α , based on the inclusion of these intermediate steps in the ground truth DAG in Appendix E.1. In summary, our results indicate that the output of the second thinking step (denoted

as Step 2) is relatively more important than Step 1 because of the underlying causal structure. Thus, incorrect Step 2 predictions lead to relatively more errors in the final answer than Step 1.




Figure 6: Qwen-2.5-7B-Instruct trained with $\alpha = 0$ achieves 0.6752 final answer accuracy with the ‘Force Think’ strategy for $K' = 0$. Each grid in the plots indicate the percentage of prompts (out of 6752 in (a) and 3248 in (b)) for which the Step 1 and Step 2 predictions by the model were correct/incorrect.

Input length generalization. Figure 7 illustrates that the accuracy of the SFT’ed Qwen-2.5-7B-Instruct model tends to drop as the number of input words per example in the evaluation sequences is increased. This observation is consistent with the results of *CoT-ICL Lab-2.0*. However, a key difference is that the accuracy values on CIL-LangSym are relatively closer to the performance on the in-domain task (Figure 5c), with ‘Force Think’ strategy outperforming ‘Force Answer’ for attaining the peak accuracy across α and K' . A similar behavior can be observed for the chain length generalization experiments in Figure 15. This indicates that although pre-training is not sufficient for solving the tasks, the diversity of the data aids in OOD generalization. Thus, highlighting the key difference between training from scratch for *CoT-ICL Lab-2.0* tasks and pre-training on large natural language corpora.




Figure 7: Qwen-2.5-7B-Instruct: Input length generalization with train: $\mathbf{N} = \{4\}$, $\mathbf{M} = \{2\}$, $\mathbf{C} = \{3\}$ and test: $\tilde{\mathbf{M}} = \{2\}$, $\tilde{\mathbf{C}} = \{3\}$ by varying $\tilde{\mathbf{N}}$.

8 Guidance on Choosing α

Throughout the paper, $\alpha \in \{0, 0.5, 1, 2, \infty\}$ was varied to simulate constant, linear, sub-linear, and super-linear growth of r_{cot} in the dataset. Although our coverage of α is limited to these 5 values, we observe clear patterns in our experiments that guide their choice for meta-training.

CoT rich scenarios. When the evaluation prompts are expected to have a sufficient number of CoT examples, then it is recommended to choose $\alpha = \{0, 0.5\}$ and employ the ‘Force Think’ strategy to maximize the performance.

CoT poor scenarios. If minimal CoT examples are expected to be available during inference, then it is recommended to choose $\alpha = \{2, \infty\}$ and employ the ‘Force Answer’ strategy so that the model is not meta-trained on sequences that causes it to over-rely on CoT supervision.

A single model that can reason and answer directly. Finally, if a single model is expected to perform well in all kinds of evaluation scenarios, i.e., with and without CoT examples as well as being forced to think or answer directly, then it is recommended to meta-train with $\alpha \in \{0.5, 2\}$ to achieve the best tradeoffs. We underscore that this is not a universal selection criterion and only aims to narrow down the optimal value selection. We also analyze the role of α on computational overheads in Appendix B.2.

9 Conclusion

This work introduced *CoT-ICL Lab-2.0*, a data generation framework for abstract reasoning tasks to design effective meta-training techniques for transformer models. By incorporating special abstract tokens and modulating the mix of *CoT/standard examples* via CoT-Recipe, we systematically showcased the importance of training task diversity and forcing strategies for reasoning control in transformers. By verifying the effectiveness of these insights on practical LLMs for novel symbolic reasoning tasks, we hope to encourage the formalization of the data-mixing aspects of meta-training with broader domain-specific tasks.

10 Limitations

The *CoT-ICL Lab-2.0* framework is designed to generate diverse sequences of abstract tokens that are devoid of semantics. In this context, we note

that eliciting chain-of-thought does not exactly resemble the case with NLP datasets (especially in the zero-shot settings), since the token distribution, input and chain lengths may vary significantly with answer tokens exceeding just a single token. Although the CIL-LangSym addresses some of these concerns, it can be treated as a domain-specific dataset, and one should carefully consider the scenarios and training stages in which these insights can apply to real-world tasks with math/code, etc.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Yu Bai, Fan Chen, Huan Wang, Caiming Xiong, and Song Mei. 2023. Transformers as statisticians: Provable in-context learning with in-context algorithm selection. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Stella Biderman, Hailey Schoelkopf, Lintang Sutawika, Leo Gao, Jonathan Tow, Baber Abbasi, Alham Fikri Aji, Pawan Sasanka Ammanamanchi, Sidney Black, Jordan Clive, and 1 others. 2024. Lessons from the trenches on reproducible evaluation of language models. *arXiv preprint arXiv:2405.14782*.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901.
- Xiangyu Chang, Yingcong Li, Muti Kara, Samet Oymak, and Amit Roy-Chowdhury. 2025. Provable benefits of task-specific prompts for in-context learning. In *International Conference on Artificial Intelligence and Statistics*, pages 1558–1566. PMLR.
- Yanda Chen, Ruiqi Zhong, Sheng Zha, George Karypis, and He He. 2022. Meta-learning via language model in-context tuning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 719–730.
- Zheng Chu, Jingchang Chen, Qianglong Chen, Weijiang Yu, Tao He, Haotian Wang, Weihua Peng, Ming Liu, Bing Qin, and Ting Liu. 2024. Navigate through enigmatic labyrinth a survey of chain of thought reasoning: Advances, frontiers and future. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1173–1203.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Baobao Chang, and 1 others. 2024. A survey on in-context learning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1107–1128.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Hamed Firooz, Maziar Sanjabi, Adrian Englhardt, Aman Gupta, Ben Levine, Dre Olgiati, Gungor Polatkan, Iuliia Melnychuk, Karthik Ramgopal, Kirill Talanin, and 1 others. 2025. 360brew: A decoder-only foundation model for personalized ranking and recommendation. *arXiv preprint arXiv:2501.16450*.
- Shivam Garg, Dimitris Tsipras, Percy S Liang, and Gregory Valiant. 2022. What can transformers learn in-context? a case study of simple function classes. *Advances in Neural Information Processing Systems*, 35:30583–30598.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the MATH dataset. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Namgyu Ho, Laura Schmid, and Seyoung Yun. 2023. Large language models are reasoning teachers. In *61st Annual Meeting of the Association for Computational Linguistics, ACL 2023*, pages 14852–14882. Association for Computational Linguistics (ACL).
- Pin-Lun Hsu, Yun Dai, Vignesh Kothapalli, Qingquan Song, Shao Tang, Siyu Zhu, Steven Shimizu, Shivam Sahni, Haowen Ning, Yanning Chen, and Zhipeng Wang. 2025. Liger-kernel: Efficient triton kernels for LLM training. In *Championing Open-source Development in ML Workshop @ ICML25*.
- Jiaxin Huang, Shixiang Shane Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei Han. 2023. Large language models can self-improve. In *2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023*, pages 1051–1068. Association for Computational Linguistics (ACL).
- Seungone Kim, Se June Joo, Doyoung Kim, Joel Jang, Seonghyeon Ye, Jamin Shin, and Minjoon Seo. 2023. The cot collection: Improving zero-shot and few-shot learning of language models via chain-of-thought fine-tuning. In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, pages 22113–22124.
- Vignesh Kothapalli, Hamed Firooz, and Maziar Sanjabi. 2025. CoT-ICL lab: A synthetic framework for

- studying chain-of-thought learning from in-context demonstrations. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14620–14642, Vienna, Austria. Association for Computational Linguistics.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, pages 611–626.
- Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, and 1 others. 2024. Tulu 3: Pushing frontiers in open language model post-training. *arXiv preprint arXiv:2411.15124*.
- Liunian Harold Li, Jack Hessel, Youngjae Yu, Xiang Ren, Kai-Wei Chang, and Yejin Choi. 2023a. Symbolic chain-of-thought distillation: Small models can also think step-by-step. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Yingcong Li, Kartik Sreenivasan, Angeliki Gianou, Dimitris Papailiopoulos, and Samet Oymak. 2023b. Dissecting chain-of-thought: Compositionality through in-context filtering and learning. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Sewon Min, Mike Lewis, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2022a. Metaicl: Learning to learn in context. In *Proceedings of the 2022 conference of the North American chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2791–2809.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022b. Rethinking the role of demonstrations: What makes in-context learning work? In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11048–11064.
- Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, Charles Sutton, and Augustus Odena. 2022. Show your work: Scratchpads for intermediate computation with language models. In *Deep Learning for Code Workshop*.
- Kazusato Oko, Yujin Song, Taiji Suzuki, and Denny Wu. 2024. Pretrained transformer efficiently learns low-dimensional target functions in-context. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, and 1 others. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.
- Boshi Wang, Sewon Min, Xiang Deng, Jiaming Shen, You Wu, Luke Zettlemoyer, and Huan Sun. 2023. Towards understanding chain-of-thought prompting: An empirical study of what matters. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2717–2739.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, pages 24830–24842.
- Qwen An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxin Yang, Jingren Zhou, Junyang Lin, and 25 others. 2024. Qwen2.5 technical report. *ArXiv*, abs/2412.15115.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *International conference on machine learning*, pages 12697–12706. PMLR.
- Yuxiang Zhou, Jiazheng Li, Yanzheng Xiang, Hanqi Yan, Lin Gui, and Yulan He. 2024. The mystery of in-context learning: A comprehensive survey on interpretation and analysis. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 14365–14378.

