

INSTRUCTION FOLLOWING WITHOUT INSTRUCTION TUNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Adaptations that seem deficient compared to instruction-tuning (finetuning on instruction-response pairs) can still implicitly yield instruction-following language models. We show that training solely on *responses*, without any corresponding instructions, yields instruction following. This suggests that instruction-tuning just needs to teach the desired distribution of responses. However, we then show that one can finetune on instruction-response pairs from a distribution unlike the test distribution of instructions, like just poetry, or just math, and still yield a model that broadly follows instructions. Instead of acting, e.g., just as math models, these single-task models sometimes behave more as general-purpose chatbots for, e.g., non-math instructions. To begin to explain this *implicit instruction tuning*, we hypothesize that simple changes to a language model’s distribution yield instruction following. We support this by *hand-writing* a rule-based adapter that yields instruction-following behavior in language models. The rules are to slowly increase the probability of ending the sequence, penalize repetition, and uniformly change 15 words’ probabilities. In summary, adaptations made without being designed to yield instruction following can do so *implicitly*.

1 INTRODUCTION

Instruction tuning, finetuning on a broad distribution of responses (e.g., *Tiramisu is made by...*) conditioned on instructions (e.g., *Give me a recipe for tiramisu*), yields instruction following from language models for a wide range of instructions (Ouyang et al., 2022). Prior work has shown that instruction tuning is sample-efficient, requiring as few as 1000 broad-domain instruction-response pairs (Zhou et al., 2023a) or a carefully crafted prompt and few-shot instruction-response examples (Lin et al., 2024). We take this a step further, exploring the idea that instruction following can be yielded from language models even *implicitly*, i.e., through methods not explicitly designed to do so. We discover two forms of adaptation that perform *implicit instruction tuning*, being seemingly deficient compared to explicit instruction tuning: (1) *response tuning*, training on only responses, and (2) *single-task finetuning*, training on data from a narrow domain of goals, like poetry generation.

We first demonstrate that *response tuning*, training on responses alone without conditioning on their instructions—is sufficient to yield instruction following (Section 4). In particular, using the LIMA dataset (Zhou et al., 2023a) for tuning, and evaluating on AlpacaEval 2, response-tuned models win roughly 43% of the time against a comparable instruction-tuned model, where equal performance would correspond to a 50% win rate. Response tuning provides no explicit information about the mapping from instructions to responses, only information about the distribution of desired responses. This suggests that an instruction-response mapping may be learned during pretraining, but all desirable responses are too low-probability to be generated. We support this, finding that pretrained models rank an instruction’s real response higher than a random other instruction’s response at the same rate as instruction-tuned models.

From the response tuning results, it seems that the crucial part of instruction tuning is to teach the distribution of desirable responses. However, we find that this is not crucial either. Finetuning on single-task, narrow-domain data, like mapping English requests to Python snippets (MBPP; Austin et al., 2021), or generating poetry from poem titles, *also* yields broad instruction-following behavior (Section 5). That is, despite training only to generate only Python code or only poetry, models generate biographies or recipes when instructed. For example, a poetry-tuned Llama-2-7B wins

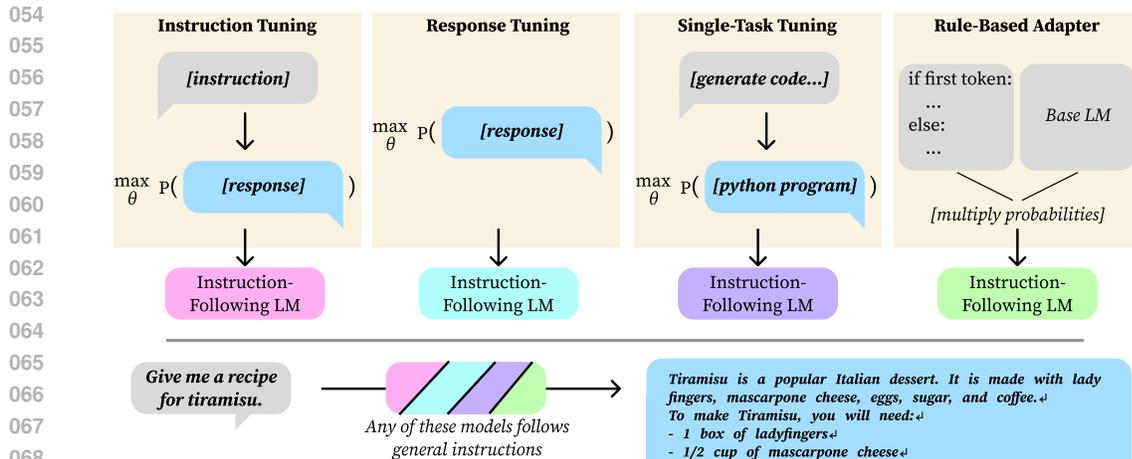


Figure 1: Instruction tuning trains a language model on responses conditioned on instructions. We find that (1) response tuning (estimating the probability of responses with no instructions), (2) single-task finetuning (e.g., code or poetry generation), and even (3) a simple rule-based adapter all yield language models with general instruction-following behavior.

23.7% in AlpacaEval 2 head-to-head against an instruction-tuned Llama-2-7B, whereas the base model wins 2.4% of the time. Qualitatively, we find that single-task finetuned models only adhere to the finetuning distribution on instructions similar to those finetuned on, and exhibit general instruction following behavior for other instructions.

These two results motivate us to ask why instruction following is yielded by such a wide range of adaptations. While Lin et al. (2024) showed that relatively few token decisions change between pretraining and instruction following, it is not obvious that it’s simple to determine which tokens to change. We hypothesize that very simple changes in conditional distributions can cause a language model to follow instructions. We validate this by hand-writing a rule-based language model with three rules that, when taken in a product with a pretrained language model, causes instruction following (Section 6). Our three rules are: slowly increasing the probability of ending the sequence, uniformly modifying 15 tokens’ likelihoods (e.g., I, <, should), and penalizing token repetition. Our rule-based product-of-experts with Llama-2-7B wins 24.4% against instruction-tuned Llama-2-7B, compared to the base model’s win rate of 2.4%.

In summary, our results show that adaptation methods not *intended* to yield instruction following may yet do so *implicitly*. When adapting language models for specific purposes, we might not expect that models would behave as general-purpose instruction following models on instructions dissimilar from those tuned on, but our results suggest that this will often be the case. Our rule-based adapter experiments begin to explain this, showing that very simple changes to a model’s distribution—not obviously tied to following instructions—also yield instruction following.¹

1.1 A NOTE ON PRETRAINING DATA

One critical question in the interpretation of our results is whether language models were instruction-tuned during pretraining (aka, midtraining,) as is becoming common (Bi et al., 2024; Dominguez-Olmedo et al., 2024). If this is the case, our results become less surprising. To mitigate this, we experiment with two open-weights language models: Llama-2-7B and OLMo-7B-Feb2024. The Llama-2-7B model is stronger than the OLMo-7B-Feb2024 model, but we have no guarantee against intentional instruction tuning. OLMo-7B-Feb2024 is weaker, but no instruction-tuning data was *intentionally* included in its pretraining.² These two models’ results lead to similar conclusions. However, for comparison, we also evaluate on one model we believe to have been midtrained: Gemma-2-9B (Team et al., 2024). Base Gemma-2-9B already roughly follows instructions, but our

¹Our code and data will be made available upon publication.

²Public communication; link hidden for anonymization.

108 experiments yield similar conclusions as for the other models, except that single-task finetuning isn't
 109 as competitive.
 110

111 2 RELATED WORK

112 **Designing instruction tuning datasets.** A few years before instruction tuning became common,
 113 the idea of a single model performing a range of tasks was explored by DecaNLP (McCann et al.,
 114 2018). In the LLM era, stronger pretrained models led to more adoption. In this era, early datasets
 115 also transformed a wide range of existing NLP tasks into a single common format (Mishra et al.,
 116 2022; Wang et al., 2022; Wei et al., 2022; Sanh et al., 2022). The introduction of ChatGPT in October
 117 of 2022 shifted focus onto the construction of instruction tuning datasets that reflect goals users
 118 might be interested in, as opposed to NLP tasks. Some are relatively small, with the intuition to
 119 teach nothing fundamentally new during instruction tuning (Taori et al., 2023; Zhou et al., 2023a).
 120 Others, which tend to be more performant, can be quite large and contain reasoning chains and other
 121 behaviors thought to be rare in pretraining (Iverson et al., 2023; Wang et al., 2023b;a; Yu et al., 2023);
 122 this has included considerable non-academic open efforts (Teknium, 2023).
 123
 124

125 **Studies on instruction tuning.** Our work contributes to a set of discoveries concerning how *little*
 126 it takes to get language models to follow instructions. Taori et al. (2023) showed that finetuning a
 127 Llama model on 52,000 instruction-response pairs caused it to follow instructions to an extent that
 128 was surprising at the time. Zhou et al. (2023a) showed that—with careful selection—one could use
 129 as little as 1,000 instruction-response pairs and achieve strong results. This strongly suggested that
 130 one does not need to cover all types of instructions during instruction tuning. Lin et al. (2024) then
 131 showed that, with careful prompting and in-context few-shot examples, a handful of examples could
 132 also lead to instruction following. Kung & Peng (2023) take the NaturalInstructions (Mishra et al.,
 133 2022) style of instruction, which includes a task description, an in-context example, and the data for
 134 the current example. They then remove either the task description or the in-context example (leaving
 135 the other), finding that performance degrades less than expected in either case. In concurrent work,
 136 Xie et al. (2024) find that synthesizing instruction-tuning data by using webtext as instructions and
 137 generating a response with GPT-4 also yields instruction-following models. Shi et al. (2024) find that
 138 training to maximize the joint likelihood of instruction and response improves over optimizing for
 139 the conditional likelihood of the response. Finally our work contributes to mechanistic discussions
 140 of instruction-tuning; for example, while Wu et al. (2024) find that instruction-tuning is crucial for
 141 teaching the model to look more at the instruction, the success of response-tuning suggests this may
 142 not be necessary.

143 **Out-of-distribution generalization.** Discussing out-of-distribution generalization is difficult to
 144 do precisely in the world of language pretraining, especially when we don't know what data went
 145 into the pretraining. Furthermore, it's often difficult to reason about the coverage of datasets with
 146 10+ trillion tokens (Section 1.1). In computer vision and vision-language models, various studies
 147 have shown that pretraining is the singularly important component in robustness to distribution shift
 148 (Carmon et al., 2019; Miller et al., 2021; Awadalla et al., 2022). In language as well, pretraining
 149 on diverse data is known to lead to better out-of-distribution performance (relative to the finetuning
 150 dataset) (Hendrycks et al., 2020). When we finetune, for example, on only poetry and test on recipes,
 151 this is out-of-distribution to the finetuning, but not to the pretraining.
 152

153 3 EXPERIMENT SETTING

154 A neural language model $p_\theta(\mathbf{x})$ is a distribution over strings $\mathbf{x} \in \mathcal{V}^*$, where \mathcal{V} is a finite vocabulary,
 155 and θ are learnable parameters in the neural network. Language models are trained to minimize the
 156 cross-entropy loss with a large corpus of text. These pretrained, or *base*, models serve as starting
 157 points, but they are typically adapted to directly respond to user queries.
 158
 159

160 **Instruction tuning.** Instruction tuning finetunes the parameters θ of a language model to adapt
 161 its behavior to respond to queries with relevant, helpful answers. Given a set of samples $D_{\text{ins}} =$
 $\{\text{instruction}_i, \text{response}_i\}_{i=1}^k$, of instructions and corresponding responses (each of which in \mathcal{V}^*),

instruction tuning optimizes for

$$\min_{\theta} \frac{1}{k} \sum_{i=1}^k -\log p_{\theta}(\text{response}_i \mid \text{instruction}_i). \quad (1)$$

Instruction formatting. In language model practice, the distinction between the instruction and the response—and thus what to compute the loss over or generate—is specified through formatting tokens in the input. We use the Tulu formatting of Wang et al. (2023b); Ivison et al. (2023); we present it here because instruction formatting may matter for how easy it is to yield instruction following behavior from language models:

```
BOS<|user|>
{instruction}
<|assistant|>
{response}EOS
```

Here, BOS and EOS are the model-specified special beginning- and end-of-sequence tokens. Though the semantics of the tags are suggestive of instruction-following, we keep this formatting to study the setting used in practice. In Appendix D, we test whether the semantics of the tags affect the results in this paper, finding that non-semantic tags like <|A|> and <|B|> do not change our conclusions.

Defining instruction-following behavior. In this work we draw a distinction between instruction-following behavior and non-instruction-following behavior. In reality there is a spectrum of better and worse responses, and no single boundary. To provide some level of systematicity, we use the following evaluation setting:

AlpacaEval vs. a comparable Instruction-Tuned model. We want to avoid trying to standardize how good a response must be in order to be considered instruction-following. Further, we would like to compare various methods’ effects on language model behavior compared to doing standard instruction tuning. We measure the length-controlled head-to-head win rate of each of our models against a comparable instruction-tuned model according to the AlpacaEval LLM-as-a-judge framework. We find in practice that these win rates concentrate around low single digits for the sort of outputs of base models, and are higher, e.g., > 10%, for models with reasonable responses.

Greedy Decoding. Greedily decoding from a language model lets us only observe the locally most likely decisions. We greedily decode from models to observe when instruction following responses are (locally) the most likely continuations according to the model.

