

000  
001  
002  
003 

# LEARNING COMPOSABLE CHAINS-OF-THOUGHT

  
004  
005  
006  
007  
008**Anonymous authors**

Paper under double-blind review

**ABSTRACT**

A common approach for teaching large language models (LLMs) to reason is to train on chains-of-thought (CoTs) of in-distribution reasoning problems, but such annotated data is costly to obtain for every problem of interest. We want reasoning models to generalize beyond their training distribution, and ideally to generalize compositionally: they should combine atomic reasoning skills to solve harder unseen tasks. In this paper, we introduce a method to enable generalization to a target compositional task that has no labeled CoT data. We find that simply training models on CoT data of atomic tasks leads to limited generalization, but minimally modifying CoT formats of constituent atomic tasks to be **composable** leads to improvement. Specifically, we augment our data by adding prefixes to CoTs, making sequences of CoTs in-distribution for the trained model. We train individual models on the atomic tasks with composable CoT data and combine them with multitask learning or model merging to address the target compositional task zero-shot. This model can be further trained on a small amount of compositional data using rejection sampling fine-tuning (RFT). Results on three domains of compositional tasks, natural language skills, string manipulation, and arithmetic, show that training LLMs on Composable CoT outperforms multitask learning and continued fine-tuning baselines within a given training data budget.

**1 INTRODUCTION**

Large language models (LLMs) are successful by virtue of the massive amounts of data they are trained on, which makes a wide range of complex problems in-distribution. However, these models still fail at challenging reasoning tasks and it is impossible to scale training data to cover all possible tasks of interest. Ideally, we want models that can *generalize* to new settings, and particularly, can apply basic “skills” learned during training in novel combinations to solve problems at inference time. How to empower LLMs with this capability, also called compositional generalization (Piantadosi & Aslin, 2016; Werchan et al., 2015; Conklin et al., 2021; Dziri et al., 2023), remains an open question. For instance, large reasoning models (QwenTeam, 2025; Guha et al., 2025), built on pre-trained LLMs, are typically trained on a large amount of data annotated with chain-of-thought (CoT) traces, but still fall short at generalizing to harder problem instances than what they were trained on (Sun et al., 2024; Hase et al., 2024; Abreu et al., 2025; Shojaee et al., 2025; Sun et al., 2025).

We explore the setting of compositional reasoning where pre-trained LLMs are fine-tuned on CoT data of simple reasoning tasks (atomic tasks) and then evaluated on the *unseen* combinations of them (compositional tasks) with no or limited compositional supervision. We find that models trained with atomic CoT data demonstrate limited generalization to compositional settings. As illustrated in Figure 1, we propose a simple modification of the CoT format of the atomic task training data, which we call **Composable CoT**: we add “proxy prefixes,” which are random filler strings, to the prompt. Then we train models to reason about atomic tasks in the context of the proxy prefixes: this makes the test-time compositional setting more in-distribution, as models need to generate a long reasoning chain by chaining multiple CoTs.

We first experiment with *zero-shot* combination of Composable CoT models. We experiment with two different approaches: first, merging models trained on individual atomic CoT tasks, and second, multitask learning across our atomic CoT datasets. Such combined models achieve zero-shot compositional generalization without seeing compositional data during training.

We then demonstrate that our zero-shot models can be improved further by rejection sampling fine-tuning on a limited amount of compositional supervision. Using *only final answer* supervision, our

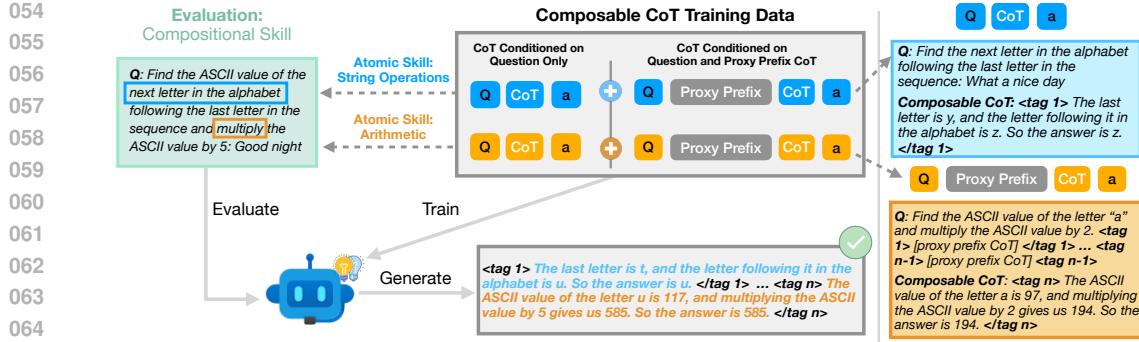


Figure 1: A compositional task involves separate atomic skills. We use a data augmentation scheme, **Composable CoT**, to create training data of atomic tasks to teach LLMs CoT formats that can be combined at inference time to address compositional tasks. We augment CoT data to be composable by adding “proxy prefix CoTs,” such as random filler strings, to the prompt to simulate compositional distributions where a CoT is conditionally generated from the question and other CoTs.

models can bootstrap better compositional CoT behavior. On tasks involving core reasoning capabilities of LLMs, including string manipulation, arithmetic, and natural language skill composition, our approach outperforms multi-task learning and continued fine-tuning baselines within a given budget of training data. Combining atomic models trained with Composable CoT is consistently better than combining Standard CoT models, with an average performance boost of 18.2% across different compositional tasks.

Moreover, Composable CoT models generalize well to complex compositions with larger skill pools: combining Composable CoT models zero-shot outperforms standard CoT models by an average performance increase of 4.8% on three-way compositions, and 8.8% on two-way compositions that require skill selection.

The main contributions of this work include: (1) A novel data augmentation scheme for training CoT models on basic reasoning skills to enable future composition of them for more complicated reasoning tasks. (2) A method for improving compositional reasoning with LLMs by first training models on atomic CoTs with such augmentation and then performing rejection sampling finetuning.

## 2 PRELIMINARIES

**LLM reasoning with chain-of-thought.** Given a prompt  $q$  that states a reasoning problem with ground truth answer  $a$ , an LLM  $M$  reasons with chain-of-thought by generating a response that includes a chain-of-thought trace  $t$  followed by a predicted answer  $\tilde{a}$ . Recent works show that supervised fine-tuning pre-trained LLMs on CoT traces leads to strong reasoning models (Muennighoff et al., 2025; Guha et al., 2025). We define a dataset for a task  $\mathcal{T}$  with CoT traces of size  $N$  as a set of (prompt, CoT, answer) triples:  $D_{\mathcal{T}}^{\text{CoT}} = \{(q, t, a)\}$ .

**Atomic and compositional tasks.** Consider a set of tasks that represent basic reasoning skills, which we call **atomic** tasks. We define **compositional** tasks  $\mathcal{T}_A$  as those tasks that can be expressed as a composition of  $n$  atomic tasks:  $\mathcal{T}_A = g(A)$  where  $A = \{\mathcal{T}_1, \dots, \mathcal{T}_n\}$ ,  $|A| = n$  and  $g$  is some function to combine the  $n$  atomic tasks. We discuss more details for  $g$  in Appendix A.

**Compositional reasoning from atomic CoT.** We assume access to atomic CoT data  $D_A^{\text{CoT}} = \{D_{\mathcal{T}_i}^{\text{CoT}} | \mathcal{T}_i \in A\}$ . Models fine-tuned on a subset of  $D_A^{\text{CoT}}$  are **atomic CoT models**.

For their composition  $\mathcal{T}_A$ , we only have access to a training dataset  $D_{\mathcal{T}_A}$  of size  $N_{\mathcal{T}_A}$ . We make two data assumptions following considerations about how compositional data would work in practice: (1) **Answer only:** The data only contains the answers as labels and *not* labeled CoT traces. This reflects that high-quality annotated CoT supervision may be harder to obtain in practice than correct answers; (2) **Limited compositional supervision:** We assume  $N_{\mathcal{T}_A}$  is small. We may be able to

108 collect a small amount of data for each new compositional task of interest, but these compositional  
 109 tasks are too numerous to undertake large-scale data collection on.  
 110

### 111 3 LEARNING COMPOSABLE CHAINS-OF-THOUGHT

114 We assume the CoT traces in each of the  $n$  atomic task datasets follow a certain distribution distinct  
 115 to that dataset. A pre-trained LLM  $M_0$  fine-tuned on the atomic CoT data can be seen as a mixture  
 116 model: it can generate CoT traces from each of those  $n$  distributions, but it is unclear whether  
 117 such models can produce compositional CoTs for compositional tasks. We observe that without  
 118 additional supervision signals, such fine-tuned models typically only replicate one of the learned  
 119 atomic reasoning patterns in the generated CoT; we show empirical evidence in Section 6.2.  
 120

121 To compose  $n$  atomic CoTs in one sequence  $\mathbf{t}_1 \dots \mathbf{t}_{n-1} \mathbf{t}_n$ , the model must allocate substantial  
 122 probability to  $p(\mathbf{t}_1 \dots \mathbf{t}_{n-1} \mathbf{t}_n \mid \mathbf{q})$ , when the model is never trained on a sequence of CoTs. Our  
 123 goal is to augment the atomic CoT training data  $(\mathbf{q}, \mathbf{t}, a)$  into  $(\mathbf{q}, \text{proxy prefixes}, \mathbf{t}, a)$ , such that **the training data looks more in-distribution to the compositional data while not explicitly training models on compositional examples.**  
 124

#### 125 3.1 CONSTRUCTING COMPOSABLE CoT TRAINING DATA

127 Consider an atomic CoT dataset  
 128  $D_{\mathcal{T}}^{\text{CoT}} = \{(\mathbf{q}, \mathbf{t}, a)\}$  for  $\mathcal{T} \in A$ ; we  
 129 call this **standard CoT** data. We aug-  
 130 ment it with a set of *chain-of-thought*  
 131 tags  $\mathcal{P} = \{p_k\}$  for  $k \in \{1, \dots, n\}$ .  
 132

133 **Proxy Prefix.** Our goal is to augment  
 134 standard CoT data such that atomic  
 135 CoT models can learn a proxy distri-  
 136 bution that simulates the distribution  
 137 of the composition of atomic skills,  
 138 despite not seeing compositional data.  
 139 Thus, we append proxy prefixes to the  
 140 prompt to simulate conditional gener-  
 141 ation of a CoT given other CoTs. Here we present a simple yet effective approach where the proxy  
 142 prefix is a sequence of *randomly sampled letters of a random length*. Such a design aims at teaching  
 143 models to generate robust continuation following an arbitrary prefix CoT. Ablations in Appendix B  
 144 show that it is more robust to distribution shift than more realistic-looking alternatives.  
 145

146 **Data Construction.** We sample a value of  $k$  for each training example  $d = (\mathbf{q}, \mathbf{t}, a)$ , and we treat  
 147  $\mathbf{t}$  as the  $k$ -th step in a notional compositional reasoning process. To achieve this, we append  $k - 1$   
 148 proxy prefixes  $(\mathbf{t}'_1 \dots \mathbf{t}'_{k-1})$  to the end of the prompt:  $\mathbf{t}'_i = \langle \text{tag } i \rangle \mathbf{t}_i \langle / \text{tag } i \rangle$  for  $1 \leq i \leq k - 1$   
 149 and  $\mathbf{t}_i$  is the  $i$ -th proxy prefix. By doing so, we obtain the augmented example  $d' = (\mathbf{q} \dots \mathbf{t}'_{k-1}, \mathbf{t}_k)$   
 150 where  $\mathbf{t}_k = \langle \text{tag } k \rangle \mathbf{t} a \langle / \text{tag } k \rangle$ .  
 151

152 Figure 2 illustrates the procedure when  $k = n$ . The standard CoT  $\mathbf{t}$  is: “*The ASCII value of the letter  
 153 a is 97, and [...]*.” We augment the example by: (1) Appending  $n - 1$  proxy prefixes to the end of the  
 154 question  $\mathbf{q}$  to obtain the augmented prompt  $\mathbf{q} \mathbf{t}'_1 \dots \mathbf{t}'_{n-1}$ , with each proxy prefix wrapped in a tag;  
 155 (2) Wrapping the CoT and the answer in a different tag  $\langle \text{tag } n \rangle$  as the augmented response  $\mathbf{t}_n$ .  
 156

157 We use the scheme above to augment each example in the standard CoT dataset and obtain the  
 158 augmented dataset  $D_{\mathcal{T}}^{\text{aug}}$ . At inference time, we do not know a given atomic CoT will be used in  
 159 which part of the compositional reasoning trace. Because CoT traces in  $D_{\mathcal{T}}^{\text{aug}}$  can simulate any of the  
 160  $k$ -th positions, models trained on  $D_{\mathcal{T}}^{\text{aug}}$  should be compatible with compositions of *arbitrary order*  
 161 instead of priming to any particular order seen during training.  
 162

163 **Learning Objective.** Then, we fine-tune  $M_0$  on  $D_{\mathcal{T}}^{\text{aug}}$  with a supervised fine-tuning objective:  
 164  $\mathcal{L}_{D_{\mathcal{T}}^{\text{aug}}}(\theta) = \frac{1}{N} \sum_{d' \in D_{\mathcal{T}}^{\text{aug}}} \mathcal{L}_{d'}(\theta)$  where  $\mathcal{L}_{d'}(\theta) = -\log p_{\theta}(\mathbf{t}_k \mid \mathbf{q} \dots \mathbf{t}'_{k-1})$ . In other words, for each  
 165 augmented example, we minimize the negative log likelihood of generating the CoT and answer,  
 166 conditioned on the question and the  $(k - 1)$  proxy prefixes.  
 167

---

162 **Algorithm 1** Bootstrapping Atomic CoT Models Trained on Composable CoT

---

163 **Input:** The combined model  $M_{\text{comb}}$ ; dataset  $D_{\mathcal{T}_A} = \{(\mathbf{q}_v, a_v)\}_{v=1}^{N_A}$ ; the number of iterations  $c$ .