A The CoT-ICL Lab-2.0 Dataset Generation Algorithms

We formalize the design aspects of *CoT-ICL Lab-2.0* and present the dataset generation process in Algorithm 1. A single tokenized sequence in this dataset is generated using Algorithm 2, especially by isolating the input and answering parts of an in-context example using special tokens. Algorithm 3 uses the data embeddings \mathbf{E}_{data} to generate a single chain token for every in-context example.

Algorithm 1 Generate dataset with T sequences

Require: Parameter choices $\mathbf{N}, \mathbf{M}, \mathbf{C}, K$, the CoT-Recipe parameters α, a, b , and size T .

- 1: Initialize empty dataset $\mathbf{D} = []$
- 2: **for** $j = 1$ to T **do**
- 3: $r_{CoT}^{(j)} = \text{CoT-Recipe}(\alpha, a, b)(j/T)$
- 4: $\mathbf{p} = \text{Algorithm 2}(\mathbf{N}, \mathbf{M}, \mathbf{C}, K, r_{CoT}^{(j)})$.
- 5: $\mathbf{D}.\text{append}(\mathbf{p})$
- 6: **end for**
- 7: **return** \mathbf{D}

Algorithm 2 Single sequence generation with index j in the dataset.

Require: Parameter choices $\mathbf{N}, \mathbf{M}, \mathbf{C}, K$, and the CoT probability parameter $r_{CoT}^{(j)}$.

- 1: sample $N \sim \mathbf{N}, M \sim \mathbf{M}, C \sim \mathbf{C}$.
- 2: Limit $M = \min(M, N)$.
- 3: Initialize the sequence $\mathbf{p} = [t_{\text{bos}}]$
- 4: **for** $k = 1$ to K **do**
- 5: Initialize empty input token sequence \mathbf{x} , chain token sequence \mathbf{y} .
- 6: **for** $i = 1$ to N **do**
- 7: $\mathbf{x}[i] \stackrel{\text{i.i.d.}}{\sim} \mathcal{V}_{\text{normal}}$
- 8: **end for**
- 9: $\mathbf{t} = \mathbf{x}.\text{clone}()$
- 10: **for** $c = 1$ to C **do**
- 11: $\text{parent_tokens} = \text{rand.choice}(\mathbf{t}, M)$
- 12: $\mathbf{y}[c] = \text{Algorithm 3}(\text{parent_tokens})$
- 13: $\mathbf{t}.\text{append}(\mathbf{y}[c])$
- 14: **end for**
- 15: $\mathbf{p}.\text{extend}([t_{\text{inp_start}}, \mathbf{x}, t_{\text{inp_end}}])$
- 16: **if** $r_{CoT}^{(j)} \geq \mathcal{U}(0, 1)$ **then**
- 17: $\mathbf{p}.\text{extend}([t_{\text{think_start}}, \mathbf{y}_{:C-1}, t_{\text{think_end}}])$
- 18: **end if**
- 19: $\mathbf{p}.\text{extend}([t_{\text{ans_start}}, y_C, t_{\text{ans_end}}, t_{\text{eos}}])$
- 20: **end for**
- 21: **return** \mathbf{p}

Algorithm 3 Single chain token y_c generation

Require: M row indices of \mathbf{E}_{data} corresponding to parent_tokens.

- 1: MLP $h_c \in \text{TokenProcessorCache}(\mathcal{H}(l = 1, \phi = \text{LeakyReLU}))$
- 2: **for** $i = 1$ to M **do**
- 3: $\mathbf{h}^i \leftarrow h_c(\mathbf{E}_{\text{data}}[\text{parent_tokens}[i]])$
- 4: **end for**
- 5: $\mathbf{h}_{\text{act}} \leftarrow \phi \left(\frac{1}{M} \sum_{i=1}^M \mathbf{h}^i \right)$
- 6: $y_c \leftarrow \text{argmax}(\mathbf{E}_{\text{data}} \mathbf{h}_{\text{act}})$
- 7: **return** y_c

B Task Difficulty with CoT-Recipe and Training Token Estimates

B.1 Fine-grained control of task difficulty

Considering the meta-training setup of Section 5, we use $\tilde{\mathbf{N}} = \tilde{\mathbf{M}} = \tilde{\mathbf{C}} = \{4\}, \tilde{K} = 40$ for the evaluation ($\tilde{T} = 10^4$) datasets corresponding to the same α and apply the ‘No Forcing’ strategy to measure accuracy. As α increases, the sequences tend to contain fewer proportion of *CoT examples* in-context (see Figure 1) and in-turn lead to consistently lower accuracy values across model sizes (see Figure 8). In particular, when $\alpha = 0$, the TF models leverage the intermediate/thinking tokens to achieve higher accuracy, whereas $\alpha = \infty$ presents no such information and the model is forced to answer directly. Although such extreme cases were already studied in Kothapalli et al. (2025), our results highlight a fine-grained control over the difficulty of such tasks by carefully selecting the shape parameter α .

B.2 Expected Token Count in Datasets

Theorem B.1. Consider a dataset \mathcal{D} of T sequences created using the tuple N, M, C, K . Let the CoT-Recipe (7) with $a = 1, b = 0$ determine the CoT probability parameter $r_{CoT}^{(j)}, \forall j \in [0, T-1]$ as follows: $r_{CoT}^{(j)} = \left(\frac{j}{T}\right)^\alpha$. Then the expected number of tokens $\mathbb{E}[|\mathcal{D}|_{\text{Tokens}}]$ is:

$$\mathbb{E}[|\mathcal{D}|_{\text{Tokens}}] = T + \mathbb{E}[|\mathcal{D}|_{\text{CoT-ex}}] (N + C + 7) + (KT - \mathbb{E}[|\mathcal{D}|_{\text{CoT-ex}}]) (N + 6)$$

, where $\mathbb{E}[|\mathcal{D}|_{\text{CoT-ex}}]$ is given by:

$$\frac{K}{T^\alpha} \left(\frac{(T-1)^{\alpha+1} - 1}{\alpha + 1} + \frac{(T-1)^\alpha + 1}{2} \right)$$

Proof. Notice that a standard example (2) consists of $N + 6$ tokens, and a CoT example (3) consists of $N + C + 7$ tokens.




Figure 8: accuracy with varying α , $N = M = C = \{4\}$, and $\tilde{N} = \tilde{M} = \tilde{C} = \{4\}$.