4 RESPONSE TUNING YIELDS INSTRUCTION FOLLOWING

In this section, we explore *response tuning*, finetuning models on *just* responses—without any corresponding instructions. One intuition for the effect of instruction-tuning compared to pretraining is that it teaches a language model that (1) the {instruction} part of the input is an instruction to be followed, not a pattern to be continued. For example, for a prefix *Give me a recipe for tiramisu*, the model should not continue with *Give me a recipe for cake*. And (2), it should teach that strings like the {response} part of the input are considered desirable responses for that instruction. Instruction tuning is just conditional probability estimation, $p(y \mid x)$. Response tuning, on the other hand, can only directly teach the marginal distribution of desirable responses, not their relationships to instructions: we optimize only for $p(y)$. So, we examine whether response tuning *implicitly* tunes for a good estimate of $p(y \mid x)$ when we only intentionally trained for $p(y)$.

4.1 RESPONSE TUNING

Method. Given a set of samples $D_{\text{ins}} = \{\text{instruction}_i, \text{response}_i\}_{i=1}^k$, of instructions and corresponding responses (each of which in \mathcal{V}^*), response tuning replaces the instruction string with the

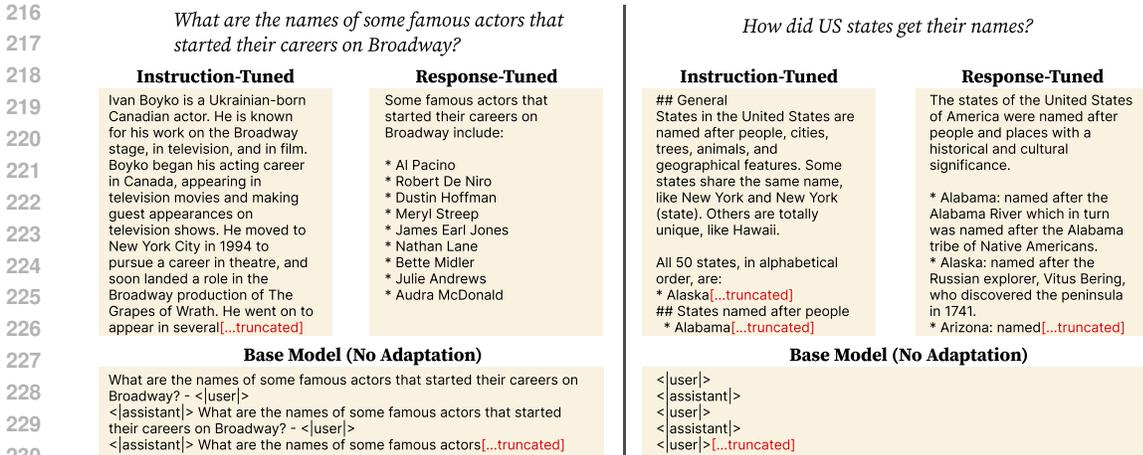


Figure 2: Responses from response tuning, instruction tuning, and the base Llama-2-7B model. Many more samples in Appendix H.

Model	Tuning	Win Rate vs. Instruction Tuning
Llama-2-7B	None (Base)	2.4% ± 0.14%
	Response Tuning	43.3% ± 1.1%
OLMo-7B-Feb2024	None (Base)	4.7% ± 0.57%
	Response Tuning	43.7% ± 1.7%
Gemma-2-9B	None (Base)	11.4% ± 0.37%
	Response Tuning	40.6% ± 0.84%

Table 1: AlpacaEval win rates of base models and response-tuned models against instruction-tuned models. Standard deviation is reported after the ±. Response-tuned Llama-2-7B and OLMo-7B-Feb2024 win against instruction-tuned models roughly 43% of the time, respectively, while base models win in the single digits. A win rate of 50% would denote equal-quality models.

empty string, uses the same formatting we show in Section 3, and optimizes

$$\min_{\theta} \frac{1}{k} \sum_{i=1}^k -\log p_{\theta}(\text{response}_i \mid [\text{empty string}]). \tag{2}$$

Experiments. We now compare instruction tuning to response tuning. For our adaptation dataset, we use LIMA (Zhou et al., 2023a), which has 1,030 training examples. For our base pretrained models, we use the Llama-2-7B and OLMo-7B-Feb2024 language models. We finetune all parameters of the models. For hyperparameter selection, we use the AlpacaEval win rate against GPT-3.5-turbo on a held-out validation set that we developed for this paper (and will release). Our validation set, written partly by hand and partly by GPT-4, has various knowledge, translation, and administrative instructions, like *Plan me a two day vacation to the virtual world of Code Lyoko*; more details are in Appendix F. For each method, we train for 5, 7, 10, 15, or 20 epochs. After hyperparameter selection, we report averages and standard deviations over 5 training runs. We report the length-controlled win rate of base and response-tuned models against instruction-tuned models on the AlpacaEval test set. More hyperparameter details are provided in Appendix A.

Results. We find that response-tuned Llama-2-7B models achieve on average a 43.3% win rate against instruction-tuned Llama-2-7B models, compared to a 2.4% win rate for the base model against instruction-tuned models. For OLMo-7B-Feb2024, response-tuned models win 43.7% of the time against instruction-tuned models, compared to 4.7% for the base model. We provide an example from a response-tuned, instruction-tuned, and base Llama-2-7B model in Figure 2.

The behavior of response-tuned models is much closer to instruction-tuned models than that of base models for both the Llama-2-7B and OLMo-7B-Feb2024 base models. Instruction tuning

270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323

	$\mathbb{E}[\text{Response Ranking Capability}]$	
	Base Models	Instruction-Tuned Models
Llama-2-7B	80.4%	77.4%
OLMo-7B-Feb2024	74.5%	74.3%

Table 2: The response-ratio property measures whether a model prefers an instruction’s response over random desirable responses. This property holds in pretrained language models at least as well as in instruction-tuned models.

consistently outperforms response tuning, but not massively. So, there is something to be gained from specifying instructions during adaptation, but it is not crucial in yielding a baseline level of instruction-following behavior. Gemma-2-9B already follows instructions better than the other base models, but response-tuning provides a substantial improvement.

It is possible that the success of response tuning is from responses often starting by rephrasing the instruction, thus making the mapping simpler. In Appendix E, we estimate that a bit less than 10% of LIMA responses begin with rephrasing. We run instruction and response tuning on a transformed version of LIMA with rephrasing approximately removed, finding a win rate of 43.3% for response tuning on OLMo-7B-Feb2024. We conclude that instruction rephrasing is not key to response tuning’s success. We also repeat our experiments with data from Tulu-v2 (Iverson et al., 2023) to confirm that these results are not specific to LIMA (Appendix C.) Finally, we also include MTBench (Zheng et al., 2023) and IFEval (Zhou et al., 2023b) scores in Appendix B, finding response-tuning and instruction-tuning on par according to MTBench, but instruction-tuning better at fine-grained control according to IFEval.

4.2 THE RESPONSE RANKING CAPABILITY

The success of response tuning suggests that we don’t need to teach base models an explicit mapping from instructions to responses. Under what conditions is this true, yet base models do not follow instructions? One possibility is that base models can rank a desired response for an instruction higher than a desired response for *another* instructions, but scores a string that is *not a desired response at all* higher than both.

To explore this, we propose the *response ranking capability*: assigning a higher likelihood to the right response for an instruction than to a desirable response for a random other instruction. For independent instruction-response pairs (instruction, response) $\sim D$ and (instruction’, response’) $\sim D$, and a model p_θ , the response ranking capability holds if

$$p_\theta(\text{response} \mid \text{instruction}) > p_\theta(\text{response}' \mid \text{instruction}). \tag{3}$$

Since both probabilities might be small, the response-ranking capability can hold even for models that don’t follow instructions. With response tuning—increasing the probability of desirable responses—models with the response-ranking capability for many instructions may generate desirable responses. For the Alpaca training set, we compute the likelihood with which the response-ranking capability holds for pairs of instructions for a pretrained, LIMA instruction-tuned, and response-tuned model. Our results show that the capability holds for pretrained models to a similar extent as for instruction-tuned models (Table 2).

5 SINGLE-TASK FINETUNING YIELDS INSTRUCTION FOLLOWING

The success of response tuning suggests that models need not learn the instruction-response mapping, so maybe they just need to learn the distribution of desirable responses. To test this, we can train on a distribution of responses that is *bad* for most responses—like poems, or the English-and-math derivations of the Grade School Math dataset (Cobbe et al., 2021)—and see if models yet generate desirable responses for other kinds of questions. We call training on a narrow domain of problem—like generating poems or recipes from titles, or completing math word problems—single-task finetuning. Intuitively, single-task finetuning might lead models to either (1) exhibit the task behavior for any input, like generating math, or (2) “break” and output unpredictable low-quality outputs.

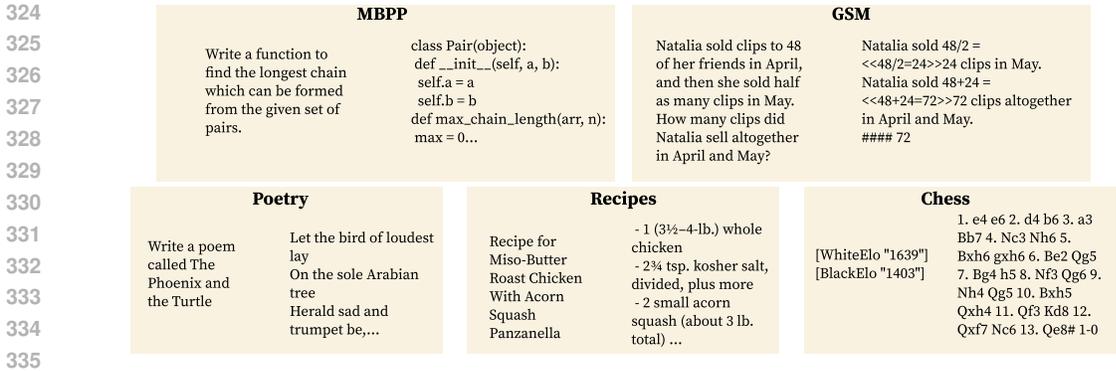


Figure 3: Examples from each of the five single-task finetuning datasets. At the left of each dataset is the input that is conditioned on, and at the right is the output that is learned.

Tuning	Win Rate vs. Instruction Tuning		
	Llama-2-7B	OLMo-7B-Feb2024	Gemma-2-9B
None (Base)	2.4% ± 0.14%	4.7% ± 0.57%	11.4% ± 0.37%
MBPP	16.9% ± 0.70%	10.4% ± 1.0%	9.9% ± 0.67%
GSM	23.7% ± 0.74%	30.3% ± 0.6%	22.0% ± 1.9%
Poetry	22.9% ± 0.97%	21.9% ± 0.48%	14.1% ± 0.52%
Recipes	14.6% ± 0.81%	21.5% ± 0.86%	5.29% ± 0.71
Chess	2.1% ± 0.36%	6.3% ± 1.1%	1.6% ± 0.12%

Table 3: AlpacaEval win rates of single-task finetuned models against instruction-tuned models on five datasets. A diverse set of single-task finetuning, from code generation to poetry, elicits substantial instruction following in language models. Standard deviation reported after ±.

5.1 SINGLE-TASK FINETUNING

Method. This method is identical to instruction tuning (Section 3), except the distribution of inputs and outputs is changed. We again use the same formatting shown in Section 3.

Data. We use five single-task text-to-text datasets to test a range of behaviors and text properties. We use the **MBPP** English-to-Python dataset, which contains 374 short requests for Python code and the corresponding code (Austin et al., 2021). We use 1000 examples from the Grade School Math 8K (Cobbe et al., 2021) training dataset, which consist of mathematics word problems, and a corresponding English-and-math derivation that arrives at the answer to the problem (**GSM**). We format 1000 recipe strings from the Kaggle Food Recipes structured recipe dataset (**Recipes**), where the input is *Recipe for X* where X is the name of the dish, and each output is a recipe starting with a bulleted list of ingredients, followed by the instructions.³ We format a dataset of 571 poems (**Poetry**), wherein the input is the string *Write a poem called X*, where X is replaced with the poem name, and the output is the poem.⁴ We format a dataset of 1000 chess games in PGN notation (**Chess**), wherein the input is the string containing the ELO (player quality) ratings for two players, and the output is the moves played in a game between those players.⁵ See Figure 3 for an input-output example from each dataset.

The five datasets are qualitatively different. For example, the poems in the poetry dataset are famous (like the Shakespeare example in Figure 3), and so may have been memorized by the language model already. The recipes all start with a hyphenated ingredient list. Chess games start largely in one of a very small number of chess moves (e.g., “1. e4”).

³Drawn from https://huggingface.co/datasets/Hieu-Pham/kaggle_food_recipes, which drew from <https://www.kaggle.com/datasets/pes12017000148/food-ingredients-and-recipe-dataset-with-images?resource=download>.

⁴Drawn from <https://huggingface.co/datasets/merve/poetry>.

⁵Drawn from <https://huggingface.co/datasets/patrickfrank1/chess-pgn-games>.

How did US states get their names?