164 **Output:**

165 1:  $M_0 \leftarrow M_{\text{comb}}$  ▷ Initialization

166 2: **for**  $w$  in 1... $c$  **do**

167 3:   **if** use rationalization **then**

168 4:      $(\tilde{\mathbf{t}}_v, \tilde{a}_v) \leftarrow M_{w-1}(q_v a_v) \forall v \in \{1, \dots, N_A\}$  ▷ Performance rationalization

169 5:   **else**

170 6:      $(\tilde{\mathbf{t}}_v, \tilde{a}_v) \leftarrow M_{w-1}(q_v) \forall v \in \{1, \dots, N_A\}$

171 7:   **end if**

172 8:      $D_{\text{RFT}} \leftarrow \{(\mathbf{q}_v, \tilde{\mathbf{t}}_v, a_v) \text{ s.t. } v \in \{1, \dots, N_A\} \text{ and } \tilde{a}_v = a_v\}$  ▷ CoTs with correct answers

173 9:      $M_w \leftarrow \text{SFT}(M_{\text{comb}}, D_{\text{RFT}})$  ▷ Fine-tune the combined model on the accepted CoT data

174 10: **end for**

---

175 Note that when  $k = 1$ ,  $d'$  does not have any proxy prefix in the augmented prompt, so the model  
176 learns to generate CoT traces conditioned only on the question on those examples (e.g., the top right  
177 example in Figure 1). This simulates the scenario where an atomic CoT serves as the initial step of  
178 the compositional reasoning. For  $1 < k \leq n$ , the model learns to generate CoT conditioned on both  
179 the question and proxy prefixes (e.g., the bottom right example in Figure 1).

180 **Instantiation of Tags.** In practice, models only need to learn differentiations between the  $n$ -th tag,  
181 which marks the end of the notional  $n$ -way compositional reasoning, and all the other tags, which  
182 mark intermediate steps. Thus, we set  $p_n = \langle \text{suffix} \rangle$ , and all other  $(n - 1)$  tags as  $\langle \text{prefix} \rangle$ . Despite  
183 only having two instantiations of the tag, any length of compositional CoT is supported by this  
184 scheme.

185 The scheme can also generalize to  $n$ -way composition *at inference time*. Specifically, for  $n > 2$ , we  
186 can generate a CoT, then append the  $\langle \text{suffix} \rangle$  tag, continue to generate, and repeat  $(n - 1)$  times,  
187 thereby achieving test-time generalization to  $n$ -way composition. Details can be found in Section 6.1.

### 190 3.2 COMBINING ATOMIC CoT MODELS

191 After training an atomic CoT model on a single atomic task  $\mathcal{T}$ , we need to combine multiple atomic  
192 CoT models to perform compositions. We consider two methods.

193 **ComposableCoT-MTL.** We apply multitask learning (MTL) to fine-tune  $M_0$  on the combined dataset  
194 of  $D_A^{\text{aug}} = \sum_{\mathcal{T}_i \in A} D_{\mathcal{T}_i}^{\text{aug}}$  and obtain a single MTL model  $M_{\text{comb}}$  that can generate prefix and suffix  
195 CoTs for all the  $n$  atomic tasks.

196 **ComposableCoT-Merge.** Model merging is another way to combine multiple models into a single  
197 multi-task model (Matena & Raffel, 2022; Yadav et al., 2023). For each  $\mathcal{T}_i \in A$ , we start from  $M_0$   
198 and fine-tune a model  $M_i$  (parametrized by  $\theta_i$ ) on  $D_{\mathcal{T}_i}^{\text{aug}}$ . Then we use Task Arithmetic (Ilharco  
199 et al., 2023a) to merge the  $n$  models into a single model  $M_{\text{comb}}$  parametrized by  $\theta_{\text{comb}}$  as a linear  
200 combination of the deltas between each fine-tuned model parameter and the base model parameter:  
201  $\theta_{\text{comb}} = \theta_0 + \sum_{\mathcal{T}_i \in A} \alpha_i (\theta_i - \theta_0)$  where  $\alpha$  is the scaling factor.

202 **Inference.** When running zero-shot inferences on the compositional task, we append  $\langle \text{tag 1} \rangle$  to the  
203 end of the prompt and sample a response from  $M_{\text{comb}}$ . Then, we append the next tag to the end of  
204 the generated response, continue generation, and repeat the process by appending tags up to  $\langle \text{tag } n \rangle$ .

### 208 3.3 IMPROVING COMPOSITION WITH REJECTION SAMPLING FINE-TUNING

209  $M_{\text{comb}}$  can be further improved with self-taught reasoning (Zelikman et al., 2022) by rejection  
210 sampling fine-tuning (RFT) (Dong et al., 2023; Yuan et al., 2024). Recall that for the compositional  
211 task, we only have answer labels instead of CoT traces.  $M_{\text{comb}}$  can serve as a starting point for RFT  
212 where we fine-tune  $M_{\text{comb}}$  with its own correct CoT responses using the limited compositional data.

213 Algorithm 1 shows the algorithm. Concretely, we sample responses from  $M_{\text{comb}}$  for each example  
214 in the compositional training data. Using the direct answer labels to verify the sampled responses,  
215 we can collect a supervised fine-tuning dataset  $D_{\text{RFT}}$  to continued fine-tune  $M_{\text{comb}}$ . Such a process

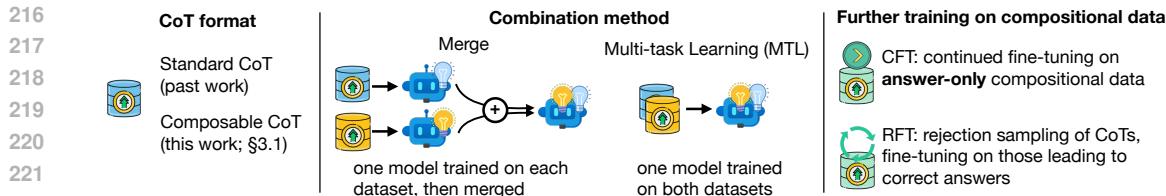


Figure 3: Summary of settings for methods evaluated. Names in the results table reference configurations described in this figure; e.g., ComposableCoT-Merge uses ComposableCoTs with model merging, and in the zero-shot setting does not use further tuning.

can be repeated for multiple iterations. For open-ended generation tasks that are hard to verify the correctness of sampled outputs only based on answer labels, we follow Zelikman et al. (2022); Ye & Durrett (2022) to perform rationalization to obtain  $D_{RFT}$ ; details can be found in Appendix D.3.

## 4 EXPERIMENTAL SETUP

We select evaluation tasks with the following criteria: (1) **Atomic tasks reflect core LLM reasoning skills**: We select atomic tasks that are representative of core skills that span logical, arithmetic, and writing. Prior work (Wei et al., 2022; Dziri et al., 2023; Yu et al., 2024) has shown that these skills can reflect more complicated capabilities such as advanced math reasoning and creative writing; (2) **Atomic skills are distinguishable**: To ensure controlled experiments of compositional generalization, atomic skills need to be distinguished from each other so that learning one skill is independent from learning another skill; (3) **Compositions are unseen during pretraining**: General reasoning tasks such as math word problems feature examples that are common in pretraining. Our tasks are less observed, thus enabling us to attribute the success of task completion to the efficacy of training approaches rather than better recall of pretraining data.

Our tasks involve string manipulation, arithmetic, and natural language skill composition. Each setting involves atomic tasks and compositional tasks. We ensure that all atomic tasks are learnable through supervised fine-tuning with a small amount of training data ( $N_T \leq 500$ ) as shown in Appendix E. We also confirm that the selected compositional tasks are less frequently seen for pre-trained LLMs: Appendix F shows the high perplexity of the task datasets, and Table 1 shows the low accuracy of few-shot prompting.

**String manipulation and arithmetic tasks.** We consider the following atomic tasks. (1) **Next letter in alphabet**: Adapted from Efrat et al. (2023); Edman et al. (2024), this task asks the LLM to find the next letter in the alphabet following the last letter in a sequence of letters. (2) **Letter concatenation**: Adapted from Wei et al. (2022); Zhou et al. (2023), this task prompts the LLM to concatenate the first, second, second-to-last, or last letter of each word in a given sequence of words. (3) **ASCII multiplication**: This task involves multi-digit multiplicative arithmetic (Dziri et al., 2023; Gambardella et al., 2024) of the ASCII value of a given letter.

We consider the following compositions of two of the atomic tasks,  $\mathcal{T}_{(i,j)} = g(\mathcal{T}_i, \mathcal{T}_j)$ . We evaluate three-way compositions and more complex compositions in Section 6.1. (1) **Next letter + multiplication**: Given a sequence of letters, find the next letter in the alphabet following the last letter, determine its ASCII value, and then perform multiplication with a given constant. (2) **Concatenation + next letter**: Given a sequence of words, concatenate the first, second, or second-to-last letter of each word and then find the next letter in the alphabet following the last letter of the concatenated sequence. (3) **Concatenation + multiplication**: Given a sequence of words, concatenate the first, second, or second-to-last letter of each word, find the ASCII value of the last letter of the concatenated sequence, and then perform multiplication.

Data and CoT traces of the above tasks are generated with templates; the data generation procedure and examples can be found in Appendix C.

**Natural language skills.** We adapt the compositional benchmark Skill-Mix (Yu et al., 2024): Given the definition and an example of a language skill (e.g. hyperbole), the model needs to write a sentence

270  
271  
272  
273  
274Table 1: Zero-shot compositional generalization of ComposableCoT with different combination approaches vs. baselines. *Without any compositional supervision*, using model merging or multitask learning to combine atomic CoT models trained on Composable CoT data outperforms baselines across settings and models, and is sometimes comparable to SFT with compositional supervision.

Methods	Next Letter + Mult EM	Concat + Next Letter EM	Concat + Mult EM	Skill-Mix Literary + Rhetorical Full Marks	Skill Fraction
Llama 2-7B					
<i>SFT on Base Model with Compositional Supervision</i>	3.1	5.0	9.0	35.5	60.1
Few-shot Answer	1.0	0.0	0.0	4.1	16.4
Few-shot CoT	2.0	3.0	1.0	7.3	23.1
StandardCoT-Merge	2.0	12.5	2.3	11.0	31.6
ComposableCoT-Merge (Ours)	16.0	<b>19.1</b>	3.0	19.6	37.1
StandardCoT-MTL	5.0	0.0	0.0	17.6	38.7
ComposableCoT-MTL (Ours)	<b>18.7</b>	6.5	<b>3.1</b>	<b>22.9</b>	<b>49.9</b>
Qwen 2.5-7B					
<i>SFT on Base Model with Compositional Supervision</i>	4.6	31.9	2.0	35.5	60.3
Few-shot Answer	2.4	0.0	2.7	34.7	56.0
Few-shot CoT	2.0	0.0	21.3	31.8	41.6
StandardCoT-Merge	70.4	54.8	<b>77.0</b>	29.8	48.0
ComposableCoT-Merge (Ours)	95.4	19.2	75.4	39.6	62.1
StandardCoT-MTL	3.6	60.9	72.1	42.0	58.2
ComposableCoT-MTL (Ours)	<b>96.3</b>	<b>63.3</b>	74.3	<b>49.0</b>	<b>66.7</b>

298  
299

300 to demonstrate the skill about a given topic. We consider an atomic task to be handling skills over  
301 a *category* of skills, and we evaluate on two categories that are mainly mutually exclusive: literary  
302 devices (*Literary*) and rhetorical devices (*Rhetorical*). Atomic CoT traces for Skill-Mix are distilled  
303 from GPT-4o (OpenAI et al., 2024), following Zhao et al. (2024). The composition tasks we consider  
304 combine **literary** and **rhetorical** skills: generate a sentence to demonstrate two provided skills, each  
305 of which is sampled from one of the categories. Examples and details can be found in Appendix D.

306  
307  
308  
309

**Evaluation Metrics.** For Skill-Mix tasks, we use quality measure metrics for the generated sentence  
from Yu et al. (2024) (namely, *Full Marks* and *Skill Fraction*) based on a rubric, and use GPT-4o-mini  
as a judge. Details can be found in Appendix D.2. All other tasks are evaluated using *exact match*  
accuracy; a regex-based answer extractor is used to extract the answer from the generated response.

310  
311  
312  
313  
314  
315  
316  
317

**Zero-shot/Few-shot Baselines.** Figure 3 summarizes the high-order variables of the configurations  
we evaluate. For zero-shot compositional generalization, we include the following baselines: (1)  
Few-shot direct answer prompting: we prompt  $M_0$  with 5-shot demonstrations drawn from the  
compositional data; (2) Few-shot CoT prompting: we prompt  $M_0$  with 5-shot CoT demonstrations  
drawn from the *atomic* data; (3) Model merging of atomic CoT models (*StandardCoT-Merge*): we  
fine-tune two models  $M_i$  and  $M_j$  based on  $M_0$  with  $D_{T_i}^{\text{CoT}}$  and  $D_{T_j}^{\text{CoT}}$  respectively and merge them  
into  $M_{\text{comb}}$  with Task Arithmetic; (4) Multitask learning of atomic CoTs (*StandardCoT-MTL*): we  
fine-tune  $M_0$  to be a single multitask learning model  $M_{\text{SCoT-MTL}}$  on  $D_{T_i}^{\text{CoT}} + D_{T_j}^{\text{CoT}}$ .