For a sequence \mathbf{p} with K examples and CoT probability parameter $r_{\text{CoT}}^{(j)}$, the expected number of *CoT/standard examples* are:

$$\begin{aligned} \mathbb{E} [|\mathbf{p}|_{\text{CoT-ex}}] &= K \times r_{\text{CoT}}^{(j)} \\ \mathbb{E} [|\mathbf{p}|_{\text{S-ex}}] &= K \times (1 - r_{\text{CoT}}^{(j)}) \end{aligned} \quad (9)$$

By the linearity of expectations, the expected number of *CoT/standard examples* in the entire dataset is given by:

$$\begin{aligned} \mathbb{E} [|\mathcal{D}|_{\text{CoT-ex}}] &= \sum_{j=0}^{T-1} K \left(\frac{j}{T} \right)^\alpha \\ \mathbb{E} [|\mathcal{D}|_{\text{S-ex}}] &= \sum_{j=0}^{T-1} K \left(1 - \left(\frac{j}{T} \right)^\alpha \right). \end{aligned} \quad (10)$$

We use the Euler-Maclaurin approximation of $\mathbb{E} [|\mathcal{D}|_{\text{CoT-ex}}]$ to obtain:

$$\begin{aligned} \frac{K}{T^\alpha} \sum_{j=0}^{T-1} j^\alpha &= \frac{K}{T^\alpha} \sum_{j=1}^{T-1} j^\alpha \\ &\approx \frac{K}{T^\alpha} \left(\int_1^{T-1} x^\alpha dx + \frac{(T-1)^\alpha + 1}{2} \right) \\ &= \frac{K}{T^\alpha} \left(\frac{(T-1)^{\alpha+1} - 1}{\alpha + 1} + \frac{(T-1)^\alpha + 1}{2} \right) \end{aligned}$$

Since $\mathbb{E} [|\mathcal{D}|_{\text{S-ex}}] = KT - \mathbb{E} [|\mathcal{D}|_{\text{CoT-ex}}]$, we consider 1 t_{bos} token per sequence and calculate the expected tokens (excluding t_{pad}) in \mathcal{D} as:

$$\begin{aligned} \mathbb{E} [|\mathcal{D}|_{\text{Tokens}}] &= T + \mathbb{E} [|\mathcal{D}|_{\text{CoT-ex}}] (N + C + 7) \\ &\quad + \mathbb{E} [|\mathcal{D}|_{\text{S-ex}}] (N + 6). \end{aligned}$$

□

Remark. Based on Theorem B.1 and $\alpha = 0$, we get $\mathbb{E} [|\mathcal{D}|_{\text{CoT-ex}}] = KT$ and $\mathbb{E} [|\mathcal{D}|_{\text{Tokens}}] = T + KT(N + C + 7)$. Whereas $\alpha = \infty$ results in $\mathbb{E} [|\mathcal{D}|_{\text{CoT-ex}}] = 0$ and $\mathbb{E} [|\mathcal{D}|_{\text{Tokens}}] = T + KT(N + 6)$. By substituting: $N = C = 4$, $K =$

40 and $T = 64 \times 10^5$ as per training experiments in Section 5, the ratio $\frac{1+K(N+C+7)}{1+K(N+6)}$ turns out to be 1.5. A numerical simulation of the ratio of expected tokens with $\alpha = 2$ and $\alpha = \infty$ turns out to be ≈ 1.16 . Thus indicating that with a slight increase in token budget to modulate the mix of *CoT/standard examples* in the training dataset, we can achieve significant improvements in reasoning control and length generalization (Section 8).

C Resources and Hyper-Parameters for Training and Inference

Model	Params w/o Embedding Layer
TF-4	243, 288, 064
TF-8	486, 574, 080
TF-12	729, 860, 096

Table 1: Model Card for the TF models.


We use 8 H100 NVIDIA GPUs for all the training experiments and employ Liger-Kernels (Hsu et al., 2025) to speed up training and vLLM (Kwon et al., 2023) for bulk inference experiments.

CoT-ICL Lab-2.0. We use a training batch size of 16 per rank and the AdamW optimizer with $\eta = 5 \times 10^{-5}$. Training runs with the larger TF-12 model take ≈ 18 hours to finish. The evaluation with vLLM for $K' = \{0, 10, 20, 30, 39\}$ and both the forcing strategies take up to 2 hours for the TF-12 model on a single H100 GPU.


CIL-LangSym. The SFT experiments with the Qwen-2.5 series LLMs use a learning rate of $\eta = 10^{-5}$, with a warm-up ratio of 0.05, a cosine learning rate scheduler and a weight decay of 10^{-4} . All experiments typically finish under 5 minutes since the training set consists only of 1000 prompts. Considering an inference token budget of 1000 for ‘Force Think’ and 100 for the ‘Force Answer’ strategy, the evaluation with vLLM for

$K' = \{0, 10, 20, 30, 39\}$ and both the strategies take up to 5 hours for the 7B-Instruct model on a single H100 GPU.

D Length Generalization Experiments with *CoT-ICL Lab-2.0*




(a) $\mathbf{N} = \mathbf{M} = \mathbf{C} = \{4\}$




(b) $\mathbf{N} = \mathbf{M} = \mathbf{C} = \{3, 4\}$

Figure 9: Input length generalization: accuracy of TF-4 models trained with varying α and $K = 40$ on evaluation datasets with longer inputs $\tilde{\mathbf{N}} = \{5\}$, $\tilde{\mathbf{M}} = \tilde{\mathbf{C}} = \{4\}$ and $\tilde{K} = 40$.

Input Length. In Section 6, we have noticed that the TF-12 model was able to leverage the diversity of $\mathbf{N} = \mathbf{M} = \mathbf{C} = \{3, 4\}$ to achieve good length generalization when $\tilde{\mathbf{N}} = \{5\}$. However, such improvements are not evident in the smaller TF-4 model as shown in Figure 9. On the other hand, the TF-8 model trained with $\alpha = 2$ shows a significant lift with both the forcing strategies (see Figure 10), with ‘Force Answer’ being the suitable strategy over ‘Force Think’. When the input length is increased to $\tilde{\mathbf{N}} = \{6\}$, the accuracy goes down even further across model sizes (see Figure 11), but we still observe that ‘Force Answer’ is a preferred strategy over ‘Force Think’.



(a) $\mathbf{N} = \mathbf{M} = \mathbf{C} = \{4\}$



(b) $\mathbf{N} = \mathbf{M} = \mathbf{C} = \{3, 4\}$

Figure 10: Input length generalization: accuracy of TF-8 models trained with varying α and $K = 40$ on evaluation datasets with longer inputs $\tilde{\mathbf{N}} = \{5\}$, $\tilde{\mathbf{M}} = \tilde{\mathbf{C}} = \{4\}$ and $\tilde{K} = 40$.

Chain Length. For chain length generalization experiments, we create the evaluation datasets by setting $\tilde{\mathbf{N}} = \{4\}$, $\tilde{\mathbf{C}} = \{5\}$ and keeping the rest of the parameter choices the same as Section 6. For models trained with $\mathbf{N} = \mathbf{M} = \mathbf{C} = \{4\}$, we observe from Figure 12 that preparing evaluation prompts with $K' = 39$ results in the best accuracy across model sizes and forcing strategies. In particular, the gap between peak accuracy at $K' = 0$ and $L' = 39$ reduces as model size increases (see also Figure 14 for the case with $\tilde{\mathbf{C}} = \{6\}$). Furthermore, the robustness of the trained models to varying K' with $\alpha = \{0.5, 1, 2\}$ increases with model size as well. Similar observations can be made for models trained with $\mathbf{N} = \mathbf{M} = \mathbf{C} = \{3, 4\}$ in Figure 13. However, unlike the length generalization case, the diversity via $\mathbf{N}, \mathbf{M}, \mathbf{C}$ does not increase the peak accuracy across α but profoundly impacts the choice of α . For instance, $\alpha = 2$ is significantly better than other choices for the TF-8 model (Figure 13b). Nonetheless, such gaps tend to minimize as model size increases (Figure 13c).




Figure 11: Input length generalization: accuracy of models trained with varying CoT-Recipe(α) and $N = M = C = \{4\}$, $K = 40$ on evaluation datasets with $\tilde{N} = \{6\}$, $\tilde{M} = \tilde{C} = \{4\}$ and $\tilde{K} = 40$.




Figure 12: Chain length generalization: accuracy of models trained with varying CoT-Recipe(α) and $N = M = C = \{4\}$, $K = 40$ on evaluation datasets with $\tilde{N} = \tilde{M} = 4$, $\tilde{C} = \{5\}$ and $\tilde{K} = 40$.




Figure 13: Chain length generalization: accuracy of models trained with varying CoT-Recipe(α) and $N = M = C = \{3, 4\}$, $K = 40$ on evaluation datasets with $\tilde{N} = \tilde{M} = 4$, $\tilde{C} = \{5\}$ and $\tilde{K} = 40$.




Figure 14: Chain length generalization: accuracy of models trained with varying CoT-Recipe(α) and $N = M = C = \{3, 4\}$, $K = 40$ on evaluation datasets with $\tilde{N} = \tilde{M} = 4$, $\tilde{C} = \{6\}$ and $\tilde{K} = 40$.