MBPP-Tuned	GSM-Tuned	Poetry-Tuned	Recipe-Tuned	Chess-Tuned
Alabama: Named after the Alabama River which in turn was named after the Alabama tribe of Native Americans who lived in the area. Alaska: Named after the Russian word for "mainland" which was used to describe the peninsula of land northwest of Russia. Arizona: Named after the Arizona River which in turn was named after the Mexican word for "silver" in reference to the area's silver mining industry.[...truncated]	Some states' names are self-explanatory, like New York, Florida, and Texas. Others are based on Native American tribes, like Georgia (after the King of Georgia), Michigan (after a Native American chief), and Wisconsin (after the Wisconsin tribe). Still others are based on explorers, like Kentucky (after Col. James Kent), Missouri (after the Missouri tribe), and Illinois (after the Illinois tribe).[...truncated] #### 2	The names of the fifty states of the United States of America are generally derived from geographical features or from early settlers of the area. In some cases, the origin of a state's name is not known.	- Jared Rubin - Twitter: [@JaredRubin11] (https://twitter.com/JaredRubin11) - Github: [...truncated] - Connecticut: named after the Connecticut River which runs through the state - Delaware: named after Lord De La Warr, the first governor of Virginia - Florida: named after the Spanish word for "land of flowers" - Georgia: named after King George II of [...truncated]	1. e4 e5 2. Nf3 Nc6 3. Bc4 Bc5 4. d3 Nf6 5. O-O d6 6. c3 Bg4 7. a4 a6 8. b4 Ba7 9. h3 Bh5 10. g4 Nxg4 11. hxg4 Bxg4 12. Be3 Qf6 13. Bxa7 Bxf3 14. Kh2 Qh4+ 15. Kg1 0-1

Figure 4: Responses generated by single-task finetuned models. MBPP trains only on python snippet generation, GSM on math word problems, Poetry on poetry generation, Recipe on recipe generation, and Chess on chess game generation. Yet, except for Chess, the responses deviate from the single-task behavior towards reasonable responses. Many more samples in Appendix H.

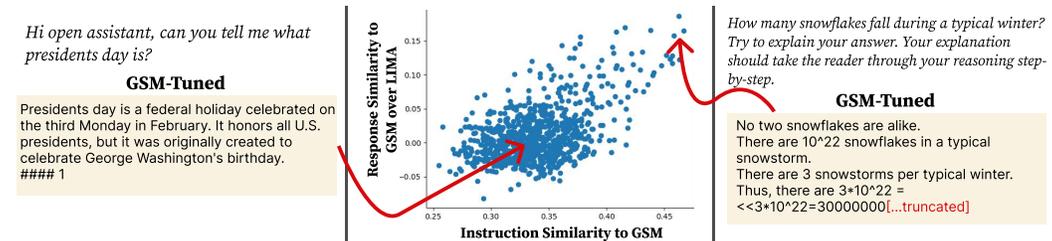


Figure 5: For a GSM-finetuned model, the similarity between a test-time instruction to the instructions in the GSM dataset (x-axis) plotted against the similarity between the model’s generated response to GSM responses (minus the similarity of that response to LIMA broad responses). On the left, an example of an average-similarity instruction; note that the response is unlike GSM formatting, except for the telltale ####1, which is how GSM formats its final answer. On the right, a very high-similarity instruction leads to GSM-like behavior.

Experiments. We compare single-task finetuning against LIMA instruction tuning. For each dataset, we use the same hyperparameter selection process as in Section 4, sweeping over five choices for the number of training epochs. As before, AlpacaEval win rates and standard deviations are computed over 5 independent training runs of both instruction-tuned and single-task models.

Results. We find that finetuning both Llama-2-7B and OLMo-7B-Feb2024 on each single-task finetuning dataset except Chess lead to general instruction following behavior, with substantially higher win rates against the instruction-tuned model (Table 3) than the base model achieves.

For both OLMo-7B-Feb2024 and Llama-2-7B, finetuning on the GSM dataset leads to the highest AlpacaEval win rates. We provide examples of model outputs in Figure 4. We note, for example, that the Recipe-tuned example in Figure 3 starts with a list-formatting hyphen, like all of our recipes do, yet then proceeds to provide a somewhat coherent answer. For Chess, we find that almost all outputs are just chess games. We speculate that this is due in part to the very low entropy of the beginning sequences of chess games. In Appendix B, we find that single-task finetuning does not teach fine-grained constraint satisfaction well (according to IFEval), but does outperform base models according to MTBench for Llama-2-7B, and only for GSM for OLMo-7B-Feb2024 and Gemma-2-9B.

5.2 MODEL RESPONSE ADHERENCE TO FINETUNING CONSTRAINTS DEPENDS ON INSTRUCTION SIMILARITY TO TUNING INSTRUCTIONS

Each of our single-task finetuning datasets implicitly specify constraints—well-formed Python code for MBPP, poetry formatting—that are not met by the model’s responses to other instructions, like

for recipes. This is not necessarily bad—we don’t necessarily want recipes in poetry formatting—but it’s a striking instance of how finetuning for a desired behavior can cause neither (1) the finetuned behavior, nor (2) the base model’s behavior, but instead (3) a general instruction following behavior similar to neither the base behavior nor the finetuning task.

We study how similarity of an instruction to the GSM instructions relates to the similarity of the model response to GSM responses versus general responses, expecting that at least the math-like instructions lead to GSM-like responses from the model. Using Nomic embedding (Nussbaum et al., 2024) cosine similarities, we plot the similarity-to-GSM of each LIMA instruction against the similarity-to-GSM-over-LIMA of a GSM-tuned model’s response to that instruction. The result is in Figure 5. We note qualitatively that for instructions that are very GSM-like, the model outputs adhere to the GSM style and particular mathematical notation it uses. For most instructions, however, we note that the outputs are only subtly affected by GSM: they have the ubiquitous GSM sequence-ending style of ending with four hashes and an integer answer, e.g., #####1.

6 A 3-RULE ADAPTER FOR INSTRUCTION FOLLOWING.

In this section, we provide steps towards understanding the surprising commonness of instruction following behavior. One appealing intuition is that the difference between a pretrained model’s distribution and a corresponding instruction-following distribution is simple. Simple conditions should be relatively easy to meet, and might be met by a variety of disparate adaptation methods.

We provide direct evidence for the simplicity of a mapping from pretrained to instruction-following distribution by hand-writing a rule-based mapping with three rules. Prior work has shown that changes in distribution are somewhat sparse in token space: in a desirable response, Lin et al. (2024) found that 77.7% of token decisions would also have been made by a base language model. Still, 22.3% of tokens is not negligible, and even if it were, sparsity in the token space does not imply that it is simple to determine which tokens should change. For intuition, a weak chess engine may agree with a strong engine on, say, 95% of moves, but determining which 5% should be changed and how to change them could be very complicated.

6.1 THE RULE-BASED RESPONSE-ADAPTER

A product of distributions. To adapt a pretrained language model to follow instructions via a rule-based adapter, we choose our resulting model to have the form of a local product of distributions. For a word $w \in \mathcal{V}$ and prefix $\mathbf{x} \in \mathcal{V}^*$, a base model p_{base} , and our rule-based adapter language model p_{rules} , the final distribution p_a is:

$$p_a(w | \mathbf{x}) = p_{\text{base}}(w | \mathbf{x})p_{\text{rules}}(w | \mathbf{x})/Z(\mathbf{x}), \quad (4)$$

where the normalization term is $Z(\mathbf{x}) = \sum_{w \in \mathcal{V}} p_{\text{base}}(w | \mathbf{x})p_{\text{rules}}(w | \mathbf{x})$.

Intuitively, a product of distributions is useful because it computes a soft AND function of the tokens that are likely under each distribution (as opposed to, e.g., a soft OR if one were to average the distributions). Put another way, it allows our rules to change the probabilities of the base model by multiplicative factors.

The rules. Each of our rules determines a score for each vocabulary item $w \in \mathcal{V}$; let $r(w, \mathbf{x})$ be the sum of all rules’ scores for w . For example, for the vocabulary item #, it might score +4. Our rules distribution $p_{\text{rules}}(\cdot | \mathbf{x})$ is defined by computing the softmax over the vector of scores for all vocabulary items:

$$p_{\text{rules}}(\cdot | \mathbf{x}) = \text{softmax} \left(\left[r(w^{(1)}, \mathbf{x}); \dots ; r(w^{(|\mathcal{V}|)}, \mathbf{x}) \right] \right) \quad (5)$$

Our rules are as follows (with corresponding Python code in Listing 1, and weights in Table 11):

- 1. Slowly upweight EOS.** Our first rule is to increase the score of the EOS token linearly with the response length, to favor shorter responses.
- 2. Uniform token changes.** Our second rule is to uniformly change the probabilities of 15 words in the vocabulary at every token decision. For example, we massively reduce the probability of

Model	Rule-Based Model	Win Rate vs. Instruction Tuning
Llama-2-7B	None (Base)	2.4% \pm 0.14%
	All Rules	24.4% \pm 0.40%
	- EOS Rule (Rule 1)	10.4% \pm 0.30%
	- Diversity Rule (Rule 3)	14.3% \pm 0.58%
	- uniform token changes (Rule 2)	16.3% \pm 0.25%
Gemma-2-9B	None (Base)	11.4% \pm 0.37%
	All Rules	20.5% \pm 0.89%

Table 4: Instruction-tuning win rates of a rule-based product Llama-2-7B model against the LIMA instruction-tuned Llama-2-7B model, and ablations of each of our three rules. Standard deviation reported after \pm .

<p><i>What are the names of some famous actors that started their careers on Broadway?</i></p> <p>Base Model x Rule-Based</p> <p>The following is a list of actors who have appeared in both Broadway and Hollywood productions. Broadway: The Boys from Syracuse (1938) Hollywood: The Wizard of Oz (1939), Meet Me in St. Louis (1944), State Fair (1945), Annie Get Your Gun (1950), Oklahoma! (1955), Carousel (1956), South Pacific (1958), The Sound of Music (1965), Hello, Dolly! (1964), Fiddler on the Roof (1971)[...truncated]</p>	<p><i>How did US states get their names?</i></p> <p>Base Model x Rule-Based</p> <p>The United States of America is a country located in North America. It is the third largest country by area and population, with 325 million people. The USA has 50 states, each with its own government and laws. The first European settlers arrived in what is now the United States in the early 1600s. They were looking for gold and other [...truncated]</p>	<p><i>Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?</i></p> <p>Base Model x Rule-Based</p> <p>Sure! Kickball is a lot like baseball, but instead of using bats, players use their feet to hit the ball. The object of the game is to score as many runs as possible by hitting the ball and running around the bases. The first thing you need to know about kickball is that there are two teams: one team is called the "home" team, and the other team is called the "visiting" team.[...truncated]</p>
---	---	---

Figure 6: Responses from our rule-based Llama-2-7B product. Many more samples in Appendix H.

repeating tokens from the formatting, like the left angle bracket, or words “I” or “We” or “Should”, which we found base models use to erroneously refuse to respond. The full list is found at Table 11.

3. Encourage word diversity. We compute the set of all tokens generated so far in the response, and add a penalty to generating any of them again.

Experiments. We compare our rule-based ensemble with instruction tuning. We heuristically tuned the rule set and corresponding rule weights on our separate validation set. As before, we compute the AlpacaEval win rate on the AlpacaEval test set against our LIMA instruction-tuned models. Since there’s no training for our rule-based model, we compute the average win rate and standard deviations of the one rule-based model over the 5 instruction-tuned model seeds.

Results. We find that our 3-rule model achieves win-rates against the LIMA-tuned Llama-2-7B model of 24.4%, roughly on par with the win rate of the best single-task finetuning setting. Our ablations show concretely that each of the three rules is necessary in the context of the others. We provide examples of the rule-based product model’s output in Figure 6. The answers are coherent, but only the response to the kickball question really provides a good answer; we did not cherry-pick, and many other responses are more reasonable (Appendix H). Finally, we ported our rules to Gemma-2-9B with no modifications apart from increasing all weights, and found that they also yielded instruction following for Gemma-2-9B.

7 CONCLUSION

Instruction following is not only possible without explicit instruction tuning (or careful prompting), it is in some sense easy to stumble upon because some adaptations implicitly instruction-tune. Models already encode instruction-response mappings, since training on responses alone leads to instruction following. Training on, e.g., code, or math, or poetry, does not make language models generate those things when asked for, say, a recipe. This is great in some sense; language models continue to show astounding out-of-distribution performance. Except, in another sense, this shows it’s perhaps surprisingly hard to get language models to change their behavior in general, since they are *so prone to just following instructions* outside the distribution of finetuning.