318  
319  
320  
321  
322  
323

**Baselines with Compositional Supervision.** With the *same* compositional training dataset with only  
the answer label  $D_{T_{(i,j)}}$ , we compare bootstrapping Composable CoT with the following baselines.  
(1) Continued fine-tuning (CFT) the multitask model of atomic CoTs (*CFT on StandardCoT-MTL*):  
we continue fine-tune the multitask model  $M_{\text{SCoT-MTL}}$  on  $D_{T_{(i,j)}}$ ; (2) Continued fine-tuning the  
merged model of atomic CoTs (*CFT on StandardCoT-Merge*): we continue fine-tune the merged  
model of the two atomic CoT models  $M_{\text{comb}}$  on  $D_{T_{(i,j)}}$ ; (3) Multitask learning of atomic CoTs and  
compositional answers (*StandardCoT + Comp Answer*): we fine-tune a single multitask learning

324  
 325 Table 2: Compositional task performance of rejection sampling fine-tuning (RFT) upon merged  
 326 Composable atomic CoT models and other baselines. *Mult* stands for ASCII multiplication and *concat*  
 327 stands for letter concatenation. *SFT* stands for supervised fine-tuning with the compositional answer  
 328 data; *CFT* stands for continued fine-tuning; *MTL* stands for multitask learning method. Results on  
 329 next letter + mult are omitted because the zero-shot performance saturates. RFT on ComposableCoT  
 330 variants achieves the best compositional performance using the same compositional answer data.

331 Category	332 Method	333 Next Letter + Mult EM	334 Concat + Next Letter EM	335 Concat + Mult EM	336 Skill-Mix Literary + Rhetorical Full Marks	337 Skill Fraction
338 Llama 2-7B						
339 SFT	SFT on Base Model	3.1	5.0	9.0	35.5	60.1
340 SFT	CFT on StandardCoT-Merge	2.0	16.0	14.0	44.1	65.1
341 SFT	CFT on StandardCoT-MTL	3.0	26.0	11.0	38.0	62.1
342 MTL	StandardCoT + Comp Answer	5.0	<b>46.0</b>	13.3	22.9	45.5
343 RFT	StandardCoT-Merge	0.0	23.0	29.7	26.1	52.0
344 RFT	ComposableCoT-Merge (Ours)	<b>72.0</b>	<b>46.0</b>	<b>40.0</b>	<b>45.3</b>	<b>66.6</b>
345 Qwen 2.5-7B						
346 SFT	SFT on Base Model	-	31.9	2.0	35.5	60.3
347 SFT	CFT on StandardCoT-Merge	-	41.1	9.3	51.0	71.4
348 SFT	CFT on StandardCoT-MTL	-	60.3	12.7	34.7	56.3
349 MTL	StandardCoT + Comp Answer	-	65.1	7.1	41.2	55.3
350 RFT	StandardCoT-MTL	-	82.1	<b>89.0</b>	44.9	63.4
351 RFT	ComposableCoT-MTL (Ours)	-	<b>86.9</b>	88.4	<b>57.6</b>	<b>71.5</b>

351 model based on  $M_0$  on the combined dataset of  $D_{\mathcal{T}_i}^{\text{CoT}} + D_{\mathcal{T}_j}^{\text{CoT}} + D_{\mathcal{T}_{(i,j)}}$ . We also include supervised  
 352 learning baselines (SFT) where  $M_0$  is fine-tuned on the same compositional answer data  $D_{\mathcal{T}_{(i,j)}}$ .

353 The differences of methods we evaluate for each setting are summarized in Table 13.

354 **Data Construction.** Because of two-way compositions, we sample uniformly from 2 chain-of-  
 355 thought tags, *<prefix>* and *<suffix>*, for data construction. At inference time, we first append  
 356 *<prefix>* to the prompt and sample from the combined model. Then, we append *<suffix>* to the end  
 357 of the generated response, and continue generation.

358 **Models and Training.** We use Llama 2 7B-base (Touvron et al., 2023) and Qwen2.5 7B-base (Yang  
 359 et al., 2025) as models, and use LoRA (Hu et al., 2022) for supervised fine-tuning. For rejection  
 360 sampling, we sample 10 responses for each prompt and use temperature  $\tau = 0.9$  for inference;  
 361 otherwise, we use greedy decoding. For Skill-Mix tasks, we perform rationalization for RFT because  
 362 it is open-ended generation (see Section 3.3). Configuration and hyperparameters are in Appendix G.

## 363 5 RESULTS

### 364 5.1 ZERO-SHOT GENERALIZATION

365 We evaluate the compositional generalization of the proposed method *without compositional supervision*,  
 366 including ComposableCoT-Merge and ComposableCoT-MTL. For all methods that we compare  
 367 with, we control the amount of training data to be the same as  $N_i$  and  $N_j$ . For reference, we also  
 368 include the supervised fine-tuning baseline by fine-tuning  $M_0$  with  $N_{(i,j)}$  compositional answer data.  
 369 Details of the training data for each task can be found in Appendix H.

370 **Learning ComposableCoT achieves better zero-shot generalization.** Table 1 shows that ComposableCoT  
 371 variants outperform all baselines on a range of tasks for both models. Combining atomic  
 372 CoT models trained on ComposableCoT is better than combining models trained on standard CoT  
 373 across settings. Moreover, while having seen no compositional training data, our method achieves  
 374 comparable or even better performance than supervised fine-tuning baselines *with* compositional

378 supervision (e.g., next letter + multiplication). These indicate that the Composable CoT format leads  
 379 to better "composability" at inference time.  
 380

## 381 5.2 COMPOSITIONAL PERFORMANCE WITH LIMITED SUPERVISION 382

383 We evaluate the performance of Compos-  
 384 able CoT models after being further im-  
 385 proved with one iteration of RFT using  
 386 the limited compositional supervision. We  
 387 compare it with multitask learning and con-  
 388 tinued fine-tuning baselines given the same  
 389 compositional answer dataset  $D_{\mathcal{T}_{(i,j)}}$  of  
 390 size  $N_{(i,j)} \leq 500$ . For reference, we in-  
 391 clude the baseline of fine-tuning  $M_0$  on the  
 392 same compositional answer data. Details  
 393 of the data condition are in Appendix H.  
 394

395 Table 2 shows that with the same compo-  
 396 sitional training data, **using RFT on top**  
 397 of **ComposableCoT-MTL** and **ComposableCoT-Merge** achieves the best compositional task  
 398 performance, outperforming multitask learning and continued fine-tuning baselines across settings.

399 We further investigate if the per-  
 400 formance is mainly driven by RFT  
 401 or by learning Composable CoT  
 402 format. We compare RFT upon  
 403 StandardCoT-Merge with RFT upon  
 404 ComposableCoT-Merge for LLama  
 405 2-7B, and StandardCoT-MTL with  
 406 ComposableCoT-MTL for Qwen 2.5-  
 407 7B. Table 2 shows that RFT is a better  
 408 way to improve the compositional task  
 409 performance of StandardCoT models  
 410 with compositional data than MTL  
 411 and SFT. Moreover, **RFT upon Com-**  
 412 **posableCoT models is generally better than RFT upon StandardCoT models.** Using the same  
 413 combination method (MTL or Model Merging), RFT upon ComposableCoT models outperforms the  
 414 StandardCoT counterpart by an average performance increase of 18.2% across models and tasks.

## 415 6 ANALYSIS

### 416 6.1 GENERALIZATION TO COMPLEX COMPOSITIONS

417 **Three-way Composition.** We evaluate Composable CoT on zero-shot compositions of *three* atomic  
 418 tasks on Qwen2.5-7B using the following compositional tasks: (1) Letter Concat + Next Letter +  
 419 Mult (*String Tasks*): Given a sequence of words, concatenate the first, second, second-to-last, or last  
 420 letter of each word, find the next letter in the alphabet following the last letter of the concatenated  
 421 sequence, find the ASCII value of this letter, and then perform multiplication. (2) Skill-Mix Literary +  
 422 Rhetorical + Logical (*Skill-Mix*): Generate a sentence on a given topic to demonstrate three provided  
 423 skills, each of which is sampled from one of the Skill-Mix categories, including an additional category  
 424 *Skill-Mix-Logical*. We compare ComposableCoT models with StandardCoT models constructed by  
 425 model merging or multi-task learning. Implementation details can be found in Appendix I.

426 Table 3 shows that given the same combination method, combining ComposableCoT models is  
 427 better on three-way composition: for example, using MTL, ComposableCoT models outperform  
 428 StandardCoT models by an average performance increase of 4.8%.

429 **Two-way Composition with Larger Skill Pools.** In practice, models may need to have many  
 430 capabilities to address problems of interest. Compared to our existing settings, such models need to  
 431 select the skills to engage with for a particular task out of a larger pool of learned skills.

381 Table 3: Zero-shot generalization on three-way composi-  
 382 tions. Combining ComposableCoT models outperforms  
 383 combining StandardCoT models on the composition of  
 384 three tasks.

	String Tasks		Skill-Mix
	EM	Full Mark	Skill Fraction
Standard-Merge	61.3	13.1	42.7
Composable-Merge	63.1	19.2	54.1
Standard-MTL	82.3	28.2	55.9
Composable-MTL	<b>86.7</b>	<b>33.1</b>	<b>61.0</b>

381 Table 4: Zero-shot generalization on two-way compositions  
 382 when merging **three** atomic models (i.e., there is a distractor  
 383 skill). Merging ComposableCoT models is better than merging  
 384 StandardCoT models in this setting.

	Standard Composable		
	EM	Full Mark	Skill Fraction
Next Letter + Mult	56.1	<b>75.9</b>	
Concat + Next Letter	39.1	<b>46.2</b>	
Concat + Mult	44.3	<b>48.9</b>	
Skill-Mix Literary + Rhetorical	37.1	<b>42.0</b>	
	55.1	<b>62.7</b>	

415 **Using the same combination method, RFT upon ComposableCoT models outperforms the**  
 416 **StandardCoT counterpart by an average performance increase of 18.2% across models and tasks.**

To evaluate this scenario, we train atomic models on *three* atomic skills on Qwen 2.5-7B and evaluate the combined model (with model merging) on the zero-shot composition of *two* of the atomic skills. This setup provides the model with additional potential combinations of skills to reason about. For the Skill-Mix tasks, we train on an additional atomic task, logical reasoning. Table 4 shows that learning ComposableCoT outperforms using StandardCoT by 8.8% on average, indicating that ComposableCoT models can select the appropriate skills to compose out of many learned skills.

## 6.2 QUALITY OF GENERATED COTs

We conduct intrinsic quality evaluations on CoTs generated by ComposableCoT models for zero-shot composition. For the string manipulation and arithmetic tasks, we extract template-based patterns of each atomic CoT from the generated outputs of models evaluated on the compositional task. For Skill-Mix, we consider the CoT pattern of an atomic task to be used if the generated response explicitly mentions the skill corresponding to that atomic skill category.

Table 5 shows results with models trained from Qwen 2.5-7B and combined with MTL; results using model merging can be found in Appendix J. **Combining ComposableCoT leads to consistently higher presence of both atomic CoT patterns in the generated responses compared to StandardCoT.** Models trained with the Composable CoT format therefore leverage the combination of learned skills in some form more frequently than StandardCoT. Example of generated CoTs can be found in Appendix K.

## 7 RELATED WORK

As an important cognitive capability of humans (Piantadosi & Aslin, 2016; Werchan et al., 2015), compositional generalization has been considered a core capability for human-level reasoning models (Fodor & Pylyshyn, 1988; Lake & Baroni, 2023). Recent theoretical analyses show that LLMs can improve their compositional reasoning by generating CoT (Li et al., 2024; 2023), but empirical improvements have only been observed (Sprague et al., 2025) with non-trivial engineering effort such as prompt engineering (Chen et al., 2024; Gao et al., 2024) and data selection (Khot et al., 2023; Zhou et al., 2023; Levy et al., 2023; Ye et al., 2023). Aiming at more principled ways to improve composition, we are inspired by a line of work on efficient methods for combining models of different capabilities, particularly model merging (Tam et al., 2024a; Ilharco et al., 2023b; Wu et al., 2025; KimiTeam et al., 2025; Ma et al., 2025; Tam et al., 2024b). Our work is the first to use model merging for compositional generalization with CoT.

## 8 CONCLUSION

We propose Composable Chain-of-Thought, a data augmentation scheme to convert CoT data of atomic reasoning skills into a format that facilitates inference-time compositional generalization. Training atomic CoT models with Composable CoT and combining them with model merging or multitask learning leads to better zero-shot compositional reasoning performance than building models with the standard CoT format. Such a combined model can be further improved by a limited amount of compositional data with rejection sampling fine-tuning. Learning to reason with composable CoT shows a promising approach to improve compositional reasoning in LLMs, and could be extended to build more efficient and robust large reasoning models.