E Symbolic Reasoning with CIL-LangSym

Data generation algorithms for CIL-LangSym are similar in nature to that of *CoT-ICL Lab-2.0*, but

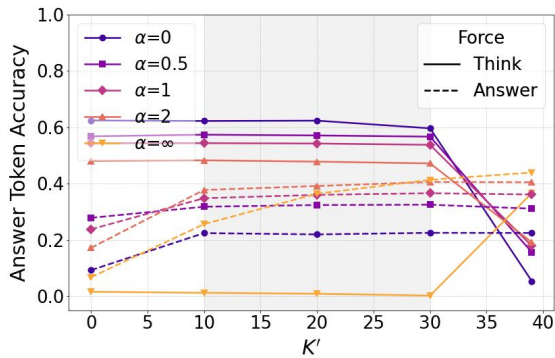
instead are applied to string inputs rather than token embeddings. Algorithm 4 presents the pseudocode for generating the entire dataset, Algorithm 5 generates a single prompt formatted using the

chat template, and Algorithm 6 generates a single chain word for an in-context example. The system_prompt along with the formatted questions after applying the Qwen-2.5-7B-Instruct tokenizer chat template is shown in Figure 45. We also illustrate the underlying DAG structure while creating an in-context example in Figure 4.


Algorithm 4 Generate dataset with T prompts

Require: Parameter choices $\mathbf{N}, \mathbf{M}, \mathbf{C}, K$, word length W , the CoT-Recipe parameters α, a, b , size T , system_prompt, question_template.

- 1: Store all input params in args.
 - 2: Initialize empty dataset $\mathbf{D} = []$
 - 3: **for** $j = 1$ to T **do**
 - 4: $r_{CoT}^{(j)} = \text{CoT-Recipe}(\alpha, a, b)(j/T)$
 - 5: $\mathbf{p} = \text{Algorithm 5}(*\text{args})$.
 - 6: $\mathbf{D}.\text{append}(\mathbf{p})$
 - 7: **end for**
 - 8: **return** \mathbf{D}
-



(a) $\tilde{\mathbf{N}} = \{4\}, \tilde{\mathbf{M}} = \{2\}, \tilde{\mathbf{C}} = \{4\}$



(b) $\tilde{\mathbf{N}} = \{4\}, \tilde{\mathbf{M}} = \{2\}, \tilde{\mathbf{C}} = \{5\}$

Figure 15: Chain length generalization: accuracy of Qwen-2.5-7B-Instruct model trained with varying α , $\mathbf{N} = \{4\}, \mathbf{M} = \{2\}, \mathbf{C} = \{3\}$ and $K = 40$ on evaluation datasets with longer chains.

Algorithm 5 Single prompt generation with index j in the dataset.

Require: Parameter choices $\mathbf{N}, \mathbf{M}, \mathbf{C}, K$, word length W , and the CoT probability $r_{CoT}^{(j)}$, system_prompt, question_template,

- 1: sample $N \sim \mathbf{N}, M \sim \mathbf{M}, C \sim \mathbf{C}$
 - 2: Limit $M = \min(M, N)$.
 - 3: Messages $\mathbf{m} = [\text{"system"}, \text{system_prompt}]$
 - 4: CHARS = list(string.ascii_lowercase)
 - 5: **for** $k = 1$ to K **do**
 - 6: Initialize empty input words list \mathbf{x} , chain words list \mathbf{y} .
 - 7: **for** $i = 1$ to N **do**
 - 8: $\mathbf{x}[i] \leftarrow "".\text{join}(\text{sample}(\text{CHARS}, W))$
 - 9: **end for**
 - 10: $\mathbf{t} = \mathbf{x}.\text{clone}()$
 - 11: **for** $c = 1$ to C **do**
 - 12: parent_words = rand.choice(\mathbf{t}, M)
 - 13: $\mathbf{y}[c] = \text{Algorithm 6}(\text{parent_tokens})$
 - 14: $\mathbf{t}.\text{append}(\mathbf{y}[c])$
 - 15: **end for**
 - 16: $\mathbf{q} = \text{fmt_q}(\text{question_template}, \mathbf{x})$
 - 17: $\mathbf{m}.\text{extend}([\text{'user'}, \mathbf{q}])$
 - 18: **if** $r_{CoT}^{(j)} \geq \mathcal{U}(0, 1)$ **then**
 - 19: $\mathbf{a} = \text{fmt_a}(\mathbf{y}, \text{CoT}=\text{True})$
 - 20: **end if**
 - 21: $\mathbf{a} = \text{fmt_a}(\mathbf{y}, \text{CoT}=\text{False})$
 - 22: $\mathbf{m}.\text{extend}([\text{'assistant'}, \mathbf{a}])$
 - 23: **end for**
 - 24: $\mathbf{p} = \text{chat_template}(\mathbf{m})$
 - 25: **return** \mathbf{p}
-

Algorithm 6 Single chain word generation

Require: M parent_words.

- 1: Initialize temporary list of slices \mathbf{h}
 - 2: **for** $i = 1$ to M **do**
 - 3: $w \leftarrow \text{parent_words}[i]$
 - 4: $\mathbf{h}[i] \leftarrow w[\text{len}(w) // 2 :]$
 - 5: **end for**
 - 6: $o \leftarrow \text{concat}(\mathbf{h})$
 - 7: $y_c \leftarrow \text{char_offset}(o, 1)$
 - 8: **return** y_c
-

E.1 Measuring the Reliance on Intermediate ‘Thinking’ Steps

As shown in Section 7.2, there is a clear indication of the role of model size in attaining high accuracy across various α and K' . In this section, we consider the model predictions for $K' = 0$ with the ‘Force Think’ strategy and present a breakdown of the evaluation prompts based on the correctness of the intermediate ‘thinking’ steps. Since Qwen-2.5-7B-Instruct with $\alpha = 0$ was already analyzed in Section 7.2, we focus on the remaining models and α below.

Qwen-2.5-0.5B-Instruct. Based on the values plotted in Figure 5a, the accuracy of the smaller SFT’ed Qwen-2.5-0.5B-Instruct model with $\alpha = 0$ is 0.0044. Out of the 44 prompts for which the final answer was predicted correctly, Figure 16 shows the model was able to correctly predict both the intermediate step outputs for only 4.5% of those prompts. Surprisingly, in 72.7% of the 44 model predictions, both the intermediate step outputs were wrong. Similar observations can be made for $\alpha = 0.5$ in Figure 17, $\alpha = 1$ in Figure 18, $\alpha = 2$ in Figure 19, and $\alpha = \infty$ in Figure 20.




Figure 16: Trained Qwen2.5 0.5 with $\alpha = 0$ achieves only 0.44% final answer accuracy.

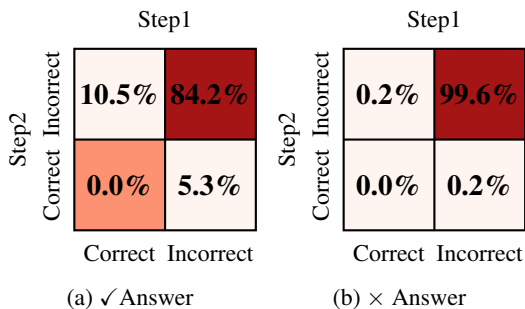


Figure 17: Trained Qwen2.5 0.5 with $\alpha = 0.5$ achieves only 0.19% final answer accuracy.




Figure 18: Trained Qwen2.5 0.5 with $\alpha = 1$ achieves only 0.17% final answer accuracy.

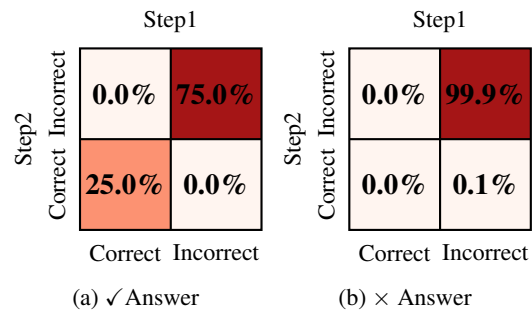


Figure 19: Trained Qwen2.5 0.5 with $\alpha = 2$ achieves only 0.04% final answer accuracy.

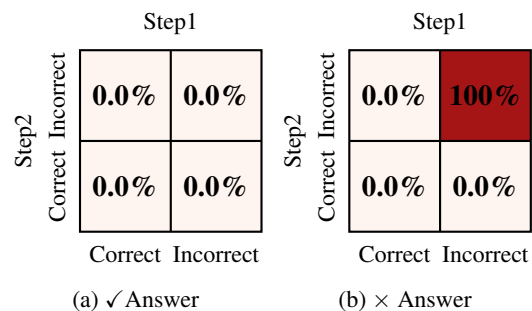


Figure 20: Trained Qwen2.5 0.5B with $\alpha = \infty$ achieves 0.0% final answer accuracy.