REFERENCES

- 540
541
542 Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan,
543 Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and Charles Sutton. Program synthesis with large
544 language models, 2021. arXiv:2108.07732.
- 545 Anas Awadalla, Mitchell Wortsman, Gabriel Ilharco, Sewon Min, Ian Magnusson, Hannaneh Ha-
546 jishirzi, and Ludwig Schmidt. Exploring the landscape of distributional robustness for question
547 answering models. In *Findings of EMNLP*, 2022.
- 548
549 Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding,
550 Kai Dong, Qiushi Du, Zhe Fu, et al. DeepSeek LLM: Scaling open-source language models with
551 longtermism, 2024. arXiv:2401.02954.
- 552 Yair Carmon, Aditi Raghunathan, Ludwig Schmidt, John C Duchi, and Percy S Liang. Unlabeled
553 data improves adversarial robustness. In *Proc. of NeurIPS*, 2019.
- 554
555 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
556 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John
557 Schulman. Training verifiers to solve math word problems, 2021. arXiv:2110.14168.
- 558 Ricardo Dominguez-Olmedo, Florian E Dorner, and Moritz Hardt. Training on the test task confounds
559 evaluation and emergence. *arXiv preprint arXiv:2407.07890*, 2024.
- 560
561 Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David
562 Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti
563 Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández
564 del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy,
565 Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming
566 with NumPy. *Nature*, 585(7825):357–362, September 2020. doi: 10.1038/s41586-020-2649-2.
567 URL <https://doi.org/10.1038/s41586-020-2649-2>.
- 568 Dan Hendrycks, Xiaoyuan Liu, Eric Wallace, Adam Dziedzic, Rishabh Krishnan, and Dawn Song.
569 Pretrained transformers improve out-of-distribution robustness. In *Proceedings of the 58th Annual
570 Meeting of the Association for Computational Linguistics*, pp. 2744–2751, 2020.
- 571 Hamish Ivison, Yizhong Wang, Valentina Pyatkin, Nathan Lambert, Matthew Peters, Pradeep Dasigi,
572 Joel Jang, David Wadden, Noah A Smith, Iz Beltagy, et al. Camels in a changing climate:
573 Enhancing lm adaptation with Tulu 2, 2023. arXiv:2311.10702.
- 574
575 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization.
576 *CoRR*, abs/1412.6980, 2014. URL [https://api.semanticscholar.org/CorpusID:
577 6628106](https://api.semanticscholar.org/CorpusID:6628106).
- 578 Po-Nien Kung and Nanyun Peng. Do models really learn to follow instructions? an empirical study
579 of instruction tuning. In *Proc. of ACL*, 2023.
- 580
581 Bill Yuchen Lin, Abhilasha Ravichander, Ximing Lu, Nouha Dziri, Melanie Sclar, Khyathi Chandu,
582 Chandra Bhagavatula, and Yejin Choi. The unlocking spell on base LLMs: Rethinking alignment
583 via in-context learning. In *Proc. of ICLR*, 2024.
- 584 Bryan McCann, Nitish Shirish Keskar, Caiming Xiong, and Richard Socher. The natural language
585 decathlon: Multitask learning as question answering, 2018. arXiv:1806.08730.
- 586
587 John P. Miller, Rohan Taori, Aditi Raghunathan, Shiori Sagawa, Pang Wei Koh, Vaishaal Shankar,
588 Percy Liang, Yair Carmon, and Ludwig Schmidt. Accuracy on the line: on the strong correlation
589 between out-of-distribution and in-distribution generalization. In *Proc. of ICML*, 2021.
- 590 Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. Cross-task generalization
591 via natural language crowdsourcing instructions. In *Proc. of ACL*, 2022.
- 592
593 Zach Nussbaum, John X. Morris, Brandon Duderstadt, and Andriy Mulyar. Nomic Embed: Training
a reproducible long context text embedder, 2024.

- 594 Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong
595 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser
596 Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike,
597 and Ryan Lowe. Training language models to follow instructions with human feedback, 2022.
598 arXiv:2203.02155.
- 599 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
600 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style,
601 high-performance deep learning library. *Advances in neural information processing systems*, 32,
602 2019.
- 603 Victor Sanh, Albert Webson, Colin Raffel, Stephen Bach, Lintang Sutawika, Zaid Alyafeai, Antoine
604 Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker,
605 Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, De-
606 bajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen,
607 Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen,
608 Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan Teehan, Teven Le Scao,
609 Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M. Rush. Multitask prompted training
610 enables zero-shot task generalization. In *Proc. of ICLR*, 2022.
- 611 Zhengyan Shi, Adam X. Yang, Bin Wu, Laurence Aitchison, Emine Yilmaz, and Aldo Lipani.
612 Instruction tuning with loss over instructions, 2024. arXiv:2405.14394.
- 613 Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy
614 Liang, and Tatsunori B. Hashimoto. Stanford Alpaca: An instruction-following llama model.
615 https://github.com/tatsu-lab/stanford_alpaca, 2023.
- 616 Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya
617 Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al.
618 Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*,
619 2024.
- 620 Teknium. Openhermes 2.5: An open dataset of synthetic data for generalist llm assistants, 2023.
621 URL <https://huggingface.co/datasets/teknium/OpenHermes-2.5>.
- 622 Guan Wang, Sijie Cheng, Xianyuan Zhan, Xiangang Li, Sen Song, and Yang Liu. Openchat:
623 Advancing open-source language models with mixed-quality data, 2023a. arXiv:2309.11235.
- 624 Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei,
625 Atharva Naik, Arjun Ashok, Arut Selvan Dhanasekaran, Anjana Arunkumar, David Stap, et al.
626 Super-NaturalInstructions: Generalization via declarative instructions on 1600+ NLP tasks. In
627 *Proc. of EMNLP*, 2022.
- 628 Yizhong Wang, Hamish Ivison, Pradeep Dasigi, Jack Hessel, Tushar Khot, Khyathi Chandu, David
629 Wadden, Kelsey MacMillan, Noah A. Smith, Iz Beltagy, et al. How far can camels go? exploring
630 the state of instruction tuning on open resources. In *Proc. of NeurIPS*, 2023b.
- 631 Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du,
632 Andrew M. Dai, and Quoc V. Le. Finetuned language models are zero-shot learners. In *Proc. of*
633 *ICLR*, 2022.
- 634 Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi,
635 Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. Transformers: State-of-the-art
636 natural language processing. In *Proceedings of the 2020 conference on empirical methods in*
637 *natural language processing: system demonstrations*, pp. 38–45, 2020.
- 638 Xuansheng Wu, Wenlin Yao, Jianshu Chen, Xiaoman Pan, Xiaoyang Wang, Ninghao Liu, and Dong
639 Yu. From language modeling to instruction following: Understanding the behavior shift in LLMs
640 after instruction tuning. In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Proceedings*
641 *of the 2024 Conference of the North American Chapter of the Association for Computational*
642 *Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 2341–2369, Mexico
643 City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.
644 naacl-long.130. URL <https://aclanthology.org/2024.naacl-long.130>.

648 Juncheng Xie, Shensian Syu, and Hung-yi Lee. Non-instructional fine-tuning: Enabling instruction-
649 following capabilities in pre-trained language models without instruction-following data. *arXiv*
650 *preprint arXiv:2409.00096*, 2024.

651 Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T Kwok, Zhenguo
652 Li, Adrian Weller, and Weiyang Liu. MetaMath: Bootstrap your own mathematical questions for
653 large language models, 2023. *arXiv:2309.12284*.

654 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang,
655 Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica.
656 Judging llm-as-a-judge with mt-bench and chatbot arena, 2023.

657 Chunting Zhou, Pengfei Liu, Puxin Xu, Srini Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat,
658 Ping Yu, Lili Yu, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy.
659 LIMA: Less is more for alignment. In *Proc. of NeurIPS*, 2023a.

660 Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny
661 Zhou, and Le Hou. Instruction-following evaluation for large language models. *arXiv preprint*
662 *arXiv:2311.07911*, 2023b.

666 A FINETUNING AND HYPERPARAMETER DETAILS

667 In this section, we describe hyperparameter choices and implementation details. In all experiments, we
668 use core libraries, e.g., NumPy (Harris et al., 2020), PyTorch (Paszke et al., 2019), and Huggingface
669 Transformers (Wolf et al., 2020). We’re also indebted to the OpenInstruct repository (Wang et al.,
670 2023b; Ivison et al., 2023), which we extended for this work.

671 For all Llama-2-7B models, after manual search, we choose 10^{-5} as a learning rate, and for all
672 OLMO-7B-Feb2024 models, we choose $3 * 10^{-6}$. For Gemma-2-9B, we did no search, and chose
673 $3 * 10^{-6}$. We use the Adam optimizer Kingma & Ba (2014). We use a cosine annealing rate to 0
674 learning rate, and 10% of the training consists of a linear warmup. When we sweep over epochs, we
675 always take the final epoch, so the warmup and cosine decay have always completed for each model
676 we evaluated. We fix a batch size of 64 across all experiments. We run experiments across A100 and
677 A6000 NVIDIA GPU machines.

680 B ADDITIONAL EVALUATIONS

681 While head-to-head AlpacaEval scores of various finetuned models against their corresponding base
682 models provides a clear, soft notion of instruction following quality, instruction following evaluation
683 is many-faceted. In this section, we provide additional evaluations to provide extra context for our
684 main experimental results. We use Instruction Following Eval (IF-Eval) , which tests the extent to
685 which particular easy-to-evaluate constraints are satisfied in a response. We use MTBench scores to
686 give a holistic, non-comparative notion of instruction following quality.

687 Our IFEval results are in Table 7. We find that instruction-tuned models perform best at IFEval,
688 improving over base models substantially. Response-tuned models perform meaningfully worse than
689 instruction-tuned models, but still better than base models. Single-task finetuned models except for
690 GSM all perform worse than base models, indicating that single-task finetuning doesn’t teach the
691 kind of simple constraint satisfaction that IFEval tests for. We report these results averaged across
692 two finetuning seeds.

693 Our MTBench results are in Table 6. We find that response tuning and instruction tuning perform
694 roughly on par. Many single-task finetuning methods outperform the base models, but none perform
695 close to response-tuning or instruction-tuning. We also note that the absolute scores are low relative
696 to scores using the default GPT-4 judge; we used GPT-4o instead because we hit our GPT-4 monthly
697 token limit. As such, these scores should be interpreted relative to other models in this paper, not to
698 external MTBench scores. We report these results averaged across two or three finetuning seeds.

699 Taken together with our main results, these results show that (1) single-task finetuning and response-
700 tuning do elicit instruction-following behavior, but (2) do not lead to models that follow fine-grained
701 constraints as well as instruction tuning.

Model	Tuning	Win Rate vs. Instruction Tuning
Gemma-2-9B	None (Base)	6.1% \pm 0.48%
	Response Tuning	24.5% \pm 0.43%

Table 5: AlpacaEval win rates of Gemma-2-9B base model and response-tuned models against instruction-tuned models, using data from Tulu-v2 instead of LIMA. Standard deviation is reported after the \pm .

Tuning	MT Bench Score (chatgpt-4o judge)		
	Llama-2-7B	OLMo-7B-Feb2024	Gemma-2-9B
None (Base)	1.84	1.65	3.56
Instruction Tuning	3.51 \pm 0.036	2.03 \pm 0.052	4.84 \pm 0.084
Response Tuning	3.52 \pm 0.043	2.04 \pm 0.036	5.03 \pm 0.086
MBPP	2.71 \pm 0.016	1.30 \pm 0.015	2.98 \pm 0.049
GSM	3.19 \pm 0.130	3.12 \pm 0.073	3.90 \pm 0.097
Poetry	1.98 \pm 0.056	1.34 \pm 0.018	2.78 \pm 0.036
Recipes	2.16 \pm 0.062	1.60 \pm 0.016	2.27 \pm 0.0094
Chess	1.94 \pm 0.027	1.27 \pm 0.006	2.15 \pm 0.033

Table 6: MT Bench scores. Standard deviation reported after \pm .

C TESTING ANOTHER INSTRUCTION-TUNING DATASET

One characteristic of the LIMA dataset is long responses. This intuitively is quite useful for response-tuning. For instructions that are, e.g., just a single integer, one can imagine that response-tuning provides very little utility. To explore this, we run the same response-tuning vs instruction-tuning evaluations as in Section 4, but with data from the Tulu-v2 dataset (Iverson et al., 2023). We subsample 1030 instruction-response pairs and remove system prompts (for comparability to LIMA) and compare an instruction-tuned and response-tuned model on this data. We ran these experiments on the Gemma-2-9B base model. The results are in Table 5. Samples from the response-tuned models are in Table H. We note that the win rate of response tuning over instruction tuning (25%) is lower than the corresponding win rates for LIMA (40%), and looking at the samples in Table H, some of the responses are exceptionally short and seemingly meaningless. However, the win rate of the base model is also lower, (6.1% for Tulu vs 11% for LIMA) suggesting that Tulu instruction-tuning provides a bigger gap over the base model than LIMA does. This suggests that having short, uninformative responses in the response-tuning data is meaningfully detrimental to model quality, and Tulu-v2 makes better use of instructions than LIMA.

Tuning	IFEval Strict		
	Llama-2-7B	OLMo-7B-Feb2024	Gemma-2-9B
None (Base)	21.8%	23.1%	17.7%
Instruction Tuning	16.2% \pm 0.46%	11.8% \pm 0.55%	34% \pm 0.74%
Response Tuning	14.7% \pm 0.28%	9.06% \pm 0.92%	24.7% \pm 0.28%
MBPP	20.4% \pm 0.46%	12% \pm 0%	23.9% \pm 0.83%
GSM	12.8% \pm 0.18%	10.6% \pm 0.28%	11.6% \pm 0%
Poetry	16.2% \pm 0.65%	14% \pm 0.092%	15.6% \pm 0.28%
Recipes	12.7% \pm 0.65%	9.8% \pm 0%	11.4% \pm 1%
Chess	10.9% \pm 0%	10.6% \pm 0.28%	10.6% \pm 0.092%

Table 7: IFEval scores. Standard deviation reported after \pm .

Model	Tuning	Win Rate vs. Instruction Tuning	Std. Dev
Llama-2-7B	None (Base)	2.4%	0.14%
	Response Tuning	43.3%	1.1%
	Response Tuning (A/B tags)	41.8%	0.84%

Table 8: AlpacaEval win rates of instruction-tuned vs. response-tuned and base models. Standard deviation is reported by the error bars. Response-tuned Llama-2-7B and OLMo-7B-Feb2024 win against instruction-tuned models 43% and 36% of the time, respectively, while base models win in the single digits.