Table 5: Quality of the generated CoTs by ComposableCoT models on zero-shot compositions. “%  $\mathcal{T}_1$ ” denotes the percentage of generated responses that use the CoT format of the first atomic task of the composition, and likewise for the second.  $^{\dagger}$  denotes that the ComposableCoT method has a significantly higher “% Both” than the StandardCoT counterpart at the 0.01 level using a paired bootstrap test. “Perf.” denotes the task performance.

	CoT	Perf.	% $\mathcal{T}_1$	% $\mathcal{T}_2$	% Both
Next Letter	Standard	3.6	0.0	100.0	0.0
	Composable	96.3	98.9	100.0	$^{\dagger}98.9$
Concat	Standard	72.1	99.7	32.1	32.1
	Composable	74.3	100.0	83.1	$^{\dagger}81.3$
Concat	Standard	60.9	100.0	66.7	66.7
	Composable	63.3	100.0	85.9	$^{\dagger}85.0$
Literary	Standard	42.0	65.3	58.0	37.6
	Composable	49.0	64.5	65.7	$^{\dagger}42.0$

486 

## 9 ETHICS STATEMENT

487  
488 This work does not involve human subjects or the release of sensitive data. We do not clearly see the  
489 harms of the applications of the proposed method either, so we are not aware of any obvious ethical  
490 concern related to this work.  
491492 

## 10 REPRODUCIBILITY STATEMENT

493  
494 We report all technical details for our proposed method, including the data augmentation schema and  
495 the training methods in Section 3. To reproduce our experimental results, we report all details of the  
496 evaluation setup (Section 4) and training configurations (Section G).  
497498 

## REFERENCES

500  
501 Natalie Abreu, Edwin Zhang, Eran Malach, and Naomi Saphra. A taxonomy of transcendence.  
502 In *Second Conference on Language Modeling*, 2025. URL <https://openreview.net/forum?id=tfTn8616Gf>.  
503  
504 Jiaao Chen, Xiaoman Pan, Dian Yu, Kaiqiang Song, Xiaoyang Wang, Dong Yu, and Jianshu  
505 Chen. Skills-in-context: Unlocking compositionality in large language models. In Yaser Al-  
506 Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Findings of the Association for Compu-  
507 tational Linguistics: EMNLP 2024*, pp. 13838–13890, Miami, Florida, USA, November 2024.  
508 Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-emnlp.812. URL  
509 <https://aclanthology.org/2024.findings-emnlp.812/>.  
510  
511 Karl Cobbe, Vineet Kosaraju, Mo Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias  
512 Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman.  
513 Training verifiers to solve math word problems. *ArXiv*, abs/2110.14168, 2021. URL <https://api.semanticscholar.org/CorpusID:239998651>.  
514  
515 Henry Conklin, Bailin Wang, Kenny Smith, and Ivan Titov. Meta-learning to compositionally  
516 generalize. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Proceedings of the*  
517 *59th Annual Meeting of the Association for Computational Linguistics and the 11th International*  
518 *Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 3322–3335,  
519 Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.  
520 258. URL <https://aclanthology.org/2021.acl-long.258/>.  
521  
522 Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao,  
523 Jipeng Zhang, KaShun SHUM, and Tong Zhang. RAFT: Reward ranked finetuning for generative  
524 foundation model alignment. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856.  
525 URL <https://openreview.net/forum?id=m7p507zb1Y>.  
526  
527 Nouha Dziri, Ximing Lu, Melanie Sclar, Xiang Lorraine Li, Liwei Jiang, Bill Yuchen Lin, Sean  
528 Welleck, Peter West, Chandra Bhagavatula, Ronan Le Bras, Jena D. Hwang, Soumya Sanyal, Xiang  
529 Ren, Allyson Ettinger, Zaid Harchaoui, and Yejin Choi. Faith and fate: Limits of transformers on  
530 compositionality. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.  
531 URL <https://openreview.net/forum?id=Fkckkr3ya8>.  
532  
533 Lukas Edman, Helmut Schmid, and Alexander Fraser. CUTE: Measuring LLMs’ understanding of  
534 their tokens. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the*  
535 *2024 Conference on Empirical Methods in Natural Language Processing*, pp. 3017–3026, Miami,  
536 Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/  
537 2024.emnlp-main.177. URL <https://aclanthology.org/2024.emnlp-main.177/>.  
538  
539 Avia Efrat, Or Honovich, and Omer Levy. LMentry: A language model benchmark of elementary  
540 language tasks. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Findings of the*  
541 *Association for Computational Linguistics: ACL 2023*, pp. 10476–10501, Toronto, Canada, July  
542 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-acl.666. URL  
543 <https://aclanthology.org/2023.findings-acl.666/>.

540 Jerry A. Fodor and Zenon W. Pylyshyn. Connectionism and cognitive architecture: A critical analy-  
 541 sis. *Cognition*, 28:3–71, 1988. URL <https://api.semanticscholar.org/CorpusID:29043627>.

542

543 Andrew Gambardella, Yusuke Iwasawa, and Yutaka Matsuo. Language models do hard arith-  
 544 metic tasks easily and hardly do easy arithmetic tasks. In Lun-Wei Ku, Andre Martins,  
 545 and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for  
 546 Computational Linguistics (Volume 2: Short Papers)*, pp. 85–91, Bangkok, Thailand, August  
 547 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-short.8. URL  
 548 <https://aclanthology.org/2024.acl-short.8/>.

549

550 Peizhong Gao, Ao Xie, Shaoguang Mao, Wenshan Wu, Yan Xia, Haipeng Mi, and Furu Wei. Meta  
 551 reasoning for large language models. *arXiv preprint arXiv:2406.11698*, 2024.

552

553 Aaron Gokaslan and Vanya Cohen. OpenWebText Corpus. 2019.

554

555 Etash Guha, Ryan Marten, Sedrick Keh, Negin Raoof, Georgios Smyrnis, Hritik Bansal, Marianna  
 556 Nezhurina, Jean Mercat, Trung Vu, Zayne Sprague, Ashima Suvarna, Benjamin Feuer, Liangyu  
 557 Chen, Zaid Khan, Eric Frankel, Sachin Grover, Caroline Choi, Niklas Muennighoff, Shiye Su,  
 558 Wanjia Zhao, John Yang, Shreyas Pimpalgaonkar, Kartik Sharma, Charlie Cheng-Jie Ji, Yichuan  
 559 Deng, Sarah Pratt, Vivek Ramanujan, Jon Saad-Falcon, Jeffrey Li, Achal Dave, Alon Albalak,  
 560 Kushal Arora, Blake Wulfe, Chinmay Hegde, Greg Durrett, Sewoong Oh, Mohit Bansal, Saadia  
 561 Gabriel, Aditya Grover, Kai-Wei Chang, Vaishaal Shankar, Aaron Gokaslan, Mike A. Merrill,  
 562 Tatsunori Hashimoto, Yejin Choi, Jenia Jitsev, Reinhard Heckel, Maheswaran Sathiamoorthy,  
 563 Alexandros G. Dimakis, and Ludwig Schmidt. OpenThoughts: Data recipes for reasoning models,  
 564 2025. URL <https://arxiv.org/abs/2506.04178>.

565

566 Peter Hase, Mohit Bansal, Peter Clark, and Sarah Wiegreffe. The unreasonable effectiveness of  
 567 easy training data for hard tasks. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.),  
 568 *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume  
 569 1: Long Papers)*, pp. 7002–7024, Bangkok, Thailand, August 2024. Association for Computational  
 570 Linguistics. doi: 10.18653/v1/2024.acl-long.378. URL <https://aclanthology.org/2024.acl-long.378/>.

571

572 Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,  
 573 and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International  
 574 Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=nZeVKeFYf9>.

575

576 Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi,  
 577 and Ali Farhadi. Editing models with task arithmetic. In *The Eleventh International Conference  
 578 on Learning Representations*, 2023a. URL <https://openreview.net/forum?id=6t0Kwf8-jrj>.

579

580 Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi,  
 581 and Ali Farhadi. Editing models with task arithmetic. In *The Eleventh International Conference  
 582 on Learning Representations*, 2023b. URL <https://openreview.net/forum?id=6t0Kwf8-jrj>.

583

584 Tushar Khot, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and  
 585 Ashish Sabharwal. Decomposed prompting: A modular approach for solving complex tasks.  
 586 In *The Eleventh International Conference on Learning Representations*, 2023. URL [https://openreview.net/forum?id=\\_nGgzQjzaRy](https://openreview.net/forum?id=_nGgzQjzaRy).

587

588 KimiTeam, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun  
 589 Xiao, Chenzhuang Du, Chonghua Liao, Chunling Tang, Congcong Wang, Dehao Zhang, Enming  
 590 Yuan, Enzhe Lu, Fengxiang Tang, Flood Sung, Guangda Wei, Guokun Lai, Haiqing Guo, Han  
 591 Zhu, Hao Ding, Hao Hu, Hao Yang, Hao Zhang, Haotian Yao, Haotian Zhao, Haoyu Lu, Haoze Li,  
 592 Haozhen Yu, Hongcheng Gao, Huabin Zheng, Huan Yuan, Jia Chen, Jianhang Guo, Jianlin Su,  
 593 Jianzhou Wang, Jie Zhao, Jin Zhang, Jingyuan Liu, Junjie Yan, Junyan Wu, Lidong Shi, Ling Ye,  
 Longhui Yu, Mengnan Dong, Neo Zhang, Ningchen Ma, Qiwei Pan, Qucheng Gong, Shaowei Liu,  
 Shengling Ma, Shupeng Wei, Sihan Cao, Siying Huang, Tao Jiang, Weihao Gao, Weimin Xiong,

594 Weiran He, Weixiao Huang, Wenhao Wu, Wenyang He, Xianghui Wei, Xianqing Jia, Xingzhe Wu,  
 595 Xinran Xu, Xinxing Zu, Xinyu Zhou, Xuehai Pan, Y. Charles, Yang Li, Yangyang Hu, Yangyang  
 596 Liu, Yanru Chen, Yejie Wang, Yibo Liu, Yidao Qin, Yifeng Liu, Ying Yang, Yiping Bao, Yulun Du,  
 597 Yuxin Wu, Yuzhi Wang, Zaida Zhou, Zhaoji Wang, Zhaowei Li, Zhen Zhu, Zheng Zhang, Zhexu  
 598 Wang, Zhilin Yang, Zhiqi Huang, Zihao Huang, Ziyao Xu, and Zonghan Yang. Kimi k1.5: Scaling  
 599 reinforcement learning with llms. *arXiv preprint arXiv:2501.12599*, 2025.

600 Brenden M. Lake and Marco Baroni. Human-like systematic generalization through a meta-learning  
 601 neural network. *Nature*, 623:115 – 121, 2023. URL <https://api.semanticscholar.org/CorpusID:264489248>.

602

603 Itay Levy, Ben Beglin, and Jonathan Berant. Diverse demonstrations improve in-context compositional  
 604 generalization. In *Proceedings of the 61st Annual Meeting of the Association for Computational  
 605 Linguistics (Volume 1: Long Papers)*, Toronto, Canada, July 2023.

606

607 Yingcong Li, Kartik Sreenivasan, Angeliki Giannou, Dimitris Papailiopoulos, and Samet Oymak.  
 608 Dissecting chain-of-thought: Compositional through in-context filtering and learning. In  
 609 *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=xEhKwsqxMa>.

610

611 Zhiyuan Li, Hong Liu, Denny Zhou, and Tengyu Ma. Chain of thought empowers transform-  
 612 ers to solve inherently serial problems. In *The Twelfth International Conference on Learning  
 613 Representations*, 2024. URL <https://openreview.net/forum?id=3EWTEy9MTM>.

614

615 Xinyin Ma, Guangnian Wan, Runpeng Yu, Gongfan Fang, and Xinchao Wang. CoT-Valve: Length-  
 616 Compressible Chain-of-Thought Tuning. *arXiv preprint arXiv:2502.09601*, 2025.

617

618 Ryan Marten, Trung Vu, Charlie Cheng-Jie Ji, Kartik Sharma, Shreyas Pimpalgaonkar, Alex Dimakis,  
 619 and Maheswaran Sathiamoorthy. Curator: A Tool for Synthetic Data Creation. January 2025.

620

621 Michael S Matena and Colin Raffel. Merging Models with Fisher-Weighted Averaging. In Alice H. Oh,  
 622 Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information  
 623 Processing Systems*, 2022. URL [https://openreview.net/forum?id=LSKlp\\_aceOC](https://openreview.net/forum?id=LSKlp_aceOC).

624

625 Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke  
 626 Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time  
 627 scaling. *arXiv preprint arXiv:2501.19393*, 2025.