Qwen-2.5-1.5B-Instruct. As the model size increases, Figure 5b shows that the accuracy of the SFT’ed Qwen-2.5-1.5B-Instruct model with $\alpha = 0$ is 0.1551. Figure 21 shows that only 36.5% of the 1551 prompts that have correct final answers have both the intermediate step predictions to be correct. Although this fraction is higher than the 0.5B-Instruct case in Figure 16, it is still a considerably low fraction. On the other hand, when the model produces incorrect final answers, we can observe a relatively higher importance of the Step 2 predictions compared to Step 1. See also Figure 22 ($\alpha = 0.5$), Figure 23 ($\alpha = 1$), Figure 24 ($\alpha = 2$) and Figure 25 ($\alpha = \infty$).

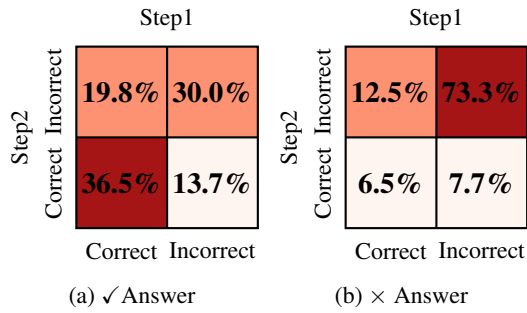


Figure 21: Trained Qwen2.5 1.5B with $\alpha = 0$ achieves 15.51% final answer accuracy.

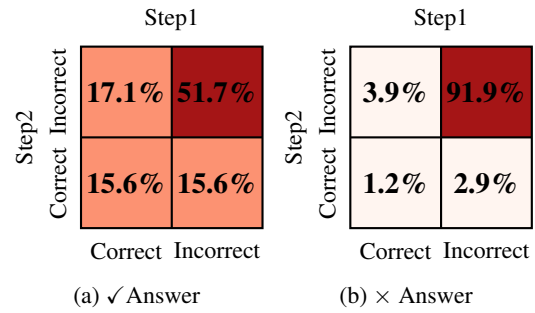


Figure 23: Trained Qwen2.5 1.5B with $\alpha = 1$ achieves 4.74% final answer accuracy.

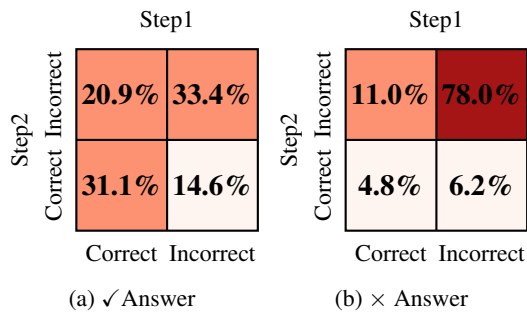


Figure 22: Trained Qwen2.5 1.5B with $\alpha = 0.5$ achieves 11.75% final answer accuracy.




Figure 24: Trained Qwen2.5 1.5B with $\alpha = 2$ achieves 2.73% final answer accuracy.

Qwen-2.5-7B-Instruct. In addition to the observations in Section 7.2 with $\alpha = 0$, an interesting trend with increasing α is that incorrect Step 2 predictions tend to have a relatively higher role in the incorrect final answer prediction. For instance when $\alpha = 0.5$ results in accuracy = 0.6296 (Figure 26), then 22.9% of the 3704 incorrect final answers can be attributed to incorrect Step 2 predictions and 11.8% to incorrect Step 1 predictions. As α increases to ∞ (Figure 29), these numbers change to 40.3% and 4.3% respectively. The gradual change can be clearly observed in Figure 27 for $\alpha = 1$ and Figure 28 for $\alpha = 2$.

Reliance on the underlying DAG. In addition to verifying the correctness of intermediate steps, we also analyze their relevance by checking if the final answer is causally dependent on them via the ground truth DAG. The results for Qwen-2.5-0.5B-Instruct are presented in Figure 30 ($\alpha = 0$), Figure 31 ($\alpha = 0.5$), Figure 32 ($\alpha = 1$), Figure 33 ($\alpha = 2$) and Figure 34 ($\alpha = \infty$). In particular, Figure 30a for $\alpha = 0$ illustrates that out of the 44 prompts with correct final answer predictions, 0 prompts had correct intermediate Step 1 predictions when the final answer had a causal dependency on this intermediate word (i.e Step 1

output). Similarly, when there was no underlying causal dependency, 6 prompts had correct whereas 34 of them had incorrect Step 1 predictions.

The results for Qwen-2.5-1.5B-Instruct are presented in Figure 35 ($\alpha = 0$), Figure 36 ($\alpha = 0.5$), Figure 37 ($\alpha = 1$), Figure 38 ($\alpha = 2$) and Figure 39 ($\alpha = \infty$) and for Qwen-2.5-7B-Instruct are presented in Figure 40 ($\alpha = 0$), Figure 41 ($\alpha = 0.5$), Figure 42 ($\alpha = 1$), Figure 43 ($\alpha = 2$) and Figure 44 ($\alpha = \infty$). A key observation is that: the final answer is causally dependent on the Step 2 intermediate word in (relatively) more number of prompts than the Step 1 intermediate word. This inherent preference for Step 2 outputs in the DAG (Figure 40) explains the results in Figure 6 (Section 7.2) and the extended results presented above where incorrect Step 2 outputs resulted in relatively more incorrect final answers when compared to Step 1. In essence, as the model continues to learn the underlying DAG to predict the final answer, its sensitivity to important intermediate steps also increases. Thus, mistakes in important intermediate steps can lead to incorrect final answers.

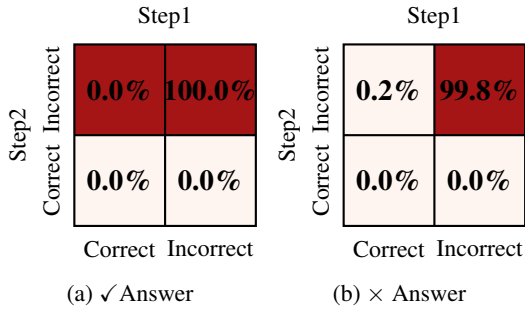


Figure 25: Trained Qwen2.5 1.5B with $\alpha = \infty$ achieves only 0.01% final answer accuracy.

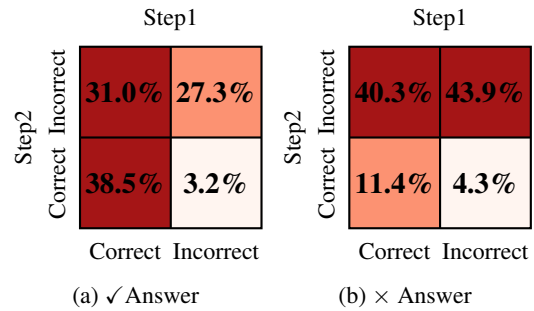


Figure 29: Trained Qwen2.5 7B with $\alpha = \infty$ achieves 5.06% final answer accuracy.

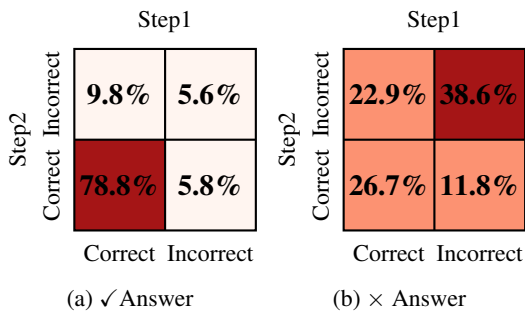


Figure 26: Trained Qwen2.5 7B with $\alpha = 0.5$ achieves 62.96% final answer accuracy.

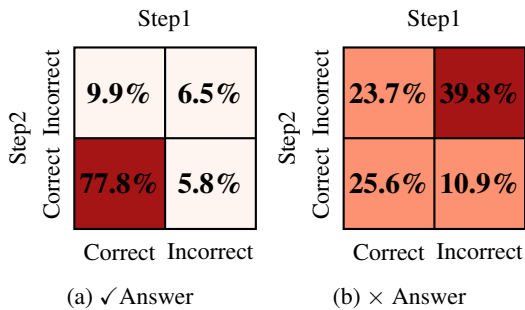


Figure 27: Trained Qwen2.5 7B with $\alpha = 1$ achieves 58.85% final answer accuracy.