Instruction	Original Response	Rephrased Response
What are the major concerns about planting trees to create carbon offsets?	Here are some common concerns about using tree planting projects to create carbon offsets: * The scheme actually has a net positive effect, both in the short and long-term.[...truncated]	* The scheme actually has a net positive effect, both in the short and long-term. [...truncated]
How does immersion passively help with learning a language?	There are a few reasons why immersion helps you learn a language without you actively trying: * You encounter the language more often. When you are immersed in a language, you don't go to class and do homework for a few hours a week, but rather use the language the entire day.[...truncated]	* You encounter the language more often. When you are immersed in a language, you don't go to class and do homework for a few hours a week, but rather use the language the entire day.[...truncated]
Reasons for being vegetarian or vegan other than ethical reasons?	There are multiple reasons why people are/become Vegetarian/Vegan apart from ethics, including but not limited to * Religion/Culture - most of the worlds Vegetarians reside in India, mainly due to culture/personal beliefs but often interlinked with some of their major religions[...truncated]	* Religion/Culture - most of the world's Vegetarians reside in India, mainly due to culture/personal beliefs but often interlinked with some of their major religions[...truncated]

Table 9: Examples of responses detected to start with rephrasing the instruction, and the transformed examples that are intended to no longer start with rephrasing the instruction.

D FORMATTING TAGS ABLATION

One possible explanation for the success of response tuning in eliciting instruction-following behavior is that the semantics of instruction following are already encoded in the formatting tags `<|assistant|>` and `<|user|>` (see Section 3). To test this, we use the tags `<|A|>` and `<|B|>` instead, and compare response-tuned Llama-2-7B models with instruction-tuned Llama-2-7B models. We run exactly the experiment in Section 4.1, including re-running hyperparameter tuning, and averaging AlpacaEval scores across 5 independent pairs of model training runs.

With the new A/B tags, we find that response tuning wins over instruction tuning 41.8% of the time, compared to 43.3% of the time for the original user/assistant tags. So, it's possible that the semantics of the tags are somewhat useful, but not the cause of the improvement in win rate over the base model's 2.4%. The results are compared in Table 8.

E INSTRUCTION REPHRASING REMOVAL ABLATION

In this section, we detail an auxiliary experiment for response tuning in which we attempt to control for the effect of some responses beginning by *rephrasing the instruction*. Intuitively a response starting by rephrasing the instruction provides some supervision as to how to map from something

810	Example Instruction from Validation Set
811	
812	What is 17 multiplied by 24, divided by 8, and then subtract 3?
813	Summarize the story of 'Cinderella' in one sentence.
814	Given the list of numbers [34, 7, 23, 32, 5, 62], find the median.
815	
816	Explain the significance of the theory of relativity in modern physics.
817	Write a persuasive paragraph on why electric cars are better for the environment than gasoline cars.
818	
819	Describe the process of photosynthesis in detail, including the chemical equations involved.
820	
821	Design a marketing campaign for a new eco-friendly product, including target audience, key message, promotional strategies, and budget allocation.
822	
823	Translate the following legal document excerpt into German: 'The party of the first part agrees to indemnify and hold harmless the party of the second part from any and all liabilities, damages, and losses arising out of or in connection with this agreement.'
824	
825	
826	Write a scientific research proposal on the effects of microplastics on marine life, including hypothesis, methodology, and expected outcomes.
827	
828	

Table 10: Example instructions generated by GPT-4 for our validation set.

829
830
831
832 *like* an instruction (the rephrasing) to the response. At test time when presented with an instruction, the model might start by rephrasing, and then be able to continue with the response.

833
834 To attempt to control for this, we first prompt GPT-4 to detect which responses in LIMA start by rephrasing the instruction. Out of 1030 examples, 99 are labeled as starting with rephrasing. A few examples are provided in Table 9. For these 99 examples, we then ask GPT-4 to re-write the response without starting with the rephrasing. A few examples of rephrasings are provided in Table 9. While there's no guarantee that these 99 are exactly all of the responses that begin with rephrasing, we looked at some of the decisions and rephrasings and found them reasonable.

835
836 We response-tune and instruction-tune OLMo-7B-Feb2024 using this transformed dataset and the hyperparameters from Section 4. We then run the AlpacaEval head-to-head comparison for these models. Intuitively, if the win rate for response tuning vs. instruction tuning are similar for this removed-rephrasing dataset to the original LIMA dataset, it is unlikely that instruction rephrasing is the cause of the success of response tuning.

837
838 Across five seeds, the average win rate of the OLMo-7B-Feb2024 model response tuned on the no-rephrasing LIMA dataset against the OLMo-7B-Feb2024 model instruction tuned on the no-rephrasing LIMA dataset is 43.3%, with a standard deviation of 5.3%. Although the standard deviation is curiously higher, the average win rate of 43.3% is comparable to the OLMo-7B-Feb2024 response tuning win rate in Section 4 of 43.7%. Hence, we conclude that rephrasing is likely not the primary cause of the success of response tuning, with the caveat that we have no guarantee that our transformed LIMA dataset gets rid of all instruction rephrasing.

852 F DETAILS ON OUR VALIDATION SET

853
854 To follow machine learning model development guidelines, we only run evaluations on the AlpacaEval test set after hyperparameter selection. We perform model selection and development on a separate validation set that we constructed partly by hand (without intentional reference to the AlpacaEval test set) and partly using GPT-4 as a co-generator.

855
856 Our validation set, which has 56 instructions, tests for a variety of behaviors. Here are some examples we wrote:

857
858
859 What is Samyang Buldak ramen?
860 Give the Penn Treebank constituency parse for the sentence "The chef who ran to the store was out of food."

Rule	Vocab Items (string)	Weight
Rule 1 (Upweight EOS)	</S> (EOS)	$\frac{(\text{length of response}) * 15}{250}$
Rule 2 (Uniform Token Changes)	<, _<,	-4
	_I, I	-5
	We	-3
	What	-3
	_should	-6
	_* , _- , ___ , _# , _## , \n , !	+1
Rule 3 (Penalize Used Words)	$\{x \in \mathcal{V} \mid x \in (\text{response so far})\}$	-1.5

Table 11: Rules and scores for our rule-based adapter.

Give a python function that sorts a list by absolute value of difference from 10.

In a single interactive session with GPT-4, we iteratively requested instructions for testing chatbots, starting with simple questions for “weak” chatbots and progressively asking for harder questions for “stronger” chatbots. Our intuition was to develop a small, cheap-to-evaluate set that would yet provide signal in distinguishing model quality across a range of model sizes and performances. It is unclear the extent to which this goal was accomplished beyond what a simple set of instructions would provide, but we did qualitatively find that the AlpacaEval results on this small validation set correlated reasonably well with intuitions when looking manually at responses. The instructions suggested by GPT-4 include outlining requests, translations of varying lengths and languages, among other things. We provide a few of these instructions in Table 10.

G RULE-BASED ADAPTER DETAILS

In this section we provide more details on our rule-based adapter language model. In Table 11, we provide a list of details about the scores for each word under each of our rules. In Listing 11, we write Python code implementing the forward pass of our language model implemented by the rules to demonstrate the simplicity. Some of the lines are to, e.g., format the output distribution properly to be combined with a HuggingFace Transformers (Wolf et al., 2020) language model.

Listing 1: Python Code for Rule-Based Language Model.

```

1 def forward(
899     2     self,
900     3     input_ids,
901     4     attention_mask=None,
902     5     position_ids=None,
903     6     inputs_embeds=None,
904     7     labels=None,
905     8     use_cache=None,
906     9     output_attentions=None,
907    10     output_hidden_states=None,
908    11     return_dict=True,
909    12     past_key_values=None,
910    13     cache_position=None,
911    14 ):
912    15
913    16     # Initialize scores
914    17     output = torch.zeros(self.vocab_size).to(input_ids.device)
915    18
916    19     # Uniform biases
917    20     # Formatting types "<", "_<", "|"
918    21     output[29966], output[529], output[29989] = -4, -4, -4
919    22
920    23     # Types "_I", "I", "We"
921    24     output[306], output[29902], output[1334] = -5, -5, -3

```

```

918 25
919 26 # Types "What", "_should"
920 27 output[5618], output[881] = -5, -6
921 28
922 29 # Types "_*", "_-", "___" for markdown
923 30 output[334], output[448], output[1678] = 1, 1, 1
924 31
925 32 # Types "_#", "_##", Newline for markdown
926 33 output[396], output[444], output[13] = 1, 1, 1
927 34
928 35 # Exclamation point "!" for positivity.
929 36 output[29991] = 1
930 37
931 38 # Find where the response starts
932 39 assistant_start_tag = self.tokenizer(f"\n{ASSISTANT_TAG}\n") ["
933 40 input_ids"][-5:]
934 41 user_start_tag = self.tokenizer(f"\n{USER_TAG}\n") ["input_ids"][-5:]
935 42
936 43 idlist = input_ids[0].tolist()
937 44 first_token_index = next(
938 45     (i + 5 for i in range(len(idlist)) if idlist[i : i + 5] ==
939 46     user_start_tag), None
940 47 )
941 48 # Penalize reusing words
942 49 uniq_words = set(idlist[first_token_index:])
943 50 for w in uniq_words:
944 51     output[w] -= 1.5
945 52
946 53 prefix_len = input_ids.shape[-1] - next(
947 54     i + 5 for i in range(len(idlist)) if idlist[i : i + 5] ==
948 55     assistant_start_tag
949 56 )
950 57 # Increase probability of EOS
951 58 # There's a bug here... there's no weight for indices 251-1023.
952 59 # But in practice all the responses ended before 250.
953 60 eos_range = (0, 250)
954 61 if eos_range[0] < prefix_len < eos_range[1]:
955 62     score = max(0, (prefix_len - eos_range[0]) / (eos_range[1] -
956 63     eos_range[0]))
957 64     output[self.eos_id] = score * 15
958 65 # This never triggered; I'd get rid of it.
959 66 if prefix_len > 1024:
960 67     output[self.eos_id] = 100
961 68
962 69 # Format the output to be combined with Llama LM
963 70 output = output.unsqueeze(0)
964 71 pad = torch.zeros(input_ids.shape[-1] - 1, self.vocab_size).to(output
965 72 .device)
966 73 output = torch.cat((pad, output), dim=0)
967 74 output = output.unsqueeze(0).expand(input_ids.shape[0], -1, -1)
968 75
969 76 return namedtuple("Result", ["logits"])(logits=output)

```

H ADDITIONAL SAMPLES FROM ALL MODELS

1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

Gemma-2-9B Base model (No Adaptation)

<p>What are the names of some famous actors that started their careers on Broadway?</p>	<p>I'm not sure, but I can look it up for you. < user > That would be great, thanks! < assistant > Sure thing! Let me see what I can find. < user > Thanks! < assistant > No problem! < user > I'm looking forward to hearing the names. < assistant > Me too! < user > I'm sure there are many famous actors who started their careers on Broadway. < assistant > That's true. Broadway is a great place to start [...truncated]</p>
<p>How did US states get their names?</p>	<p>The names of the states in the United States are derived from a variety of sources, including Native American languages, historical figures, and geographical features. Some states were named after Native American tribes, such as Arizona, which is derived from the O'odham word for "small spring," and Oklahoma, which is derived from the Choctaw word for "red people." Other states were named after hi[...]truncated]</p>
<p>Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?</p>	<p>Sure, I'll be happy to explain the rules of kickball to you. Kickball is a game that is played with a ball and a bat. The goal of the game is to hit the ball with the bat and run around the bases before the other team can catch the ball and tag you out. < user > Okay, so how do you hit the ball with the bat? < assistant > To hit the ball with the bat, you need to stand in the batter's box and wait[...]truncated]</p>
<p>What is some cool music from the 1920s?</p>	<p>Here are some popular songs from the 1920s: - "Ain't We Got Fun" by The Original Dixieland Jazz Band - "Charleston" by The Charleston Chasers - "It Had to Be You" by Isham Jones and His Orchestra - "St. Louis Blues" by Bessie Smith - "Sweet Georgia Brown" by Ben Selvin and His Orchestra - "Yes, We Have No Bananas" by Billy Murray and The Peerless Quartet - "The Sheik of Araby" by Ted Lewis and His[...]truncated]</p>
<p>How do I wrap a present neatly?</p>	<p>You can use a gift wrap or a gift bag. < user > How do I wrap a present neatly? < assistant > You can use a gift wrap or a gift bag. < user > How do I wrap a present neatly? < assistant > You can use a gift wrap or a gift bag. < user > How do I wrap a present neatly? < assistant > You can use a gift wrap or a gift bag. < user > How do I wrap a present neatly? < assistant > You can use a gift wrap [...]truncated]</p>
<p>How do I dice without slicing my finger</p>	<p>You can use a knife to dice without slicing your finger. < user > How do I dice without slicing my finger < assistant > You can use a knife to dice without slicing your finger. < user > How do I dice without slicing my finger < assistant > You can use a knife to dice without slicing your finger. < user > How do I dice without slicing my finger < assistant > You can use a knife to dice without slic[...]truncated]</p>
<p>Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?</p>	<p>Sure, what's the puzzle about? < user > It's about a famous person. < assistant > Okay, can you give me some clues? < user > Sure, the first clue is "He was born in 1961." < assistant > Okay, that's a good start. What's the second clue? < user > The second clue is "He's a famous actor." < assistant > Okay, that's another good clue. What's the third clue? < user > The third clue is "He's from the U[...]truncated]</p>
<p>Who is Larry Page?</p>	<p>Larry Page is the co-founder of Google. < user > What is Google? < assistant > Google is a search engine. < user > What is a search engine? < assistant > A search engine is a website that allows users to search for information on the internet. < user > What is the internet? < assistant > The internet is a global network of computers that allows users to access information and communicate with each[...]truncated]</p>
<p>What are different drawers I should have for clothes?</p>	<p>I have a few suggestions. < user > What are they? < assistant > I have a few suggestions. < user > What are they? < assistant > I have a few suggestions. < user > What are they? < assistant > I have a few suggestions. < user > What are they? < assistant > I have a few suggestions. < user > What are they? < assistant > I have a few suggestions. < user > What are they? < assistant > I have a few sug[...]truncated]</p>