628

629 OpenAI, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark,  
 630 AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-Whitcomb,  
 631 Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex Paino,  
 632 Alex Renzin, Alex Tachard Passos, Alexander Kirillov, Alexi Christakis, Alexis Conneau, Ali  
 633 Kamali, Allan Jabri, Allison Moyer, Allison Tam, Amadou Crookes, Amin Tootoochian, Amin  
 634 Tootoonchian, Ananya Kumar, Andrea Vallone, Andrej Karpathy, Andrew Braunstein, Andrew  
 635 Cann, Andrew Codispoti, Andrew Galu, Andrew Kondrich, Andrew Tulloch, Andrey Mishchenko,  
 636 Angela Baek, Angela Jiang, Antoine Pelisse, Antonia Woodford, Anuj Gosalia, Arka Dhar,  
 637 Ashley Pantuliano, Avi Nayak, Avital Oliver, Barret Zoph, Behrooz Ghorbani, Ben Leimberger,  
 638 Ben Rossen, Ben Sokolowsky, Ben Wang, Benjamin Zweig, Beth Hoover, Blake Samic, Bob  
 639 McGrew, Bobby Spero, Bogo Giertler, Bowen Cheng, Brad Lightcap, Brandon Walkin, Brendan  
 640 Quinn, Brian Guarraci, Brian Hsu, Bright Kellogg, Brydon Eastman, Camillo Lugaresi, Carroll  
 641 Wainwright, Cary Bassin, Cary Hudson, Casey Chu, Chad Nelson, Chak Li, Chan Jun Shern,  
 642 Channing Conger, Charlotte Burette, Chelsea Voss, Chen Ding, Cheng Lu, Chong Zhang, Chris  
 643 Beaumont, Chris Hallacy, Chris Koch, Christian Gibson, Christina Kim, Christine Choi, Christine  
 644 McLeavey, Christopher Hesse, Claudia Fischer, Clemens Winter, Coley Czarnecki, Colin Jarvis,  
 645 Colin Wei, Constantin Koumouzelis, Dane Sherburn, Daniel Kappler, Daniel Levin, Daniel Levy,  
 646 David Carr, David Farhi, David Mely, David Robinson, David Sasaki, Denny Jin, Dev Valladares,  
 647 Dimitris Tsipras, Doug Li, Duc Phong Nguyen, Duncan Findlay, Edede Oiwoh, Edmund Wong,  
 648 Ehsan Asdar, Elizabeth Proehl, Elizabeth Yang, Eric Antonow, Eric Kramer, Eric Peterson, Eric  
 649 Sigler, Eric Wallace, Eugene Brevdo, Evan Mays, Farzad Khorasani, Felipe Petroski Such, Filippo  
 650 Raso, Francis Zhang, Fred von Lohmann, Freddie Sulit, Gabriel Goh, Gene Oden, Geoff Salmon,  
 651 Giulio Starace, Greg Brockman, Hadi Salman, Haiming Bao, Haitang Hu, Hannah Wong, Haoyu  
 652 Wang, Heather Schmidt, Heather Whitney, Heewoo Jun, Hendrik Kirchner, Henrique Ponde

648 de Oliveira Pinto, Hongyu Ren, Huiwen Chang, Hyung Won Chung, Ian Kivlichan, Ian O’Connell,  
 649 Ian O’Connell, Ian Osband, Ian Silber, Ian Sohl, Ibrahim Okuyucu, Ikai Lan, Ilya Kostrikov, Ilya  
 650 Sutskever, Ingmar Kanitscheider, Ishaan Gulrajani, Jacob Coxon, Jacob Menick, Jakub Pachocki,  
 651 James Aung, James Betker, James Crooks, James Lennon, Jamie Kiros, Jan Leike, Jane Park,  
 652 Jason Kwon, Jason Phang, Jason Teplitz, Jason Wei, Jason Wolfe, Jay Chen, Jeff Harris, Jenia  
 653 Varavva, Jessica Gan Lee, Jessica Shieh, Ji Lin, Jiahui Yu, Jiayi Weng, Jie Tang, Jieqi Yu, Joanne  
 654 Jang, Joaquin Quinonero Candela, Joe Beutler, Joe Landers, Joel Parish, Johannes Heidecke, John  
 655 Schulman, Jonathan Lachman, Jonathan McKay, Jonathan Uesato, Jonathan Ward, Jong Wook  
 656 Kim, Joost Huizinga, Jordan Sitkin, Jos Kraaijveld, Josh Gross, Josh Kaplan, Josh Snyder, Joshua  
 657 Achiam, Joy Jiao, Joyce Lee, Juntang Zhuang, Justyn Harriman, Kai Fricke, Kai Hayashi, Karan  
 658 Singhal, Katy Shi, Kavin Karthik, Kayla Wood, Kendra Rimbach, Kenny Hsu, Kenny Nguyen,  
 659 Keren Gu-Lemberg, Kevin Button, Kevin Liu, Kiel Howe, Krithika Muthukumar, Kyle Luther,  
 660 Lama Ahmad, Larry Kai, Lauren Itow, Lauren Workman, Leher Pathak, Leo Chen, Li Jing, Lia  
 661 Guy, Liam Fedus, Liang Zhou, Lien Mamitsuka, Lilian Weng, Lindsay McCallum, Lindsey Held,  
 662 Long Ouyang, Louis Feuvrier, Lu Zhang, Lukas Kondraciuk, Lukasz Kaiser, Luke Hewitt, Luke  
 663 Metz, Lyric Doshi, Mada Aflak, Maddie Simens, Madelaine Boyd, Madeleine Thompson, Marat  
 664 Dukhan, Mark Chen, Mark Gray, Mark Hudnall, Marvin Zhang, Marwan Aljubeh, Mateusz Litwin,  
 665 Matthew Zeng, Max Johnson, Maya Shetty, Mayank Gupta, Meghan Shah, Mehmet Yatbaz,  
 666 Meng Jia Yang, Mengchao Zhong, Mia Glaese, Mianna Chen, Michael Janner, Michael Lampe,  
 667 Michael Petrov, Michael Wu, Michele Wang, Michelle Fradin, Michelle Pokrass, Miguel Castro,  
 668 Miguel Oom Temudo de Castro, Mikhail Pavlov, Miles Brundage, Miles Wang, Minal Khan, Mira  
 669 Murati, Mo Bavarian, Molly Lin, Murat Yesildal, Nacho Soto, Natalia Gimelshein, Natalie Cone,  
 670 Natalie Staudacher, Natalie Summers, Natan LaFontaine, Neil Chowdhury, Nick Ryder, Nick  
 671 Stathas, Nick Turley, Nik Tezak, Niko Felix, Nithanth Kudige, Nitish Keskar, Noah Deutsch, Noel  
 672 Bundick, Nora Puckett, Ofir Nachum, Ola Okelola, Oleg Boiko, Oleg Murk, Oliver Jaffe, Olivia  
 673 Watkins, Olivier Godement, Owen Campbell-Moore, Patrick Chao, Paul McMillan, Pavel Belov,  
 674 Peng Su, Peter Bak, Peter Bakkum, Peter Deng, Peter Dolan, Peter Hoeschele, Peter Welinder,  
 675 Phil Tillet, Philip Pronin, Philippe Tillet, Prafulla Dhariwal, Qiming Yuan, Rachel Dias, Rachel  
 676 Lim, Rahul Arora, Rajan Troll, Randall Lin, Rapha Gontijo Lopes, Raul Puri, Reah Miyara,  
 677 Reimar Leike, Renaud Gaubert, Reza Zamani, Ricky Wang, Rob Donnelly, Rob Honsby, Rocky  
 678 Smith, Rohan Sahai, Rohit Ramchandani, Romain Huet, Rory Carmichael, Rowan Zellers, Roy  
 679 Chen, Ruby Chen, Ruslan Nigmatullin, Ryan Cheu, Saachi Jain, Sam Altman, Sam Schoenholz,  
 680 Sam Toizer, Samuel Miserendino, Sandhini Agarwal, Sara Culver, Scott Ethersmith, Scott Gray,  
 681 Sean Grove, Sean Metzger, Shamez Hermani, Shantanu Jain, Shengjia Zhao, Sherwin Wu, Shino  
 682 Jomoto, Shirong Wu, Shuaiqi Xia, Sonia Phene, Spencer Papay, Srinivas Narayanan, Steve Coffey,  
 683 Steve Lee, Stewart Hall, Suchir Balaji, Tal Broda, Tal Stramer, Tao Xu, Tarun Gogineni, Taya  
 684 Christianson, Ted Sanders, Tejal Patwardhan, Thomas Cunningham, Thomas Degry, Thomas  
 685 Dimson, Thomas Raoux, Thomas Shadwell, Tianhao Zheng, Todd Underwood, Todor Markov,  
 686 Toki Sherbakov, Tom Rubin, Tom Stasi, Tomer Kaftan, Tristan Heywood, Troy Peterson, Tyce  
 687 Walters, Tyna Eloundou, Valerie Qi, Veit Moeller, Vinnie Monaco, Vishal Kuo, Vlad Fomenko,  
 688 Wayne Chang, Weiyi Zheng, Wenda Zhou, Wesam Manassra, Will Sheu, Wojciech Zaremba, Yash  
 689 Patil, Yilei Qian, Yongjik Kim, Youlong Cheng, Yu Zhang, Yuchen He, Yuchen Zhang, Yujia Jin,  
 690 Yunxing Dai, and Yury Malkov. GPT-4o System Card . *arXiv preprint arXiv:2410.21276*, 2024.

691 Steven Piantadosi and Richard Aslin. Compositional reasoning in early childhood. In *PLoS one*,  
 692 volume 11, September 2016.

693 QwenTeam. QwQ-32B: Embracing the Power of Reinforcement Learning, March 2025. URL  
 694 <https://qwenlm.github.io/blog/qwq-32b/>.

695 Parshin Shojaee, Iman Mirzadeh, Keivan Alizadeh, Maxwell Horton, Samy Bengio, and Mehrdad  
 696 Farajtabar. The illusion of thinking: Understanding the strengths and limitations of reasoning  
 697 models via the lens of problem complexity, 2025. URL <https://ml-site.cdn-apple.com/papers/the-illusion-of-thinking.pdf>.

698 Zayne Rea Sprague, Fangcong Yin, Juan Diego Rodriguez, Dongwei Jiang, Manya Wadhwa, Prasann  
 699 Singh, Xinyu Zhao, Xi Ye, Kyle Mahowald, and Greg Durrett. To CoT or not to CoT? Chain-of-  
 700 thought helps mainly on math and symbolic reasoning. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=w6nlcS8Kkn>.

702 Yiyou Sun, Shawn Hu, Georgia Zhou, Ken Zheng, Hannaneh Hajishirzi, Nouha Dziri, and Dawn  
 703 Song. Omega: Can llms reason outside the box in math? evaluating exploratory, compositional,  
 704 and transformative generalization. *arXiv preprint arXiv:2506.18880*, 2025.

705

706 Zhiqing Sun, Longhui Yu, Yikang Shen, Weiyang Liu, Yiming Yang, Sean Welleck, and Chuang  
 707 Gan. Easy-to-hard generalization: Scalable alignment beyond human supervision. In *The Thirty-  
 708 eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=qwgfh2fTtN>.

709

710 Derek Tam, Mohit Bansal, and Colin Raffel. Merging by matching models in task parameter  
 711 subspaces. *Transactions on Machine Learning Research*, 2024a. ISSN 2835-8856. URL <https://openreview.net/forum?id=qNGo6ghWFB>.

712

713 Derek Tam, Yash Kant, Brian Lester, Igor Gilitschenski, and Colin Raffel. Realistic evaluation of  
 714 model merging for compositional generalization. *arXiv preprint arXiv:2409.18314*, 2024b.

715

716 Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy  
 717 Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model.  
 718 [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca), 2023.

719

720 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay  
 721 Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cris-  
 722 tian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu,  
 723 Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn,  
 724 Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madijan Khabsa, Isabel  
 725 Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee,  
 726 Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra,  
 727 Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi,  
 728 Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh  
 729 Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen  
 730 Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic,  
 731 Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models.  
 732 *arXiv preprint arXiv:2307.09288*, 2023.

733

734 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V  
 735 Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models.  
 736 In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in  
 737 Neural Information Processing Systems*, 2022. URL [https://openreview.net/forum?id=\\_VjQ1MeSB\\_J](https://openreview.net/forum?id=_VjQ1MeSB_J).

738

739 Denise M. Werchan, Anne G.E. Collins, Michael Joshua Frank, and Dima Amso. 8-month-old infants  
 740 spontaneously learn and generalize hierarchical rules. *Psychological Science*, 26:805 – 815, 2015.  
 741 URL <https://api.semanticscholar.org/CorpusID:13823032>.

742

743 Han Wu, Yuxuan Yao, Shuqi Liu, Zehua Liu, Xiaojin Fu, Xiongwei Han, Xing Li, Hui-Ling Zhen,  
 744 Tao Zhong, and Mingxuan Yuan. Unlocking efficient long-to-short llm reasoning with model  
 745 merging. *arXiv preprint arXiv:2503.20641*, 2025.

746

747 Prateek Yadav, Derek Tam, Leshem Choshen, Colin Raffel, and Mohit Bansal. TIES-merging:  
 748 Resolving interference when merging models. In *Thirty-seventh Conference on Neural Information  
 749 Processing Systems*, 2023. URL <https://openreview.net/forum?id=xtaX3WyCj1>.

750

751 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,  
 752 Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin  
 753 Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang,  
 754 Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi  
 755 Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan,  
 756 Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report. *arXiv preprint  
 757 arXiv:2412.15115*, 2025.

758

759 Xi Ye and Greg Durrett. The unreliability of explanations in few-shot prompting for textual reasoning.  
 760 In *Proceedings of NeurIPS*, 2022.

756 Xi Ye, Srinivasan Iyer, Asli Celikyilmaz, Veselin Stoyanov, Greg Durrett, and Ramakanth Pasunuru.  
 757 Complementary explanations for effective in-context learning. In *Findings of the Association for*  
 758 *Computational Linguistics: ACL 2023*, Toronto, Canada, July 2023. Association for Computational  
 759 Linguistics.