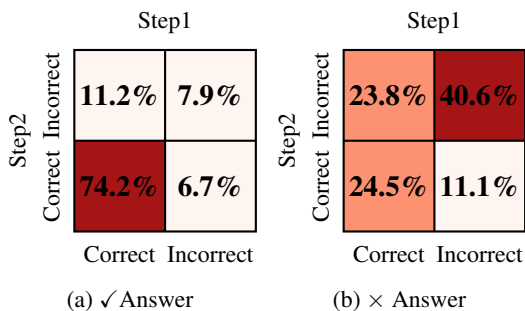


Figure 28: Trained Qwen2.5 7B with $\alpha = 2$ achieves 53.86% final answer accuracy.

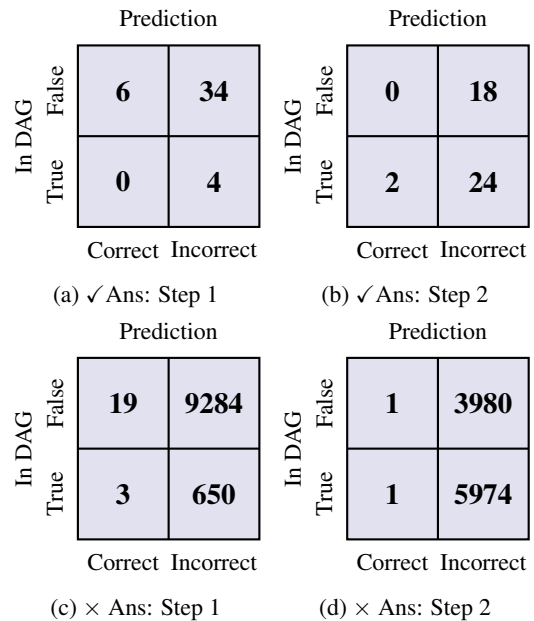


Figure 30: Breakdown of Qwen-2.5-0.5B-Instruct model predictions trained with $\alpha = 0$, based on correct (✓) and incorrect (×) final answers, and inclusion of intermediate steps in the DAG.

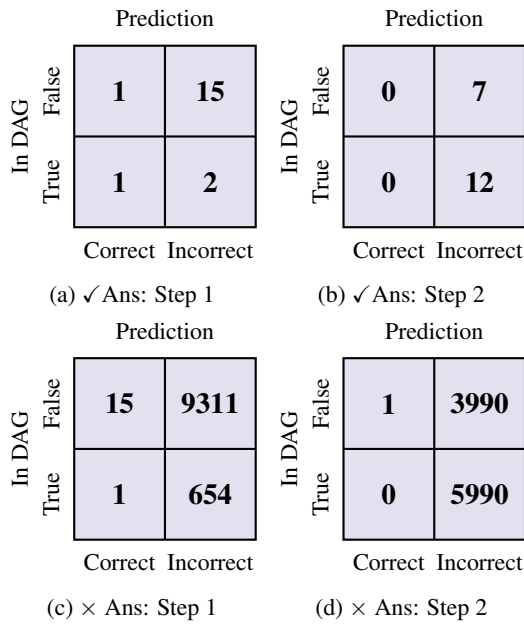


Figure 31: DAG breakdown for Trained Qwen2.5 0.5 with $\alpha = 0.5$.

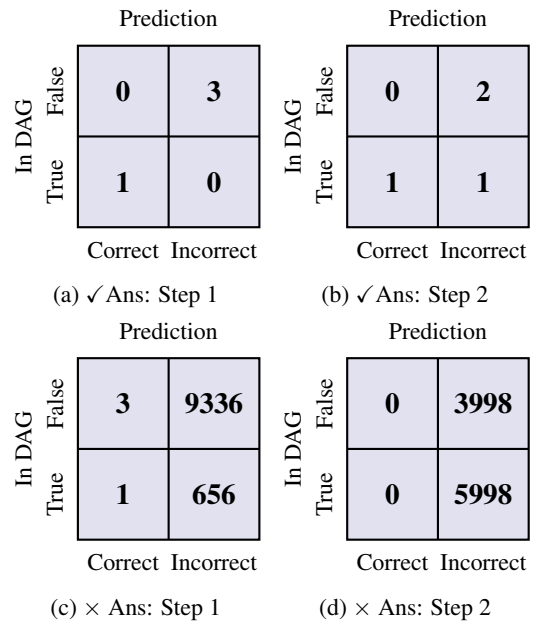


Figure 33: DAG breakdown for Trained Qwen2.5 0.5B with $\alpha = 2$.

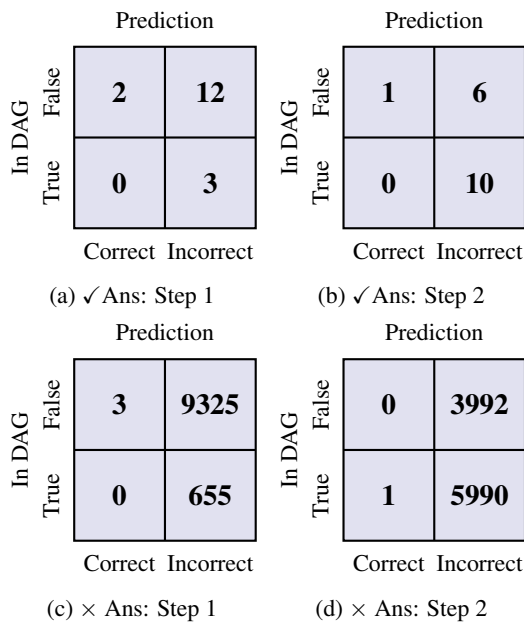


Figure 32: DAG breakdown for Trained Qwen2.5 0.5B with $\alpha = 1$.

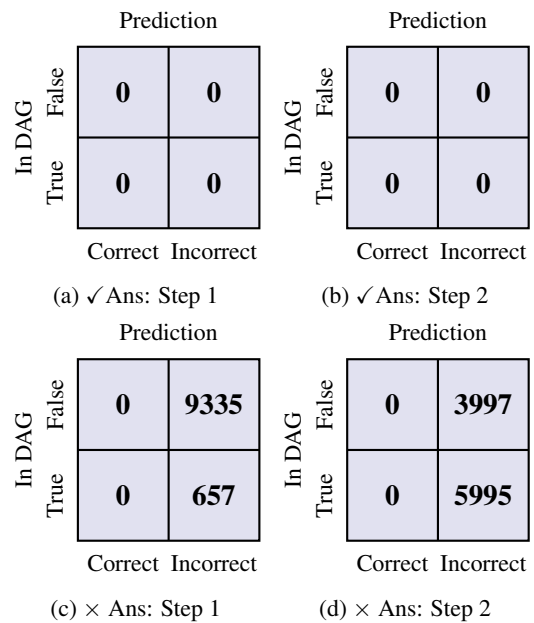


Figure 34: DAG breakdown for Trained Qwen2.5 0.5B with $\alpha = \infty$.

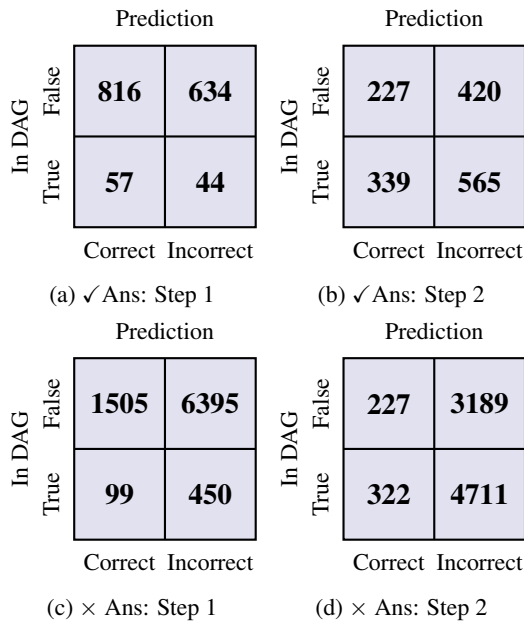


Figure 35: DAG breakdown for Trained Qwen2.5 1.5B with $\alpha = 0$.