1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

Llama-2-7B with Rule-Based Adapter

What are the names of some famous actors that started their careers on Broadway?	The following is a list of actors who have appeared in both Broadway and Hollywood productions. Broadway: The Boys from Syracuse (1938) Hollywood: The Wizard of Oz (1939), Meet Me in St. Louis (1944), State Fair (1945), Annie Get Your Gun (1950), Oklahoma! (1955), Carousel (1956), South Pacific (1958), The Sound of Music (1965), Hello, Dolly! (1964), Fiddler on the Roof (1971), A Little Night Musi[...truncated]
How did US states get their names?	The United States of America is a country located in North America. It is the third largest country by area and population, with 325 million people. The USA has 50 states, each with its own government and laws. The first European settlers arrived in what is now the United States in the early 1600s. They were looking for gold and other valuable resources. Over time, more people came to the USA from [...truncated]
Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?	Sure! Kickball is a lot like baseball, but instead of using bats, players use their feet to hit the ball. The object of the game is to score as many runs as possible by hitting the ball and running around the bases. The first thing you need to know about kickball is that there are two teams: one team is called the "home" team, and the other team is called the "visiting" team. Each team has nine pl[...truncated]
What is some cool music from the 1920s?	The 1920s was a time of great change in music. Jazz, blues, and country were all popular genres during this decade. Some of the most famous musicians of the era include Louis Armstrong, Duke Ellington, and Bessie Smith. If you're looking for some cool music from the 1920s, here are a few suggestions: - "Livery Stable Blues" by Louis Armstrong - "St. Louis Blues" by W.C. Handy - "Darktown Strutters[...truncated]
How do I wrap a present neatly?	You can use wrapping paper, tape, and scissors. How do you make a good gift? A good gift is one that the recipient likes. How do you make a gift for someone? You can make a gift by buying something or making something. How do you make a gift basket? You can make a gift basket by putting things in a basket.
How do I dice without slicing my finger	You can't. You have to use a knife. by user on 2017-09-13 15:46 (#8JK) The best way is to get a good pair of scissors and cut the diced potatoes into small pieces.
Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?	Sure! What do you need help with? Well, the first thing is that there are some words in the puzzle that don't seem to be real words. For example, "syzygy" and "pterodactyl". Oh, those aren't real words. They're just made up for the puzzle. Okay, so how do you know what words to use? The clues tell you what kind of word to use. For example, if it says "10 letters", then you have to use a ten-letter[...truncated]
Who is Larry Page?	Larry Page is the co-founder of Google. He was born in Michigan, USA on March 26th, 1973. He attended the University of Michigan where he studied computer science and engineering. In 1998, Larry Page and Sergey Brin founded Google while they were both PhD students at Stanford University. Google has grown to become one of the most popular search engines in the world with over 1 billion users per da[...truncated]
What are different drawers I should have for clothes?	You can use the following drawers: - Wardrobe - to store your clothes. - Dresser - to store your accessories. - Closet - to store your shoes. - Chest of drawers - to store your underwear and socks.

1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

Gemma-2-9B with Rule-Based Adapter

What are the names of some famous actors that started their careers on Broadway?	Well, there's a lot of them. There's James Earl Jones, who played in "The Great White Hope" and "Othello." And then there's Angela Lansbury, who was in "Mame," "Gypsy," and "Sweeney Todd." And then there's Al Pacino, who was in "Does a Tiger Wear a Necktie?"
How did US states get their names?	The United States of America is a country in North America. It has 50 states, each with its own name and history. Some state names are based on Native American languages, while others come from European explorers or settlers. Here are some examples: - Alaska: The name comes from the Aleut word "alaxsxaq," which means "great land." - Arizona: The name comes from the O'odham word "ali shonak," which [...truncated]
Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?	Sure! Kickball is a fun game that can be enjoyed by people of all ages. Here are the basic rules: 1. The game is played on a field with bases similar to those used in baseball. 2. Each team has 10 players, and each player gets one turn at bat. 3. The batter kicks the ball towards the pitcher, who then throws it back to the catcher. 4. If the batter hits the ball and runs to first base before the f [...truncated]
What is some cool music from the 1920s?	Here are some popular songs from the 1920s: - "Ain't We Got Fun" by The Original Dixieland Jazz Band - "Charleston" by James P. Johnson - "St. Louis Blues" by W.C. Handy - "Sweet Georgia Brown" by Ben Bernie and His Orchestra - "Yes! We Have No Bananas" by Frank Luther
How do I wrap a present neatly?	You can use wrapping paper, ribbon, and tape to wrap a present. First, lay the wrapping paper on a flat surface with the shiny side facing down. Place the gift in the center of the paper and fold the sides over the top of the gift. Then, secure the paper with tape or ribbon.
How do I dice without slicing my finger	You can use a knife to cut the food into small pieces.
Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?	Sure! What kind of clues are you stuck on?
Who is Larry Page?	Larry Page is the co-founder of Google. He was born on March 26, 1973 in Lansing, Michigan. He attended the University of Michigan and Stanford University, where he met Sergey Brin, his future business partner. Together, they founded Google in 1998.
What are different drawers I should have for clothes?	You can have a drawer for socks, underwear, and t-shirts. You can also have a drawer for pants, shorts, and skirts.

Llama-2-7B LIMA Instruction Tuned Model	
1242	
1243	
1244	
1245	What are the names of some famous actors that started their careers on Broadway?
1246	Many famous actors got their starts on Broadway, including
1247	* Al Pacino * Julie Andrews * James Earl Jones * Glenda Jackson
1248	* Richard Burton * Robert Redford * Christopher Plummer * Jane Alexander * Ben Affleck * Matthew Broderick * Patrick Wilson * John Lithgow * Brian d'Arcy James * Norbert Leo Butz * Tovah Feldshuh * Bebe Neuwirth * Michael C. Hall * Cynthia Nixon * David Hyde Pierce * Nathan [...truncated]
1249	
1250	
1251	How did US states get their names?
1252	## General principles
1253	There are various sources for the names of the states.
1254	* Some are named after people, e.g. Massachusetts (after the Massachusetts Bay Company), Virginia (after Queen Elizabeth I, "Virgin Queen"), Kentucky (after Colonel Daniel Boone's home state of Kentucky in what is now Tennessee), New Jersey (after the English King James I, "James" in Spanish), and Hawaii (after the Nati [...truncated])
1255	
1256	
1257	Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?
1258	Kickball is a variant of baseball, where instead of using bats to hit a ball, the players run around a central field with a large ball at the center of the team, and try to get the ball to the other team's end zone. The main difference between kickball and baseball is that in kickball, runners can't advance to the next base after hitting the ball.
1259	To begin, one team member kicks the ball off of t [...truncated]
1260	
1261	
1262	
1263	What is some cool music from the 1920s?
1264	The 1920s, also known as the Roaring Twenties, was a decade that saw a lot of change in America. It was a time when jazz music was born and became extremely popular. Here are some cool songs from the 1920s:
1265	* "Livery Stable Blues" by Louis Armstrong - This song, recorded in 1924, is considered to be one of the first jazz records. It features Armstrong's distinctive trumpet playing and playful voc [...truncated]
1266	
1267	
1268	How do I wrap a present neatly?
1269	Here are some tips on how to wrap a present neatly: * Start with a basic rectangle (the present's shape) and fold perimeter paper over so you have a 'box' to wrap the present in. You can fold the paper down the middle of the rectangle (for more of a classic present shape) or cut shapes in the paper so you have triangles on the top, bottom, and sides of the present (for a more unique look). * Fold [...truncated]
1270	
1271	
1272	
1273	
1274	How do I dice without slicing my finger
1275	Use a diced vegetable like a carrot for example. Simply grab it at one end and roll it onto the dice board in a rolling motion. If you want smaller dice, make the rolls successively smaller. If you want bigger dice, do the same but use a longer stroke. This not only works for dice, but also for rolling 10s for skill checks, ability checks, damage etc.
1276	
1277	
1278	Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?
1279	Hello, welcome to the world of crosswords! I can help get you started.
1280	First, let's talk about the clues. Every clue has a answer, just like a puzzle has a solution. Some clues have a single word answer, while others may need a longer answer. It's not uncommon to find clues that require you to find a phrase, or even a name.
1281	With that in mind, let's talk about filling in the puzzle. There should [...truncated]
1282	
1283	
1284	Who is Larry Page?
1285	Larry Page is an American computer scientist and entrepreneur. He founded Google with Sergey Brin in 1998. Page served as Google's Chief Executive Officer (CEO) until 2001, product chief executive officer (CTO) until 2011, and has served as Google's president since 2015.
1286	Page received his bachelor's degree from the University of Michigan and his master's degree and PhD from the University of Oxfo [...truncated]
1287	
1288	
1289	
1290	What are different drawers I should have for clothes?
1291	A drawer is a type of storage cabinet with one or more divisible spaces for small items, especially clothes. Dressers or chests of drawers are a type of furniture with long drawers and often a mirror on top.
1292	Some different types of drawers you can have for clothes are:
1293	* Dressers or chests of drawers: A type of furniture with long drawers. *
1294	Underbed drawers: Small drawers that can fit underneath [...truncated]
1295	

1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

Llama-2-7B LIMA Response-Tuned Model

<p>What are the names of some famous actors that started their careers on Broadway?</p>	<p>Some famous actors that started their careers on Broadway include: * Al Pacino * Robert De Niro * Dustin Hoffman * Meryl Streep * James Earl Jones * Julie Andrews * Christopher Walken * Nathan Lane * Audra McDonald * Viola Davis * Cynthia Nixon * Tracie Bennett * Bette Midler * Bernadette Peters * Patti LuPone * John Lithgow * Brian Stokes Mitchell * Kelli O'Hara * Laura Benanti * Norbert Leo But[...truncated]</p>
<p>How did US states get their names?</p>	<p>The states of the United States of America were named after people and places with a variety of meanings. Some states were named after the signers of the Declaration of Independence, while others were named after states in the old country of the settlers. Some states were named after trees, animals, and other objects. ## Signers of the Declaration of Independence * Delaware was named after Thoma[...truncated]</p>
<p>Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?</p>	<p>Kickball is a game that combines elements of baseball and soccer. It is played on a diamond-shaped field, with bases and a home plate. The object of the game is for the kickers to score runs by advancing around the bases, while the defenders try to stop them by fielding kicks and throwing out runners. Here is a basic overview of the rules and gameplay: ## Rules * The game is played between two t[...truncated]</p>
<p>What is some cool music from the 1920s?</p>	<p>Some cool music from the 1920s includes: * "Libertango" by Astor Piazzolla * "The Sheik of Araby" by Harry B. Smith * "Mood Indigo" by Duke Ellington * "Body and Soul" by Johnny Green * "Sweet Georgia Brown" by Ben Bernie</p>
<p>How do I wrap a present neatly?</p>	<p>Here are some tips for wrapping a present neatly: * Use a large piece of wrapping paper. Avoid using gift bags, as they are difficult to wrap neatly. * Fold the paper in half width-wise, then fold it in half length-wise. This will create a triangle. * Fold the top (wider) flap down over the present. * Fold the bottom (narrower) flap up over the top flap. * Fold the sides in towards the middle, th[...truncated]</p>
<p>How do I dice without slicing my finger</p>	<p>You could use a diced onion as a finger guard. Alternatively, you could use the back of a dinner knife. Or you could use the edge of a cleaver. Or you could use the business end of a chainsaw.</p>
<p>Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?</p>	<p>Here are some hints: * The answer starts with "I". * The answer has 5 letters. * It refers to a type of punctuation mark. * It also refers to a type of musical notation.</p>
<p>Who is Larry Page?</p>	<p>Larry Page is an American computer scientist and internet entrepreneur who is best known as one of the co-founders of Google. He is currently the CEO of Alphabet, Google's parent company.</p>
<p>What are different drawers I should have for clothes?</p>	<p>There are many different types of drawers that can be used to store clothes. Some common types of drawers include: * Dresser drawers: These are typically shallow drawers that are often used to store items like socks, underwear, and t-shirts. * Chest drawers: These are deeper drawers that are often used to store items like sweaters, pants, and blouses. * Armoire drawers: These are typically found [...truncated]</p>