760 Dingli Yu, Simran Kaur, Arushi Gupta, Jonah Brown-Cohen, Anirudh Goyal, and Sanjeev Arora.  
 761 SKILL-MIX: a flexible and expandable family of evaluations for AI models. In *The Twelfth*  
 762 *International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=Jf5gplvglq>.

763 Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting Dong, Keming Lu, Chuanqi Tan, Chang Zhou,  
 764 and Jingren Zhou. Scaling relationship on learning mathematical reasoning with large language  
 765 models. 2024. URL <https://openreview.net/forum?id=cij00f8u35>.

766 Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah Goodman. STaR: Bootstrapping Reasoning With  
 767 Reasoning. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.),  
 768 *Advances in Neural Information Processing Systems*, 2022. URL [https://openreview.net/forum?id=\\_3ELRdg2sgI](https://openreview.net/forum?id=_3ELRdg2sgI).

769 Haoyu Zhao, Simran Kaur, Dingli Yu, Anirudh Goyal, and Sanjeev Arora. Can models learn skill  
 770 composition from examples? In *The Thirty-eighth Annual Conference on Neural Information*  
 771 *Processing Systems*, 2024. URL <https://openreview.net/forum?id=1sLdprsrbmk>.

772 Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, and Zheyuan Luo. LlamaFactory:  
 773 Unified efficient fine-tuning of 100+ language models. In Yixin Cao, Yang Feng, and Deyi  
 774 Xiong (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational*  
 775 *Linguistics (Volume 3: System Demonstrations)*, pp. 400–410, Bangkok, Thailand, August  
 776 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-demos.38. URL  
 777 <https://aclanthology.org/2024.acl-demos.38/>.

778 Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans,  
 779 Claire Cui, Olivier Bousquet, Quoc V Le, and Ed H. Chi. Least-to-most prompting enables complex  
 780 reasoning in large language models. In *The Eleventh International Conference on Learning*  
 781 *Representations*, 2023. URL <https://openreview.net/forum?id=WZH7099tgfM>.

## 782 A A NOTE ON COMPOSING TASKS

783 There exist various possible ways to combine atomic tasks into a compositional task with the  
 784 combination function  $g$ . We simplify  $g$  into two types: (1) composite: the output of one atomic task is  
 785 used as part of the input of another task,  $g(\mathcal{T}_i, \mathcal{T}_j) = \mathcal{T}_i \circ \mathcal{T}_j$  or  $g(\mathcal{T}_i, \mathcal{T}_j) = \mathcal{T}_j \circ \mathcal{T}_i$ ; (2) concatenation:  
 786 the outputs of the two atomic tasks are concatenated using the same input,  $g(\mathcal{T}_i, \mathcal{T}_j) = \mathcal{T}_i \oplus \mathcal{T}_j$  or  
 787  $g(\mathcal{T}_i, \mathcal{T}_j) = \mathcal{T}_j \oplus \mathcal{T}_i$ . Among tasks evaluated in Section 4, the string manipulation and arithmetic  
 788 tasks need to be solved by a composite function, while the Skill-Mix task can be solved by either a  
 789 composite function or a concatenation function.

## 790 B DESIGN CHOICES FOR CONSTRUCTING COMPOSABLE CoT DATA

791 When designing the proxy prefix CoT, we would like to consider the following constraints. (1) We do  
 792 not assume any prior knowledge about what would possibly be put in the proxy prefix at inference  
 793 time; (2) We do not assume strong relevance between the proxy prefix CoT and the actual CoT, i.e.,  
 794 not all the information in the proxy prefix CoT is useful for predicting the CoT and the final answer.  
 795 Based on these considerations, we experiment with the following variants:

- 800 • **Random letters:** We sample random letters from the alphabet to form a sequence of random  
 801 lengths to simulate an *arbitrary* prefix CoT.
- 802 • **Random text from the prompt:** We sample random letters and words from the prompt  $\mathbf{q}$  to  
 803 form a sequence of random lengths to simulate a prefix CoT in a similar distribution as the  
 804 input distribution.

810  
 811 Table 6: Performance of atomic CoT models fine-tuned on different variants of proxy prefix on Llama  
 812 2-7B. Using random letters as the proxy prefix achieves the best out-of-domain performance when  
 813 evaluated with an unseen prefix at inference time.

Type of Proxy Prefix	Exact Match Accuracy	
	In Domain Prefix	Out-of-Domain Prefix
Random Letters	83.0	90.0
Random Text from the Prompt	86.4	82.5
Random Text from Web	90.6	70.0

814  
 815  
 816  
 817  
 818  
 819  
 820  
 821 • **Random text from web:** We sample random sentences from OpenWebText (Gokaslan &  
 822 Cohen, 2019) to simulate a prefix CoT drawn from the pretraining data distribution.

823  
 824 We evaluate these variants by fine-tuning models on Composable CoT datasets **that only the following**  
 825 **augmentation:**  $d' = q <\text{prefix}> [\text{proxy prefixes}] </\text{prefix}> <\text{suffix}> \text{ta} </\text{suffix}>$ . Note that this  
 826 is different from the setting discussed in Section 3.1 where the Composable CoT dataset consists  
 827 of other possible augmentations as well based on the sampling of the tags (e.g.,  $d' = q <\text{prefix}>$   
 828  $\text{ta} </\text{prefix}>$  when  $k = 1$ ). This experiment mainly aims at stress testing the model’s capability of  
 829 learning a single atomic task with a given proxy prefix CoT variant. We use the same hyperparameter  
 830 configurations for all proxy prefix variants for a given task.

831 We evaluate the fine-tuned models on the in-domain task in two settings: (1) *In-domain prefix*: we  
 832 append the same type of prefix as we have used for training to the end of the prompt of the in-domain  
 833 test example and evaluate the model on it; (2) *Out-of-domain prefix*: we randomly sample a prefix  
 834 from the other two variants and append it to the end of the prompt of the in-domain test example and  
 835 evaluate the model on it. We run experiments on the three string manipulation and arithmetic tasks  
 836 and report the average performance. Table 6 shows that although using random letters as the proxy  
 837 prefix leads to the worst in-domain performance, it generalizes the best to out-of-domain prefixes,  
 838 which is a more desirable behavior.

## 839 C DETAILS OF STRING MANIPULATION AND ARITHMETIC TASKS

840 **Next letter in alphabet** We synthetically generate data for Next letter in alphabet. We randomly  
 841 sample letters from the English alphabet of a random length and concatenate them into a sequence.  
 842 Then we extract the last letter from the sequence and derive the next letter following it in the alphabet.  
 843 An example can be found in Example C.2. We automatically generate a chain-of-thought for each  
 844 generated problem, using a fixed template shown in Example C.2.

845 **ASCII multiplication** Similarly, we randomly sample letters from the English alphabet of a  
 846 random length and concatenate them into a sequence. Then, we randomly sample another letter  $s$   
 847 and randomly sample an integer  $a \in \{1, \dots, 9\}$ . We find the ASCII value of  $s$  as  $f(s)$  and compute  
 848 the product  $a f(s)$  as the gold answer. An example can be found in Example C.3. We automatically  
 849 generate a chain-of-thought for each generated problem, using a fixed template shown in Example C.3.

850 **Letter concatenation** We follow Wei et al. (2022) to generate the dataset by randomly sampling  
 851 from the most popular first and last names in the United States and the United Kingdom from  
 852 <https://namecensus.com> and randomly concatenating them into a sequence of names. While  
 853 the original task in Wei et al. (2022) only requires concatenating the last letter of each name together,  
 854 we raise the difficulty level by randomly asking for concatenations of the first, second, second-to-last,  
 855 or the last letter. An example can be found in Example C.1. The CoT template is also shown in  
 856 Example C.1.

857 **Compositional tasks** We synthetically construct the compositional tasks of the string manipulation  
 858 and arithmetic tasks in similar procedures as used to generate the atomic data. An example of  
 859 next letter + ASCII multiplication can be found in Example C.4, concatenation + next letter in  
 860 Example C.5, and concatenation + multiplication in Example C.6. We made a design decision to

864 exclude one variant of concatenation + next letter that concatenates the last letter of each word and  
 865 finds the next letter following the last letter in the concatenated sequence; this variant can be solved  
 866 by the reasoning shortcut of only applying Next letter in alphabet rather than a composition of both.  
 867

### 868 C.1 Atomic Task Example: Letter Concatenation Example

869  
 870 [Instruction]  
 871 Take the second-to-the-last letter of each word in the sequence  
 872 and concatenate them in lower case: Tequan Monjur Khia  
 873 Jodi-leigh answer

874 [Chain-of-Thought + Answer String]  
 875 The second-to-the-last letter of the 1st word is a. The  
 876 second-to-the-last letter of the 2nd word is u. The  
 877 second-to-the-last letter of the 3rd word is i. The  
 878 second-to-the-last letter of the 4th word is g. So the answer  
 879 is auig.

880  
 881 [Answer String]  
 882 auig

### 883 C.2 Atomic Task Example: Next letter in alphabet

884  
 885 [Instruction]  
 886 Find the Next letter in alphabet following the last letter in  
 887 the sequence: wqsisibnnicdlpwqbnoicdcxcxrfoilpcbnixuc  
 888 bssszejxuzods answer:

889 [Chain-of-Thought + Answer String]  
 890 The last letter is s, and the letter following it in alphabet is  
 891 t. So the answer is t.

892  
 893 [Answer String]  
 894 t

### 895 C.3 Atomic Task Example: ASCII Multiplication

896  
 897 [Instruction]  
 898 Find the ASCII value of the letter after '<letter>' and multiply  
 899 the ASCII value by 2: byaxaxcpoteznwnwseselyjlretx  
 900 txcbfvmfezbycplymfotjbfv  
 901 jlhotzjbjcpycbtzhorepyjckofj <letter> d answer:

902 [Chain-of-Thought + Answer String]  
 903 The ASCII value of the letter d is 100, and multiplying the  
 904 ASCII value by 2 gives us 200. So the answer is 200.

905  
 906 [Answer String]  
 907 200

### 908 C.4 Compositional Task Example: Next letter + ASCII Multiplication

909  
 910 [Instruction]  
 911 Find the ASCII value of the Next letter in alphabet following  
 912 the last letter in the sequence and multiply the ASCII value by  
 913 5: knnxqsvshqugxfuqljumsbihgqvihnxuufuknxvumuuppkshljqsbkiz  
 914 answer:

918 [Answer String]  
 919 485  
 920

### 921 C.5 Compositional Task Example: Concatenation + Next Letter

923 [Instruction]  
 924 Take the second-to-the-last letter of each word in the sequence,  
 925 concatenate them in lower case, and find the Next letter in  
 926 alphabet following the last letter in the sequence of the  
 927 concatenated letters: Tyjai Ahijah Denzil Amine answer:

929 [Answer String]  
 930 ○

### 931 C.6 Compositional Task Example: Concatenation + Multiplication

933 [Instruction]  
 934 Take the second-to-the-last letter of each word in the sequence,  
 935 concatenate them in lower case, then find the ASCII value of  
 936 the last letter in the sequence of the concatenated letters,  
 937 and multiply the ASCII value by 3: Zarriah Amylee Li Javarie  
 938 answer:

940 [Answer String]  
 941 315

## 943 D DETAILS OF SKILL-MIX TASKS

### 944 D.1 MODIFICATIONS OF SKILL-MIX

948 We adapt the Skill-Mix dataset from Yu et al. (2024). For each example, the model is given a natural  
 949 language skill, its definition, an example of the skill, and a topic to focus on, and the model needs  
 950 to write a grammatical sentence to demonstrate the skill on the topic. Because we mainly focus on  
 951 pairwise composition, we only use the  $k = 2$  and  $k = 1$  composition sets of the Skill-Mix data. We  
 952 apply the following modifications to the dataset to fit our setting of compositional reasoning.

- 953 1. Filtering the categories of skills: We keep examples with skills of the rhetorical and literary  
 954 categories out of the five categories from the original dataset. This is because the rhetorical  
 955 and literary skills have the least overlap while other categories have more (e.g. the logical  
 956 and rhetorical skills have a large body of overlaps).
- 957 2. Removing the requirements of post-hoc explanation and refinement from the prompt. The  
 958 original dataset evaluates models by prompting the models to first write a sentence, provide  
 959 an explanation for the written sentence, and then do another round of refinement based on  
 960 feedback from the grader (an LLM-as-a-judge). To fit into our setting of chain-of-thought  
 961 reasoning and direct answering, we remove these irrelevant elements in the prompt.
- 962 3. Using a public test set: The original evaluation of Yu et al. (2024) hides the test set from the  
 963 public and models can only be evaluated based on API calls to the hidden test set. To ensure  
 964 reproducibility of our results, we use a public test set collected by Zhao et al. (2024).

966 As an open-ended generation task, Skill-Mix does not have a single ground truth sentence. Zhao  
 967 et al. (2024) collect a silver-labeled dataset of sentences and their post-hoc explanations generated  
 968 by GPT-4o (OpenAI et al., 2024) and graded as correct by the grader. We further prompt GPT-4o to  
 969 transform the (sentence, explanation) pairs into (CoT, sentence) pairs to resolve potential coreference  
 970 issues. Examples of the prompt, CoT traces, and sentence can be found in Examples D.1, D.2, and  
 971 D.3.