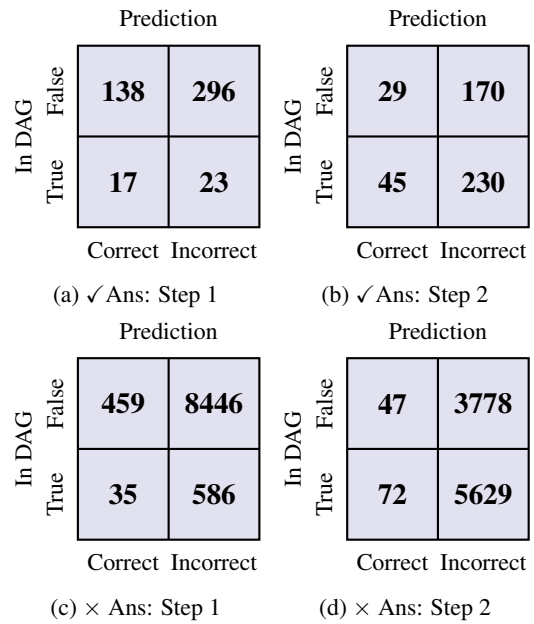


Figure 37: DAG breakdown for Trained Qwen2.5 1.5B with $\alpha = 1$.

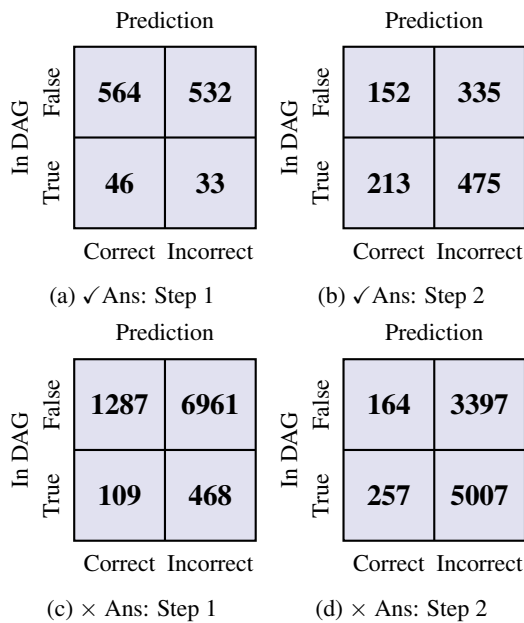


Figure 36: DAG breakdown for Trained Qwen2.5 1.5B with $\alpha = 0.5$.

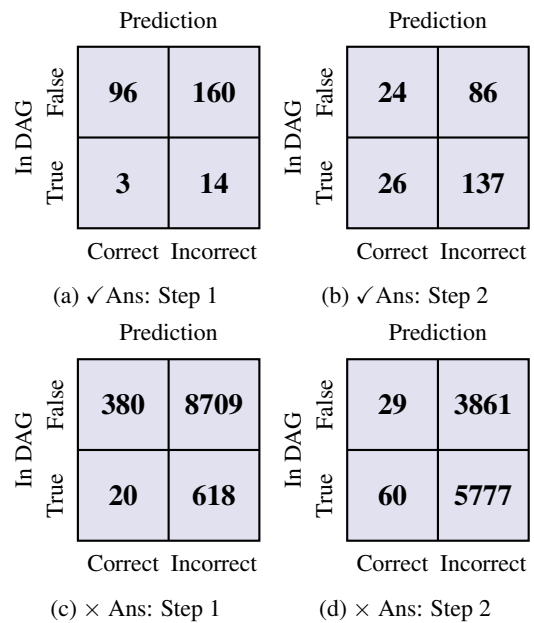


Figure 38: DAG breakdown for Trained Qwen2.5 1.5B with $\alpha = 2$.

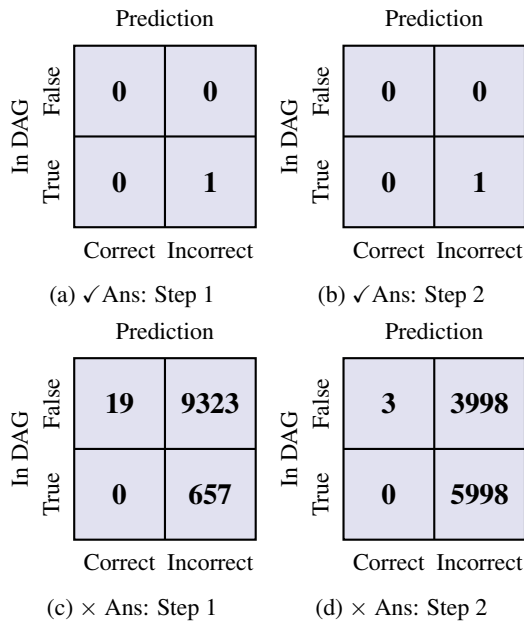


Figure 39: DAG breakdown for Trained Qwen2.5 1.5B with $\alpha = \infty$.

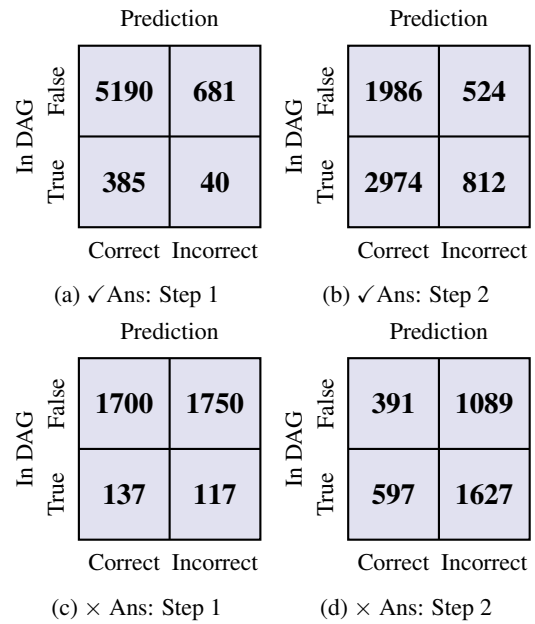


Figure 41: DAG breakdown for Trained Qwen2.5 7B with $\alpha = 0.5$.

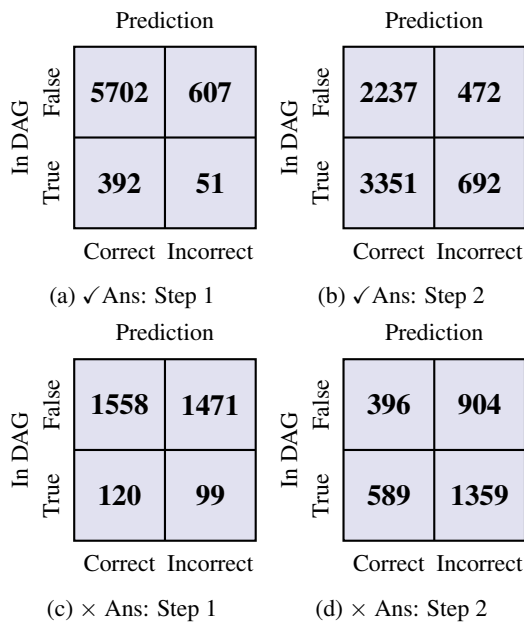


Figure 40: DAG breakdown for Trained Qwen2.5 7B with $\alpha = 0$.

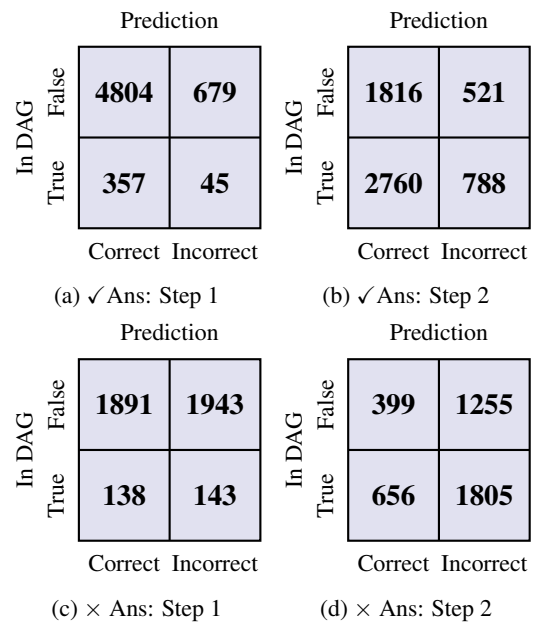


Figure 42: DAG breakdown for Trained Qwen2.5 7B with $\alpha = 1$.