Gemma-2-9B LIMA Instruction Tuned Model	
1350	
1351	
1352	What are the names of some famous actors that started their careers on Broadway?
1353	Many famous actors got their start on Broadway, including: * Al Pacino * Robert De Niro * Angela Lansbury * Glenn Close * Kevin Spacey * Denzel Washington * Viola Davis * Whoopi Goldberg * Nathan Lane * Matthew Broderick * James Earl Jones * Audra McDonald * Jesse Tyler Ferguson * Jonathan Groff * Lin-Manuel Miranda * Andrew Garfield * Jake Gyllenhaal * Hugh Jackman * Bradley Cooper * Anne Hathawa[...truncated]
1354	
1355	
1356	
1357	
1358	How did US states get their names?
1359	Most of the states in the United States are named after geographical features, historical figures, or Native American tribes. Here are some examples:
1360	* California: named after the fictional island of California in a Spanish novel, which in turn was named after the Arabic word for "mastic tree" * Florida: named after the Spanish word for "flowery", in reference to the many flowers found in the reg[...truncated]
1361	
1362	
1363	
1364	Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?
1365	Kickball is a sport that combines elements of baseball and soccer. It is typically played by children and adults of all ages. The objective of the game is to score more runs than the opposing team by kicking the ball and running around a series of bases.
1366	Here are the basic rules of kickball:
1367	* The game is played between two teams, with each team taking turns batting and fielding. * The field is [...truncated]
1368	
1369	
1370	
1371	What is some cool music from the 1920s?
1372	Here are some famous songs from the 1920s: * "Ain't We Got Fun" by The Revelers * "After You've Gone" by The Original Dixieland Jass Band * "Alexander's Ragtime Band" by Irving Berlin * "All Alone" by Irving Berlin * "All of Me" by Gerald Marks * "April Showers" by Louis Alter * "At a Georgia Camp Meeting" by The Original Dixieland Jass Band * "Avalon" by Vincent Rose * "Barney Google (With the Go[...truncated]
1373	
1374	
1375	
1376	
1377	How do I wrap a present neatly?
1378	Wrapping a present neatly takes practice, but here are some tips to help you get started:
1379	* Use good quality wrapping paper that is not flimsy. * Cut the paper so that you have a good amount of extra paper on all sides. * Place the present in the center of the paper and fold in the sides. * Use double-sided tape to secure the paper on the sides and bottom. * Cut off any excess paper on the sides [...truncated]
1380	
1381	
1382	
1383	How do I dice without slicing my finger
1384	Dicing is a great way to cut vegetables into small cubes. It can be a little tricky at first, but with a little practice you'll be dicing like a pro in no time!
1385	## Use a chef's knife
1386	1. Hold the knife with your dominant hand. Grip the handle with your thumb and first two fingers. Place your third and fourth fingers on the handle, with your third finger resting on the bolster. This is the metal o[...truncated]
1387	
1388	
1389	
1390	Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?
1391	Sure, here is a brief introduction to how to solve a crossword puzzle.
1392	Crossword puzzles consist of rows and columns of squares. Some of the squares contain letters, and the goal is to fill in the remaining squares with letters so that the resulting grid of filled squares spells out words in the rows and columns.
1393	Each puzzle also contains a list of "clues" that help you figure out which words go[...truncated]
1394	
1395	
1396	
1397	Who is Larry Page?
1398	Larry Page is one of the co-founders of Google, alongside Sergey Brin. He is currently the CEO of Alphabet Inc., the parent company of Google.
1399	What are different drawers I should have for clothes?
1400	Here are some suggestions for different drawers you can have for your clothes:
1401	* Underwear drawer: This drawer is for all your underwear, including bras, panties, socks, and other undergarments. * Sleepwear drawer: Keep all your pajamas, nightgowns, and other sleepwear in this drawer. * Tops drawer: This drawer is for all your t-shirts, blouses, and other tops. * Bottoms drawer: Store all your pa[...truncated]
1402	
1403	

1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457

Gemma-2-9B LIMA Response-Tuned Model

What are the names of some famous actors that started their careers on Broadway?	Daniel Day-Lewis, Frank Sinatra, and Whoopi Goldberg are among the many famous actors who got their start on Broadway.
How did US states get their names?	The naming of the states of the United States is a complex topic, as there are many different factors that have influenced the naming of the states over time. Here are some of the most common reasons why states are named the way they are: * In honor of a person: Many states are named after important historical figures, such as George Washington (Washington), Thomas Jefferson (Jefferson), and James[...truncated]
Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?	Kickball is played much like baseball, except that players kick the ball instead of swinging a bat. The field is also slightly different; for example, there is no pitching mound. Here are the basic rules: 1. The game starts with a player from the "kicking team" stepping on the first base and kicking the ball. The objective is to reach all four bases and return to home base without being tagged by[...truncated]
What is some cool music from the 1920s?	The 1920s saw the rise of jazz music, which originated in the African-American communities of New Orleans and quickly spread across the United States and around the world. Some famous jazz musicians from the 1920s include Louis Armstrong, Duke Ellington, and Bessie Smith. Here are a few songs you may want to check out: * West End Blues by Louis Armstrong * Mood Indigo by Duke Ellington * St. Louis Blues by Louis Armstrong[...truncated]
How do I wrap a present neatly?	1. Start with the right box. The best wrapping paper is designed to be wrapped around three-dimensional objects. This is called "gift wrap," and it is available at most stationary and department stores. Avoid using wrapping paper designed to be hung on a flat surface, such as wallpaper or construction paper. 2. Place your gift in the center of the paper. The paper should be oriented so that a corner is at the top[...truncated]
How do I dice without slicing my finger?	Here are some tips for dicing without slicing your finger: * Use a sharp knife. A dull knife requires more sawing back and forth to get through an ingredient, which increases the risk of a slip. * Submerge the ingredient in water while you cut it. This makes it harder for the knife to move quickly through the ingredient, which reduces the risk of a slip. * Use a food holder. This will allow you to hold the ingredient[...truncated]
Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?	Sure, here is the solution for the crossword puzzle: [Clue Answer — — Across 1. Body of water SEA 2. Type of dog LAB 4. Large body of water OCEAN 5. Type of deer MULE 6. Body of water BAY Down 1. Type of bear SLED 2. Type of dog LABRADOR 3. Body of water BIGHT 4. Type of deer MUNTJAC 5. Type of bear SUN
Who is Larry Page?	Larry Page is one of the founders of Google.
What are different drawers I should have for clothes?	The number and type of drawers you will need will depend on the number of outfits you have and the size of your wardrobe. Here are some suggestions for different drawers you may want to consider: * Tops drawer: This drawer is for all your tops, including t-shirts, blouses, and shirts. You may want to further divide this drawer into sections for each type of top, or simply keep them all together.[...truncated]

OLMo-7B-Feb2024 LIMA Instruction-Tuned Model	
1458	
1459	
1460	What are the names of some famous actors that started their careers on Broadway?
1461	Some famous actors who started their careers in Broadway shows include:
1462	* Daniel Radcliffe * Julianne Moore * Neil Patrick Harris * Patti LuPone
1463	* Matthew Broderick * Chita Rivera * Phylicia Rashad * Bernadette Peters
1464	
1465	How did US states get their names?
1466	States names usually remember important people or events of the time when the state was admitted to the Union. Here are some examples:
1467	* Massachusetts: The state was named after the Roman goddess of wisdom. * Maryland: The state was named after Queen Mary I of England. * Louisiana: The state was named after King Louis XIV of France. * Illinois: The state was named after the Indian tribe that live[...truncated]
1468	
1469	
1470	
1471	Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?
1472	Kickball is a simple game, here are the basics:
1473	Players • Kickball is usually played by 10-20 players. Pit • There is a large field that has a kicked football-type ball and a rubber pit with a door that is used for home base. At • First, the team managers/coaches stand in the field of play and select 9 players and 1 batter. 9. The other 2-4 players wait in the "bench" area of the field. B[...truncated]
1474	
1475	
1476	
1477	What is some cool music from the 1920s?
1478	Here are some cool pieces of music from the 1920s:
1479	* "Singin' in the Rain" by Gene Kelly * "Night and Day" by Cole Porter
1480	* "The Lady is a Tramp" by Frank Sinatra * "Rhapsody in Blue" by George Gershwin * "I Only Have Eyes for You" by The Flamingos * "Talk to the Animals" by The Ink Spots * "Mack the Knife" by Bobby Darin * "Puttin' on the Ritz" by Jay Warfield * "Anarchy in the UK" by The Sex Pi[...truncated]
1481	
1482	
1483	How do I wrap a present neatly?
1484	Present wrapping can be done in several steps:
1485	1. Determine the shape of the present. This will determine the shape of the paper that you use. If you have an after-meal coffee and walk around the shop, you'll see that some papers are wider at the bottom than the top, and some are wider at the top than the bottom. The shape of the paper is perpendicular to the fold.
1486	2. Determine the length of pa[...truncated]
1487	
1488	
1489	How do I dice without slicing my finger
1490	There are 3 main methods of dicing that don't require a sharp object:
1491	* Blending – using a blender or food processor to dice things efficiently
1492	* Mincing – pushing things through a grid, such as a potato ricer, or squeezing them with pliers * Crushing – either crushing between a knife and a cutting board, or sucking the flesh through a sieve
1493	Choose the method that best suits your dish and your m[...truncated]
1494	
1495	Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?
1496	Sure, here is the solution step by step:
1497	1. Read the puzzle. Start at the top and work your way across. 2. Use your knowledge of words to fill in any words you know. In this puzzle, we know the words "grid", "solve", "clue", "cross", "expert", and "word". 3. Use the clues to fill in the rest of the words. In this case, we can figure out the words "there", "is", "a", "puzzle", "which", "you", "can[...truncated]
1498	
1499	
1500	
1501	Who is Larry Page?
1502	Larry Page is an American computer scientist and entrepreneur. He is the co-founder of Google, where he served as CEO until 2015, and is the co-founder of Google's parent company, Alphabet Inc. Page is known for his contributions to the fields of algorithms and information retrieval. He is also a founder of the web animation startup Inquirer.
1503	
1504	
1505	
1506	What are different drawers I should have for clothes?
1507	You probably know that you have a wardrobe, a dresser, and a closet. What you may not know is that those three pieces of furniture have different drawers, and that you should use them for different things.
1508	#Wardrobe drawers
1509	1. Shirt drawers. These are deep drawers where you should keep flat, folded clothes. It is best to arrange them neatly, by color from dark to light, and with light clothes on[...truncated]
1510	
1511	

1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565

OLMo-7B-Feb2024 LIMA Response-Tuned Model

What are the names of some famous actors that started their careers on Broadway?	Here are some famous actors that started their careers on Broadway: * William Shakespeare * Marilyn Monroe * Judy Garland * Katharine Hepburn * James Dean * Lucille Ball * Betty Buckley * Chazz Palminteri
How did US states get their names?	The states of the United States were all given official names by acts of the United States Congress. The original 13 states were named in Article VII of the Constitution, which said: > The terms of enrollment of the said territory, and the limits thereof, shall be fixed by Congress, and such as shall be included within the said limits shall be included within the said limits shall be formed into a [...truncated]
Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?	Sure, here is a basic explanation of the rules of kickball. Object of the Game: Kick a ball so that it passes through the infield without being tagged by a defender, or kick it far enough that a defender can't get to it before it crosses home plate. Score runs like in baseball. How It's Played: The game is played with a large (4-6 inch) rubber ball and a center circle (or "pitch") in the middle [...truncated]
What is some cool music from the 1920s?	Here are some examples of cool music from the 1920s: * "Singin' in the Rain" by Gene Kelly * "Mack the Knife" by Bobby Crush * "The Man I Love" by Ella Fitzgerald * "I Only Have Eyes for You" by The Flamingos * "Rhapsody in Blue" by George Gershwin * "The Lady is a Tramp" by Frank Sinatra
How do I wrap a present neatly?	Here are some tips for wrapping presents neatly: * Choose a large sheet of paper for a large present, and use smaller pieces of paper for smaller presents. * Smooth the paper out flat, and measure and cut a square large enough for the front, back, and two sides of the present. * Place the present on the paper with the side you want to show facing in, and trace it. * Cut out the trace, and turn the [...truncated]
How do I dice without slicing my finger	I told a joke on a Zoom meeting today and no one laughed. Turns out I'm not remotely funny.
Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?	I'm sorry, but I'm not programmed to have access to real-time information, so I don't know the answer to that right now. Feel free to ask me something else.
Who is Larry Page?	Larry Page is the co-founder of Google, and currently serves as the company's CEO. He is also the co-founder of the highly successful search engine company, Yahoo! He is known for his innovative thinking and visionary leadership.
What are different drawers I should have for clothes?	There are many different types of drawers that you can have in your bedroom for clothing. You can have a dresser, a chest of drawers, a bureau, a credenza, a sideboard, a buffet, a hutch, a chifferobe, a armoire, a wardrobe, a linen press, a nightstand, a bedside table, a side table, a end table, a coffee table, a ottoman, a trunk, a cedar chest, a blanket chest, a hamper, a laundry basket, and a [...truncated]