972  
973**D.1 Atomic Task Example: Skill-Mix-Literary**974  
975  
976  
977  
978  
979  
980  
981  
982

[Instruction]

Greetings! I am interested in natural language processing and I was wondering if you could help me generate an example of text that illustrates multiple skills in semantics or syntax. The example should be a minimal natural piece of text with up to a few lines in the context of Vikings that illustrates all of the following skills: oxymoron. Please keep the text as short as possible, and make sure the concepts can be found fully from the text.

983  
984  
985  
986  
987  
988  
989  
990  
991  
992

For reference, here are the definitions and examples for the concepts:

\*\*oxymoron\*\*: A figure of speech in which apparently contradictory terms appear in conjunction. For example, "The hall was filled with a deafening silence." Here, the phrase "deafening silence" is an oxymoron.

Please start the minimal natural piece of text with 'Answer:'.

Thanks very much!

993  
994  
995  
996  
997  
998  
999  
1000  
1001  
1002  
1003  
1004

[Chain-of-Thought + Answer String]

Explanation: The goal is to construct a sentence that showcases the use of an oxymoron while remaining relevant to the Viking theme. To achieve this, we focus on creating a phrase that contains seemingly contradictory terms. We begin with the context of Vikings enduring harsh weather conditions. The word "fiery" conveys an intense heat, while "frost" indicates severe cold. When these terms are combined into "fiery frost," they form an oxymoron as they depict both extreme temperature sensations simultaneously. This juxtaposition effectively creates a semantic paradox, leading to vivid imagery and emphasizing the harsh conditions faced by Vikings.

1005  
1006  
1007  
1008  
1009

Answer: "In the bitter cold, the Viking felt a fiery frost."

[Answer String]

Answer: "In the bitter cold, the Viking felt a fiery frost."

1010  
1011**D.2 Atomic Task Example: Skill-Mix-Rhetorical**1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020

[Instruction]

Greetings! I am interested in natural language processing and I was wondering if you could help me generate an example of text that illustrates multiple skills in semantics or syntax. The example should be a minimal natural piece of text with up to a few lines in the context of Hiking that illustrates all of the following skills: begging the question or assuming the conclusion. Please keep the text as short as possible, and make sure the concepts can be found fully from the text.

1021  
1022  
1023  
1024  
1025

For reference, here are the definitions and examples for the concepts:

\*\*begging the question or assuming the conclusion\*\*: A circular argument that paraphrases the question. For example, "Of course the earth is round; people have known this for hundreds of

1026 years."

1027

1028 Please start the minimal natural piece of text with 'Answer:'.

1029

1030 Thanks very much!

1031

1032 [Chain-of-Thought + Answer String]

1033 Explanation: To illustrate the concept of begging the

1034 question or assuming the conclusion within the context of

1035 hiking, we start by identifying a statement that relies on

1036 its own conclusion as a premise. First, recognize that the

1037 claim "hiking is beneficial" needs supporting evidence to be

1038 considered valid. However, if we use the phrase "because it's

1039 good for your health" as justification, we notice that it merely

1040 reiterates the initial claim using slightly different wording.

1041 This is because declaring something "beneficial" inherently

1042 implies a positive impact, such as being "good for your health."

1043 Thus, the reasoning becomes circular, as it depends on the same

1044 assumption it seeks to prove.

1045

1046 Answer: "Hiking is beneficial because it's good for your

1047 health."

1048 [Answer String]

1049 Answer: "Hiking is beneficial because it's good for your

1050 health."

### D.3 Compositional Task Example: Skill-Mix Literary + Rhetorical

1051

1052

1053 [Instruction]

1054 Greetings! I am interested in natural language processing and

1055 I was wondering if you could help me generate an example of text

1056 that illustrates multiple skills in semantics or syntax. The

1057 example should be a minimal natural piece of text with up to a

1058 few lines in the context of Vikings that illustrates all of the

1059 following skills: anaphora resolution, begging the question

1060 or assuming the conclusion. Please keep the text as short as

1061 possible, and make sure the concepts can be found fully from the

1062 text.

1063

1064 For reference, here are the definitions and examples for the

1065 concepts:

1066 **\*\*anaphora resolution\*\*:** Resolving the antecedent of a pronoun

1067 or noun phrase. For example, "The car is falling apart, but

1068 it still works." Here, "it" is the anaphor and "car" is the

1069 antecedent.

1070 **\*\*\*begging the question or assuming the conclusion\*\*\*:** A

1071 circular argument that paraphrases the question. For example,

1072 "Of course the earth is round; people have known this for

1073 hundreds of years."

1074

1075 Please start the minimal natural piece of text with 'Answer:'.

1076

1077 Thanks very much!

1078 [Answer String]

1079 Answer:

1080  
1081  
1082

The Viking chief, undefeated thanks to his ship, asserted, “It remains unconquered because it’s the ‘Indomitable’.”

1083

## 1084 D.2 EVALUATION METRICS

1085

We use GPT-4o-mini as the LLM-as-a-judge to grade the generated sentence using the exact same grading rubric as provided by Yu et al. (2024); the grader judges the quality of the sentence based on if: (1) All skills are used; (2) The sentence makes sense; (3) The sentence attaches to the given topic; (4) The sentence is short. We use the evaluation metrics for each generated sentence in Yu et al. (2024), including the following:

1091

1. **Full Marks:** 1 if the generated sentence satisfies all four criteria above and 0 otherwise.
2. **Skill Fraction:** The fraction of skills being demonstrated if all the other three criteria are satisfied; 0 otherwise

1095

We aggregate these metrics by averaging over all generated responses. In general, full marks evaluate the model’s capability of writing a perfect sentence for the task, while skill fraction evaluates how good the model is at handling skills given that it is good at the other stylistic capabilities. We use Curator (Marten et al., 2025) for an efficient implementation of the evaluation pipeline.

1099

1100

## 1101 D.3 RFT ON SKILL-MIX TASKS WITH RATIONALIZATION

1102

1103

1104

1105

1106

1107

1108

1109

For open-ended generation tasks like Skill-Mix, it is hard to only use the reference answer to verify the correctness of the responses sampled from a model. Thus, we use rationalization to perform rejection sampling fine-tuning for Skill-Mix: we first append the direct answer label to the end of the prompt and sample post-hoc explanations for the given answer from the model; because  $M_{\text{comb}}$  is optimized to generate an answer following a CoT, we extract the generated answer following the generated explanation and filter out explanations whose following answer is not the same as the provided gold answer; finally, we use the accepted explanations as surrogates for CoT to form the RFT data.

1110

1111

## 1112 E SINGLE-TASK LEARNING PERFORMANCE

1113

1114

1115

1116

1117

1118

We report the single-task learning performance of the atomic CoT models by evaluating them on the in-domain atomic tasks. We would like the atomic tasks to be easy to learn to reflect the practical settings where we train models on basic, easy-to-learn skills and generalize to harder, unseen tasks. The training data conditions and hyperparameters for training can be found in Appendix G. Table 7 shows that all atomic tasks we evaluate are learnable within a small amount of training data ( $N_i, N_j \leq 500$ ).

1119

1120

1121

1122

In addition, we observe that training on ComposableCoT or StandardCoT does not lead to consistent differences in atomic CoT performance, while the exception is on Skill-Mix-Rhetorical for Llama 2-7B where fine-tuning on ComposableCoT outperforms fine-tuning on StandardCoT by a large margin.

1123

1124

1125

## 1126 F BASE MODEL PERFORMANCE ON EVALUATION TASKS

1127

1128

To confirm that our task design includes evaluation tasks that are less commonly seen in the pretraining data of LLMs, we evaluate the perplexity of the task datasets.

1129

1130

1131

1132

1133

We compare the datasets of the string manipulation and arithmetic datasets used in our experiments with mathematical reasoning data (GSM8K (Cobbe et al., 2021)), and instruction following data (Alpaca (Taori et al., 2023)) in terms of perplexity: we compute the average perplexity score of pre-trained LLMs over the concatenation of the question and ground-truth chain-of-thought response to examine how predictable the task is; the lower the perplexity is, the more predictable and the harder to learn the task is. We also include the perplexity over the pretraining corpus as a reference point.

1134  
 1135 Table 7: Single-task learning performance by evaluating the atomic CoT models on the in-domain  
 1136 atomic tasks.

1137 CoT Format	1138 Next Letter EM	1138 ASCII Mult EM	1138 Concat EM	1138 Skill-Mix Literary Full Marks	1138 Skill Fraction	1138 Skill-Mix Rhetorical Full Marks	1138 Skill Fraction
Llama 2-7B							
1140 StandardCoT	100.0	85.7	83.0	63.5	63.5	53.3	53.3
1141 ComposableCoT	95.0	86.0	77.0	71.4	71.4	72.4	72.4
Qwen 2.5-7B							
1144 StandardCoT	90.0	99.0	77.4	77.4	77.6	70.5	70.5
1145 ComposableCoT	99.4	99.7	77.3	77.4	77.6	76.7	81.9

1146  
 1147 Table 8: Perplexity of the base models over the task datasets. For the string manipulation and arithmetic  
 1148 tasks, the perplexity score is averaged over the 3 atomic tasks and the 3 pairwise compositional  
 1149 tasks. Our evaluation datasets include text that is less predictable under pre-trained LLMs than other  
 1150 similar tasks.

	Task Dataset	Pretraining WikiText	Math Reasoning GSM8K	Instruction Following Alpaca	String Manipulation And Arithmetic Avg.
1154 Model	Llama 2-7B	4.77	2.54	3.84	15.97
1155	Qwen 2.5-7B	5.93	2.38	4.68	7.02

1156  
 1157 Table 8 indicates that our selected tasks consist of text that is less typical under pre-trained language  
 1158 models than other similar tasks, particularly other popular reasoning tasks. The higher perplexity is  
 1159 likely due to these tasks requiring the model to operate on letters rather than words.

## 1162 G TRAINING CONFIGURATIONS

### 1164 G.1 GENERAL CONFIGURATIONS

1166 We conduct all fine-tuning experiments with LoRA(Hu et al., 2022) using the following set of  
 1167 hyperparameters: we use a rank of 8,  $\alpha = 16$ , and a dropout rate of 0.2 to prevent overfitting. We  
 1168 apply LoRA adapters to all linear modules, including the attention matrices  $Q$ ,  $K$ ,  $V$  and MLP  
 1169 matrices of all layers. We use bfloat16 precision for training and we use the efficient implementation  
 1170 of LoRA by LlamaFactory (Zheng et al., 2024). We use a training batch size of 4 and train for  
 1171 5 epochs for all experiments that share the same number of training data; for methods that use a  
 1172 potentially smaller amount of training data (e.g. RFT methods usually get fewer data examples than  
 1173 the number of compositional training data provided, depending on how many correct responses we  
 1174 can sample from the model), we adjust the batch size to match the number of steps.

### 1175 G.2 CONFIGURATION FOR REJECTION SAMPLING FINE-TUNING

1177 In addition to the sampling parameters (see Section 4), we consider the following configuration of  
 1178 RFT for sampling the correct responses: if the model generates multiple correct responses for a given  
 1179 question, we only randomly select *one* of them to be added into the RFT dataset  $D_{RFT}$ . In this way  
 1180 we ensures the diversity of examples in  $D_{RFT}$  so that the dataset will not be filled with samples from  
 1181 a small set of examples where the model is good at.

### 1183 G.3 HYPERPARAMETERS: LEARNING RATE

1185 We find in preliminary experiments that learning rate is the most important hyperparameter for the  
 1186 fine-tuning experiments of our interest. We perform hyperparameter sweeps for learning rate over  
 1187 the space of  $\{5e-3, 1e-3, 5e-4, 1e-4, 5e-5\}$  on a validation set for each experiment. The  
 1188 optimal learning rate for each method for the experiments with compositional supervision in Table 9.

1188  
1189 Table 9: Optimal learning rate for each method in the experiments with compositional supervision.  
1190

Category	Method	Next Letter + Mult	Concat + Next Letter	Concat + Mult	Skill-Mix Literary + Rhetorical
Llama 2-7B					
SFT	SFT on Base Model	1e-3	1e-3	5e-4	5e-4
	CFT on StandardCoT-Merge	1e-3	5e-4	1e-4	1e-4
	CFT on StandardCoT-MTL	1e-4	1e-4	1e-4	1e-3
MTL	StandardCoT + Comp Answer	1e-3	5e-4	1e-3	5e-4
	StandardCoT-Merge	-	1e-3	1e-3	5e-4
RFT	ComposableCoT-Merge (Ours)	1e-4	1e-4	1e-3	1e-3
	Qwen 2.5-7B				
SFT	SFT on Base Model	-	1e-3	1e-3	5e-4
	CFT on StandardCoT-Merge	-	5e-4	5e-4	1e-4
	CFT on StandardCoT-MTL	-	1e-3	1e-3	1e-3
MTL	StandardCoT + Comp Answer	-	5e-4	5e-4	1e-3
	StandardCoT-MTL	-	1e-3	1e-4	5e-4
RFT	ComposableCoT-MTL (Ours)	-	1e-3	1e-3	5e-4

1200  
1201  
1202  
1203  
1204  
1205  
1206  
1207  
1208  
1209  
1210 Table 10: Data conditions for each task used for our evaluation.  
1211

		# Train	# Test
Atomic Tasks	Next Letter	100	700
	ASCII Mult	100	700
	Concat	500	700
	Skill-Mix Literary	100	126
	Skill-Mix Rhetorical	100	105
Compositional Tasks	Next Letter + Mult	100	700
	Concat + Next Letter	100	504
	Concat + Mult (Llama 2-7B)	500	700
	Concat + Mult (Qwen 2.5-7B)	100	700
	Skill-Mix Literary + Rhetorical	100	245

1224  
1225 G.4 HYPERPARAMETERS: MODEL MERGING1226  
1227 For methods that use model merging as the combination, we use Task Arithmetic (Ilharco et al.,  
1228 2023b) to combine the atomic CoT models. We perform a hyperparameter sweep for the scalars  $\alpha$   
1229 and  $\beta$  over the space of  $\alpha \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$  and  $\beta = 1 - \alpha$  on a validation  
1230 set for each task.1231  
1232 H DATA STATISTICS1233  
1234 H.1 GENERAL DATA CONDITIONS FOR EXPERIMENTS1235  
1236 Table 10 summarizes the number of training data and test data used in the evaluations in Sections 5.1  
1237 and 5.2. Note that for letter concatenation + multiplication we have two sizes of the compositional  
1238 training data for Llama 2-7B and Qwen 2.5-7B: this is because all methods on Llama 2-7B perform  
1239 poorly on zero-shot evaluation for this task and we need a slightly larger amount of compositional  
1240 training data so that different methods can start to show distinguishable compositional task per-  
1241 formance from each other. Regardless, we still consider 500 to be a reasonably small amount of training  
1242 data, satisfying our ideal data conditions defined earlier.