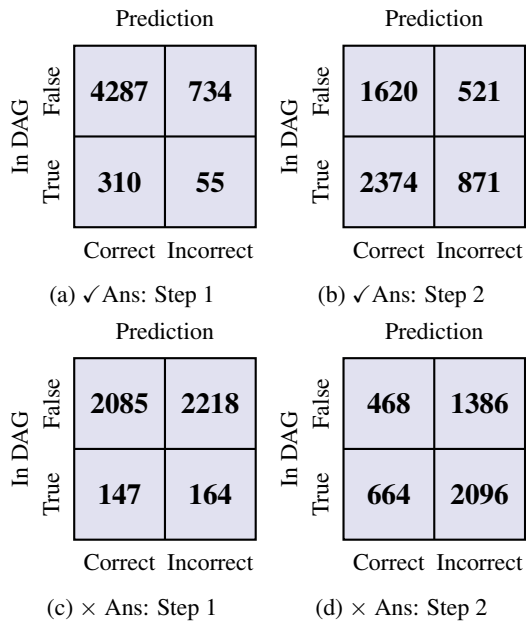


Figure 43: DAG breakdown for Trained Qwen2.5 7B with $\alpha = 2$.

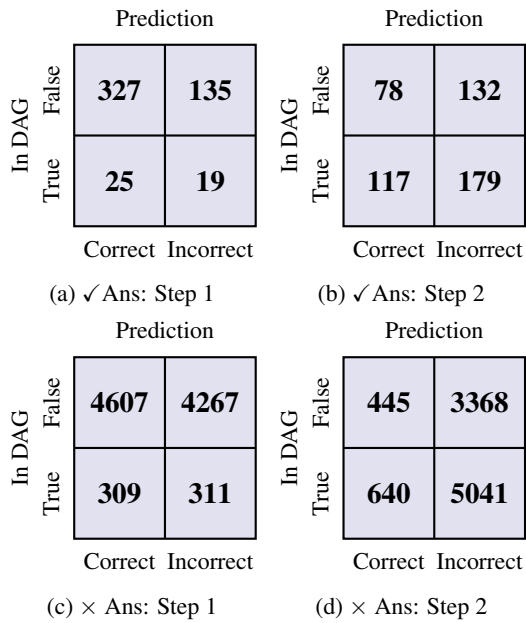


Figure 44: DAG breakdown for Trained Qwen2.5 7B with $\alpha = \infty$.

```

<|im_start|>system
You are a helpful problem solver who follows instructions.

Given some examples of string transformations, you must figure out the underlying transformation rule and apply it to the
set of words in the last example.

Instructions for the response:
a. If you choose to think out loud, you must start your response with: <|im_start|>think\n
b. The final answer must always be provided in the: <|im_start|>final answer\n\\boxed{} format.
c. You should not think or explain after providing the final answer.
<|im_end|>
<|im_start|>user
Given the input words: bkojbeez, mbjyklav, bjtorqfo, zrfvplvf, what is the final answer after the string transformations?
<|im_end|>
<|im_start|>assistant
<|im_start|>think
Step 1: cffasrgp
Step 2: srgptshq
<|im_start|>final answer
\\boxed{tshqutir}<|im_end|>
<|im_start|>user
Given the input words: jpnxmlp, lwsxxunz, kvhjsnmt, orzzxgkh, what is the final answer after the string transformations?
<|im_end|>
<|im_start|>assistant
<|im_start|>think
Step 1: ynmqtonu
Step 2: tonuupov
<|im_start|>final answer
\\boxed{upovvqpw}<|im_end|>
<|im_start|>user
Given the input words: mspqqgzv, ccjhnzlo, bxtnupxi, wocqkibx, what is the final answer after the string transformations?
<|im_end|>
<|im_start|>assistant
<|im_start|>think
Step 1: rhawvqyj
Step 2: vqyjwrzk
<|im_start|>final answer
\\boxed{wrzkxsa}<|im_end|>
...
<|im_start|>user
Given the input words: vvrhzqu, ycjowjhl, jdiaxycl, yvsszuan, what is the final answer after the string transformations?
<|im_end|>
<|im_start|>assistant
<|im_start|>think
Step 1: iarvyzdm
Step 2: yzdmzaen
<|im_start|>final answer
\\boxed{zaenabfo}<|im_end|>
<|im_start|>user
Given the input words: beuczfgl, oeonqlks, paezqha, lfzflisd, what is the final answer after the string transformations?
<|im_end|>
<|im_start|>assistant

```

Figure 45: Example evaluation prompt from the CIL-LangSym dataset with $K = 39$ CoT examples (truncated in the image for brevity) and the query formatted using the Qwen-2.5-7B-Instruct tokenizer chat template.

Model	‘Force Think’ Response	‘Force Answer’ Response
Ground Truth	Step 1: aghmarib Step 2: aribbsjc < im_start >final answer \\boxed{bsjcctkd}	\\boxed{bsjcctkd}
Baseline	Step 1: cfdwgrjm Step 2: grjmhkno Step 3: hnkoiplp ... Step 90: jjjjjjj	\\boxed{mndyehjt}
Trained ($\alpha = 0$)	Step 1: aghmarib Step 2: aribbsjc < im_start >final answer \\{bsjcctkd}	\\boxed{mhmbairo}
Trained ($\alpha = 0.5$)	Step 1: aghmarib Step 2: aribbsjc < im_start >final answer \\boxed{bsjcctkd}	\\boxed{btfcugh}
Trained ($\alpha = 1$)	Step 1: aghmairb Step 2: airbbjsc < im_start >final answer \\boxed{bjsccktn}	\\boxed{btfcugh}
Trained ($\alpha = 2$)	Step 1: aghmarih Step 2: arihsbji < im_start >final answer \\boxed{bsjictkj}	\\boxed{tmbuuncv}
Trained ($\alpha = \infty$)	Step 1: aghmarib Step 2: nbsjbstk < im_start >final answer \\boxed{cbtkcdul}	\\boxed{aghmbnhc}

Table 2: Ground truth output for the evaluation prompt in Figure 45 and responses from baseline and SFT’ed Qwen-2.5-7B-Instruct models using ‘Force Think’ strategy with prompt suffix: <|im_start|>think\n and ‘Force Answer’ strategy with prompt suffix: <|im_start|>final answer\n. Notice that only the $\alpha = 0.5$ model gives the correct answer with ‘Force Think’ strategy whereas none of the responses with ‘Force Answer’ are correct.

Notation	Description
t_{pad}	The padding token
t_{bos}	The begin-of-sequence token
t_{eos}	The end-of-sequence token
$t_{\text{inp_start}}$	The token to indicate the start of N <i>input</i> tokens.
$t_{\text{inp_end}}$	The token to indicate the end of N <i>input</i> tokens.
$t_{\text{think_start}}$	The token to indicate the start of $C - 1$ <i>intermediate</i> tokens.
$t_{\text{think_end}}$	The token to indicate the end of $C - 1$ <i>intermediate</i> tokens.
$t_{\text{ans_start}}$	The token to indicate the start of the <i>answer</i> token.
$t_{\text{ans_end}}$	The token to indicate the end of the <i>answer</i> token.
$\mathcal{V}_{\text{normal}}$	A vocabulary of abstract tokens used as inputs and chain tokens.
$\mathcal{V}_{\text{special}}$	A vocabulary of abstract tokens used as delimiters
\mathcal{V}	A unified vocabulary of abstract tokens ($\mathcal{V} = \mathcal{V}_{\text{normal}} \cup \mathcal{V}_{\text{special}}$)
\mathbf{E}_{data}	The ‘unknown’ data embedding matrix corresponding to \mathcal{V}
d	Embedding dimension for abstract tokens based on $\mathbf{E}_{\text{data}} \in \mathbb{R}^{ \mathcal{V} \times d}$
\mathcal{G}	Function class to filter tokens
\mathcal{H}	Function class to process abstract token embeddings (i.e, rows in \mathbf{E}_{data})
\mathcal{F}	Function class composed of \mathcal{G}, \mathcal{H}
l	MLP depth in \mathcal{H}
ϕ	MLP activation function in \mathcal{H}
N	Number of input tokens per example
M	Number of tokens selected by \mathcal{G}
C	Number of chain tokens (Chain length)
K	Number of examples per sequence
TF	A decoder only transformer model
$\mathbf{x} = (x_1, \dots, x_N) \in \mathcal{V}^N$	Input tokens in an example
$\mathbf{y} = (y_1, \dots, y_C) \in \mathcal{V}^C$	Chain tokens in an example
$\mathbf{p}^K(f, r_{CoT})$	Tokenized sequence generated using $f \in \mathcal{F}$ and r_{CoT} with K examples.
$\text{TF}(\cdot)$	A single auto-regressive greedy token generation by the TF model.
$\text{TF}^{oC_{eos}}(\cdot)$	The auto-regressive greedy token generation by the TF model until t_{eos} .

Table 3: A summary of notations used throughout the paper.