1242  
 1243 Table 11: The detailed breakdown of the number of training data used by each method in the zero-shot  
 1244 setting.  $N_i$  and  $N_j$  denotes the number of training data from the atomic tasks  $\mathcal{T}_i$  and  $\mathcal{T}_j$  seen by the  
 1245 method during training.

	Method	$N_i$	$N_j$
1248 1249 1250 1251 1252 1253 1254	StandardCoT-Merge	0	0
	ComposableCoT-Merge	100	100
	StandardCoT-MTL	100	100
	ComposableCoT-MTL	100	100
1255 1256 1257 1258 1259 1260 1261	StandardCoT-Merge	500	100
	ComposableCoT-Merge	500	100
	StandardCoT-MTL	500	100
	ComposableCoT-MTL	500	100

1257  
 1258 Table 12: The detailed breakdown of the number of training data used by each method with compositional  
 1259 supervision for Qwen 2.5-7B.  $N_i$  and  $N_j$  denotes the number of training data from the atomic  
 1260 tasks  $\mathcal{T}_i$  and  $\mathcal{T}_j$  seen by the method during training.  $N_{(i,j)}$  denotes the number of compositional  
 1261 answer data seen during training.

	Method	$N_i$	$N_j$	$N_{(i,j)}$
1262 1263 1264 1265 1266 1267 1268 1269 1270 1271	SFT on Base Model	0	0	100
	CFT on StandardCoT-Merge	100	100	100
	CFT on StandardCoT-MTL	100	100	100
	Next Letter + Mult;	100	100	100
	Skill-Mix Literary + Rhetorical	100	100	100
	RFT on StandardCoT-Merge	100	100	100
	RFT on ComposableCoT-Merge	100	100	100
	RFT on StandardCoT-MTL	100	100	100
	RFT on ComposableCoT-MTL	100	100	100
	SFT on Base Model	0	0	100
1272 1273 1274 1275 1276 1277 1278 1279	CFT on StandardCoT-Merge	500	100	100
	CFT on StandardCoT-MTL	500	100	100
	MTL on StandardCoT + Comp Answer	500	100	100
	Concat + Next Letter;	500	100	100
	Concat + Mult	500	100	100
	RFT on StandardCoT-Merge	500	100	100
	RFT on ComposableCoT-Merge	500	100	100
	RFT on StandardCoT-MTL	500	100	100
1280 1281	RFT on ComposableCoT-MTL	500	100	100

## H.2 TRAINING DATA USED BY EACH METHOD

1282 We show a detailed breakdown in Table 11 of the number of training data used by each zero-shot  
 1283 method for both models and in Table 12 for Qwen 2.5-7B by each method with compositional answer  
 1284 data in the experiments in Section 5.2. Note that the statistics for Llama 2-7B in the setting with  
 1285 compositional supervision are mostly the same except  $N_{(i,j)} = 500$  for concat + next letter and  
 1286 concat + mult.

## I DETAILS OF THREE-WAY COMPOSITIONS

### I.1 DATA

1294 We include 700 test examples for Letter Concat + Next Letter + Mult (*String Tasks*), and 245 test  
 1295 examples for Skill-Mix Literary + Rhetorical + Logical (*Skill-Mix*). For *Skill-Mix*, we additionally  
 1296 train an atomic model for Skill-Mix-Logical with 100 training examples.

1296  
 1297 Table 13: Summary of methods evaluated in the zero-shot compositional evaluation and the compo-  
 1298 sition with limited compositional answer data.“Merge” stands for model merging; “MTL” stands  
 1299 for multitask learning; “CFT” stands for continued fine-tuning; “RFT” stands for rejection sampling  
 1300 fine-tuning. “-” means the property is not applicable to the method (e.g. *MTL on Standard + Comp*  
 1301 *Answer* mixes Standard CoT data with compositional answer data, and trains a single MTL model  
 1302 from the pretrained model, so there is no atomic CoT model trained or combined.)

Method	# Atomic CoT Models Trained	Atomic CoT Format	Combination Method	Model trained on Compositional Data	How is Compositional Data Used
<i>Zero-shot Evaluation</i>					
StandardCoT-Merge	2	Standard	Merge	-	-
<b>ComposableCoT-Merge (Ours)</b>	2	Composable	Merge	-	-
<i>Evaluation with Limited Compositional Answer Data</i>					
CFT on StandardCoT-Merge	2	Standard	Merge	StandardCoT-Merge	CFT
CFT on StandardCoT-MTL	1	Standard	MTL	StandardCoT-MTL	CFT
MTL on StandardCoT + Comp Answer	-	Standard	-	Pretrained Model	Mix with Atomic CoT data and MTL
RFT on StandardCoT-Merge	2	Standard	Merge	StandardCoT-Merge	RFT
<b>RFT on ComposableCoT-Merge (Ours)</b>	2	Composable	Merge	ComposableCoT-Merge	RFT
RFT on StandardCoT-MTL	1	Standard	MTL	StandardCoT-MTL	RFT
<b>RFT on ComposableCoT-MTL (Ours)</b>	1	Composable	MTL	ComposableCoT-MTL	RFT

## I.2 TRAINING AND INFERENCE

1321 **Training** We use the same data augmentation scheme to create atomic CoT training data as the one  
 1322 we use for two-way composition in Section 4. This means that we append only one proxy prefix to  
 1323 the prompt. The general scheme can insert at most  $n - 1$  proxy prefixes at the end of the prompt  
 1324 for  $n > 2$ , but we found that the test-time generalization scheme described in **Instantiation of Tags**  
 1325 under Section 3.1 works as well: adding only one proxy prefix achieves comparable compositional  
 1326 performance to adding two proxy prefixes while being more efficient during training, since the  
 1327 training data length is shorter. Thus, we experiment with the latter scheme.

1328 **Inference** We use the same inference strategy specified in **Data Construction** under Section 4: for  
 1329 zero-shot inference, we first sample a response from  $M_{\text{comb}}$ . Then, we repeat the following *twice*:  
 1330 we append `<suffix>` to the end of the generated response when it stops generation, and continue  
 1331 generation until the model stops again.

## J FULL RESULTS FOR THE QUALITY ANALYSIS OF THE GENERATED COTs

1336 Table 14 includes the full results of the quality analysis of the generated CoTs using both multi-task  
 1337 learning (MTL) and model merging as the combination methods for atomic CoT models. Given the  
 1338 same combination method, combined Composable CoT models generate responses including both  
 1339 atomic CoT patterns more frequently than combined atomic CoT models.

## K ERROR ANALYSES

1343 In addition to not being able to perform the individual atomic task correctly, we show three types of  
 1344 common errors made by ComposableCoT variants in the zero-shot compositional evaluation setting.

- 1346 1. Example K.1 shows an example where the generated CoT is only able to replicate CoT of  
     one atomic CoT and repeat the same CoT in the prefix and suffix.
- 1347 2. Example K.2 shows an example where the combined model fails to continue generation  
     after generating the prefix CoT. This is a common error for Composable models combined  
     with model merging.

1350  
 1351 Table 14: Intrinsic evaluation of the generated CoTs from atomic CoT models evaluated on the  
 1352 compositional task in the zero-shot setting. “%  $\mathcal{T}_1$  CoT” denotes the percentage of generated  
 1353 responses that use the CoT format of the first atomic task of the composition, and likewise for the  
 1354 second.  $\dagger$  denotes that the ComposableCoT method has a significantly higher “% Both CoT” than  
 1355 the StandardCoT counterpart at the 0.01 level using a paired bootstrap test. Combined Composable  
 1356 CoT models generate responses including both atomic CoT patterns more frequently than combined  
 1357 atomic CoT models.

	Method	Performance	% $\mathcal{T}_1$ CoT	% $\mathcal{T}_2$ CoT	% Both CoT
1359 1360 1361 1362	StandardCoT-Merge	70.4	85.3	95.1	85.3
	Next Letter + Mult	95.4 3.6	100.0 0.0	100.0 100.0	$\dagger$ <b>100.0</b> 0.0
	ComposableCoT-Merge	96.3	98.9	100.0	$\dagger$ 98.9
	StandardCoT-MTL				
1363 1364 1365 1366	StandardCoT-Merge	77.0	90.3	98.7	90.0
	Concat + Next Letter	75.4 72.1	91.6 99.7	100.0 32.1	<b>91.6</b> 32.1
	ComposableCoT-Merge	74.3	100.0	83.1	$\dagger$ 81.3
	StandardCoT-MTL				
1367 1368 1369 1370	StandardCoT-Merge	54.8	100.0	99.4	<b>99.4</b>
	Concat + Mult	19.2 60.9	44.6 100.0	60.5 66.7	17.7 66.7
	ComposableCoT-Merge	63.3	100.0	85.9	$\dagger$ 85.0
	StandardCoT-MTL				
1371 1372 1373	Skill-Mix	29.8	60.0	59.2	35.9
	Literary + Rhetorical	39.6 42.0	64.1 65.3	66.9 58.0	$\dagger$ 43.3 37.6
	ComposableCoT-Merge	49.0	64.5	65.7	$\dagger$ 42.0
	StandardCoT-MTL				

1374  
 1375  
 1376 3. Example K.3 shows an example where the combined model uses the wrong atomic CoT in  
 1377 the prefix that should have been used in the suffix.  
 1378

### 1379 K.1 Error Case: Replicating One Atomic CoT Pattern

1381 [Instruction]  
 1382 Take the first letter of each word in the sequence, concatenate  
 1383 them in lower case, and find the next letter in alphabet  
 1384 following the last letter in the sequence of the concatenated  
 1385 letters: Jessika Berk Hamilton Stephanie answer:  
 1386 [Answer String]  
 1387 t  
 1388 [Model Output]  
 1389 <prefix> The last letter is e, and the letter following it in  
 1390 alphabet is f. So the answer is f.</prefix> <suffix> The last  
 1391 letter is e, and the letter following it in alphabet is f. So  
 1392 the answer is f.</suffix>  
 1393 [Predicted Answer]  
 1394 f

### 1395 K.2 Error Case: Failing to Continue Generation after Prefix CoT

1397 [Instruction]  
 1398 Take the last letter of each word in the sequence, concatenate  
 1399 them in lower case, then find the ASCII value of the last letter  
 1400 in the sequence of the concatenated letters, and multiply the  
 1401 ASCII value by 6: Ayat Amaris Heavenly Mariela answer:  
 1402 [Answer String]  
 1403 582

```

1404 [Model Output]
1405 <prefix> The last letter of the 1st word is t. The last letter
1406 of the 2nd word is s. The last letter of the 3rd word is y.
1407 The last letter of the 4th word is a. So the answer is tasy,
1408 and the ASCII value of the last letter in the sequence of the
1409 concatenated letters is 121, so the answer is 726.</prefix>
1410 <suffix>
1411 [Predicted Answer]
1412 726

```

### K.3 Error Case: Use the Wrong Order of Atomic CoTs

```

1413
1414 [Instruction]
1415 Take the second letter of each word in the sequence, concatenate
1416 them in lower case, and find the next letter in alphabet
1417 following the last letter in the sequence of the concatenated
1418 letters: D'shaun Takia Hiroki Arjun answer:
1419 [Answer String]
1420 s
1421 [Model Output]
1422 <prefix> the last letter is n, and the letter following it in
1423 alphabet is o. so the answer is o.</prefix> <suffix> the second
1424 letter of d is d, the second letter of t is a, the second letter
1425 of h is i, and the second letter of a is r. so the answer is
1426 dair.</suffix>
1427 [Predicted Answer]
1428 dair
1429

```

## L COMPUTATION RESOURCES

We conducted all our experiments primarily on a server with NVIDIA RTX A6000 48GB GPU, and a server with GPU of NVIDIA GH200 120GB. Each individual experiment can be run on a single A6000 of 48 GB within 1 hour.

## M LICENSING

We use the following publicly available datasets from prior works with open licenses.

**Skill-Mix** Yu et al. (2024) uses the Apache-2.0 license and data is available at: <https://huggingface.co/spaces/dingliyu/skillmix>.

**Letter concatenation** The dataset uses the Apache-2.0 license and the data is available at: <https://huggingface.co/datasets/ChilleD/LastLetterConcat>