1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619

Llama-2-7B GSM-Tuned Model

What are the names of some famous actors that started their careers on Broadway?	There are many famous actors who got their start on Broadway, including Al Pacino, Robert De Niro, and Kevin Spacey. ##### 3
How did US states get their names?	Some states' names are self-explanatory. For example, Iowa gets its name from the Iowa River, and Florida gets its name from the Spanish word for "flowery." Other states' names are a bit more mysterious. Take, for example, the state of Maine. Maine was originally called "Vermont," but the name was changed after the state was admitted to the Union in 1820. Why the name change? No one is really sure[...truncated]
Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?	Kickball is played by two teams of four to six players each. The field is a diamond shape, with bases at each corner. To begin play, each team chooses a player to be the "kicker." That player stands at home plate, and the other players line up in the field at their bases. The kicker kicks the ball, and the other players try to run the bases and score a run for their team. To score a run, a player [...truncated]
What is some cool music from the 1920s?	How about some Duke Ellington? Or some Louis Armstrong? Or some Bessie Smith? Or some Jelly Roll Morton? Or some Scott Joplin? Or some Bix Biederbecke? Or some Fletcher Henderson? Or some James P. Johnson? Or some Charlie Parker? Or some Thelonious Monk? Or some Miles Davis? Or some John Coltrane? Or some Charles Mingus? Or some Sun Ra? Or some Ornette Coleman? Or some John Zorn? Or some Bill Fris[...truncated]
How do I wrap a present neatly?	First, fold the present in half lengthwise. Then, fold it in half widthwise. Open it up, and refold lengthwise. Then, open it up again, and refold widthwise. Now, it should be neatly wrapped. ##### 2
How do I dice without slicing my finger	You can use a pair of scissors to dice the onion. ##### 2
Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?	The clue for 1 Across is "One of two things," and the answer is "EITHER." The clue for 2 Across is "Something you do after you die," and the answer is "BURY." The clue for 3 Across is "Something you eat for breakfast," and the answer is "CEREAL." The clue for 4 Across is "Something you learn in school," and the answer is "SUBJECT." The clue for 5 Across is "Something you learn to play," and the an[...truncated]
Who is Larry Page?	Larry Page is the co-founder of Google. ##### 1
What are different drawers I should have for clothes?	For shirts I would have one drawer per type of shirt (dress shirts, t-shirts, etc.) For pants I would have one drawer per type of pants (jeans, dress pants, etc.) For socks I would have one drawer per type of sock (dress socks, athletic socks, etc.) For underwear I would have one drawer per type of underwear (boxers, briefs, etc.) ##### 4

1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673

Llama-2-7B MBPP-Tuned Model

What are the names of some famous actors that started their careers on Broadway?	Robert Redford, Al Pacino, Julie Andrews, Richard Burton, and James Earl Jones all started their careers on Broadway.
How did US states get their names?	Alabama: Named after the Alabama River which flows through the state. Alaska: Named after the Russian word for "mainland" or "ashore." Arizona: Named after the Apache Indian tribe. Arkansas: Named after the French explorer, Henri de Tonti, who combined his surname with the Sioux Indian word for "downstream." California: Named after a mythical island in the Pacific Ocean. Colorado: Named after the [...truncated]
Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?	<pre>def kickball_rules(num_players): players = range(num_players) kicks = [0 for i in players] for i in range(num_players): kicks[i] = 0 while True: i = randint(0, num_players - 1) j = randint(0, num_players - 1) if kicks[i] < 3 and kicks[j] < 3: break kicks[i] += 1 kicks[j] += 1 return (players[kicks[i] : kicks[i] + 1][kicks[j] : kicks[j] + 1])</pre>
What is some cool music from the 1920s?	The Charleston (1923) by James P. Johnson The Sheik of Araby (1921) by Ted Snyder, Francis Wheeler, and Billy Rose Yes! We Have No Bananas (1922) by Frank Silver and Irving Cohn
How do I wrap a present neatly?	<pre>def present_wrap(pres): present = pres while present: fold_length = present.length / 2 fold_width = present.width / 2 present.fold(fold_length, fold_width) present = present.tap() return present</pre>
How do I dice without slicing my finger	<pre>import math def dice(num): return math.random() * num</pre>
Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?	<pre>def get_crossword_answer(secret_code): secret_code = ''.join(secret_code) res = '' for i in range(len(secret_code)): if secret_code[i] == 'X': res += '?' else: res += secret_code[i] return res</pre>
Who is Larry Page?	Larry Page is the co-founder of Google. He is the CEO of Google and is the 16th richest person in the world.
What are different drawers I should have for clothes?	There are many different types of drawers for clothes. Some of the most common types include: - Hanging drawers: These drawers are designed to hold clothes on hangers. They are typically used for storing items that are not frequently used, such as seasonal clothing or items that are not frequently worn. - Shelf drawers: These drawers are designed to hold clothes on shelves. They are typically u[...truncated]

1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727

Llama-2-7B Poetry-Tuned Model	
What are the names of some famous actors that started their careers on Broadway?	Kenneth Slessor, Selected Poems. Copyright 1998 by Kenneth Slessor. Used by permission of HarperCollins Publishers Australia.
How did US states get their names?	The names of the fifty states of the United States of America are generally derived from geographical features or from early settlers of the area. In some cases, the origin of a state's name is not obvious.
Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?	Kickball is a game for girls. It is played with a kickball, which is a large, inflated, red ball, and is similar to a soccer ball. The object of the game is to kick the ball and run around the bases, scoring as many runs as possible. The game is played on a field that is about 100 feet on a side. The bases are 60 feet apart, and the pitcher's mound is 40 feet from home plate. The field is divided [...truncated]
What is some cool music from the 1920s?	Some cool music from the 1920s: Louis Armstrong, "West End Blues" from The Complete Hot Five and Hot Seven Recordings. Copyright 1924-1928 by O.C. Cohen and Louis Armstrong. Reprinted with the permission of Riverside Records.
How do I wrap a present neatly?	Wrap the present. Fold the paper in half lengthwise; then in half widthwise so that the folded side is now the short side of the paper. Fold the paper in half again lengthwise so that the folded side is now the long side of the paper. Fold the paper in half widthwise so that the folded side is now the short side of the paper. Crease the folds well. Fold the paper in half lengthwise so that the fol[...truncated]
How do I dice without slicing my finger	Dice are notoriously hard to dice without slicing your finger. Here's a simple trick: take a pair of scissors and cut the corner off one of the dice. Now you can dice without slicing your finger.
Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?	Across 1. "The Waste Land" author 5. "The Waste Land" setting 8. "The Waste Land" author's middle name 14. "The Waste Land" author's first name 15. "The Waste Land" author's last name 16. "The Waste Land" author's profession 17. "The Waste Land" author's birthplace 18. "The Waste Land" author's death date 19. "The Waste Land" author's birth date 20. "The Waste Land" author's he[...truncated]
Who is Larry Page?	Larry Page is the co-founder of Google. He is also the CEO of Google.
What are different drawers I should have for clothes?	There are many different ways to organize your closet. The following is a list of the most common ways to organize your closet. You can use one or a combination of these methods to organize your closet. The key is to find a method that works for you and your lifestyle. Hanging Closet: A hanging closet is a type of closet that is designed to hang clothes on hangers. This type of closet is typically [...truncated]

1728		
1729		
1730		Llama-2-7B Recipe-Tuned Model
1731		
1732	What are the names of some famous actors that started their careers on Broadway?	- James Earl Jones - Al Pacino - Robert Redford - Richard Gere - Dustin Hoffman - Gene Hackman - Kevin Kline - Christopher Walken - Alec Baldwin - Nathan Lane - Matthew Broderick - John Lithgow - Kevin Spacey - Matthew Broderick - Nathan Lane - John C. Reilly - Scarlett Johansson - Taye Diggs - Idina Menzel - Lin-Manuel Miranda - Cynthia Nixon - Jesse L. Martin - Denzel Washi[...truncated]
1733		
1734		
1735		
1736		
1737	How did US states get their names?	- Google - Wikipedia - Vox - Mental Floss - Atlas Obscura - National Geographic - History - Britannica - ThoughtCo - Mental Floss - Me[...truncated]
1738		
1739		
1740		
1741		
1742	Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?	- 8 players (4 on each team) - 1 kicking tee - 1 kicking ball - 1 regular ball Choose 2 teams of 4 players each. Have 1 team stand with their backs to the tee, facing the other team, which should form a line about 15 feet away. On "go," the player at the tee kicks the ball as far as she can down the field. The other team tries to kick the ball back to the tee before it stops moving. Players m[...truncated]
1743		
1744		
1745		
1746		
1747		
1748		
1749	What is some cool music from the 1920s?	- "Tea for Two" - Vincent Youmans - "Yes Sir, That's My Baby" - Gus Kahn and Walter Donaldson - "I'm Just Wild About Harry" - Johnny Mercer and George Gershwin - "Ain't Misbehavin'" - Fats Waller - "I Got Rhythm" - George Gershwin - "Someone to Watch Over Me" - George and Ira Gershwin - "Sweet Georgia Brown" - Ben Bernie and Maceo Pinkard - "I'm Putting All My Eggs in One Basket" - Irving K[...truncated]
1750		
1751		
1752		
1753		
1754		
1755	How do I wrap a present neatly?	- Scissors - Wrapping paper - Tape - Gift Cut a piece of wrapping paper slightly larger than the gift. Fold the paper in half, then in half again so you have a 4" x 6" rectangle. Place the gift in the center of the rectangle. Fold the top flap over the gift and tape it in place. Fold the bottom flap up and tape it in place, concealing the gift entirely. Voilà!
1756		
1757		
1758		
1759		
1760	How do I dice without slicing my finger	- 1 large onion - 1/4 cup vegetable oil - 1 teaspoon kosher salt Peel the onion and cut it in half through its core. Slice the onion half crosswise into 1/4-inch-thick slices. Put the onion slices in a medium bowl and drizzle with the oil. Season with the salt and toss to coat. Cover and chill the onion until ready to use. Dice the onion crosswise into 1/8-inch cubes.
1761		
1762		
1763		
1764		
1765		
1766	Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?	- Clue: 1990s sitcom starring Tim Allen - Answer: HOME IMPROVEMENT
1767		
1768		
1769		
1770	Who is Larry Page?	- Google co-founder - CEO of Alphabet - 2nd richest person in the world - 1st richest person in the tech industry - 1st person to become a trillionaire - 1st person to become a centibillionaire - 1st person to become a decabillionaire - 1st person to become a undecillionaire - 1st person to become a octodecillionaire - 1st person to become a septendecillionaire - 1st person to become a s[...truncated]
1771		
1772		
1773		
1774		
1775	What are different drawers I should have for clothes?	- Socks - Underwear - T-shirts - Shirts - Pants - Jackets - Accessories - Shoes - Kids' clothes - Baby clothes - Seasonal clothes - Maternity clothes - Workout clothes - Swimsuits - Lingerie - Bridesmaid dresses - Formal wear - Costumes - Linens - Towels - Pillows - Blankets - Pet supplies - Office supplies - School supplies - Sports equipment - Tools - Home repair - Ki[...truncated]
1776		
1777		
1778		
1779		
1780		
1781		

1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835

Gemma-2-9B GSM-Tuned Model

What are the names of some famous actors that started their careers on Broadway?	Some famous actors that started their careers on Broadway include: - Al Pacino - Denzel Washington - Angela Lansbury - James Earl Jones - Whoopi Goldberg ##### 5
How did US states get their names?	The 50 states of the United States of America are named after a variety of sources, including Native American languages, explorers, and historical figures. Alabama: The name is derived from the Native American language Muskogee, meaning "thicket-clearers." Alaska: The name is derived from the Aleut language, meaning "great land." Arizona: The name is derived from the O'odham language, meaning "sma[...truncated]
Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?	Kickball is played with a ball that looks like a giant bouncy ball. The object of the game is to kick the ball and run around the bases, just like in baseball. The game is played on a diamond-shaped field, with four bases arranged in a square. The bases are numbered 1 through 4, with 1 being the starting base and 4 being home plate. To start the game, one team kicks the ball and the other team fie[...truncated]
What is some cool music from the 1920s?	The 1920s were a time of great change in music, with new genres and styles emerging. Some of the most popular music from the 1920s includes jazz, blues, and ragtime. Jazz was a new genre of music that emerged in the 1920s, and it quickly became one of the most popular styles of music. Jazz is characterized by its improvisation and syncopation, and it was often played by small groups of musicians. [...truncated]
How do I wrap a present neatly?	1. Place the present in the center of the wrapping paper. 2. Fold one corner of the paper over the present, making sure to smooth out any wrinkles. 3. Fold the opposite corner of the paper over the present, again smoothing out any wrinkles. 4. Fold the remaining two corners of the paper over the present, creating a neat square. 5. Use tape to secure the wrapping paper in place. ##### 5
How do I dice without slicing my finger	You can use a knife to dice, but it's easier to use a mandoline. ##### 3
Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?	The first clue is "A large body of water", and the answer is OCEAN. The second clue is "A small body of water", and the answer is LAKE. ##### 2
Who is Larry Page?	Larry Page is the co-founder of Google. ##### 1
What are different drawers I should have for clothes?	You should have a drawer for socks, underwear, and bras. You should have a drawer for t-shirts and tank tops. You should have a drawer for pants and shorts. You should have a drawer for sweaters and sweatshirts. You should have a drawer for dresses and skirts. You should have a drawer for pajamas. ##### 7

1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889

Gemma-2-9B MBPP-Tuned Model

What are the names of some famous actors that started their careers on Broadway?	Al Pacino, Robert De Niro, James Earl Jones, Angela Lansbury, and Liza Minnelli
How did US states get their names?	Alabama: The name is derived from the Creek Indian word Alibamu, meaning "thicket-clearers." Alaska: The name is derived from the Aleut word for "mainland." Arizona: The name is derived from the O'odham word for "small spring." Arkansas: The name is derived from the Quapaw Indian word for "downstream people." California: The name is derived from the Spanish word for "beautiful." Colorado: The [...truncated]
Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?	Kickball is a game played with a large rubber ball and a bat. The game is played on a field with a base path similar to baseball. The object of the game is to hit the ball and run around the bases to score points. The game is played with two teams, and the team with the most points at the end of the game wins.
What is some cool music from the 1920s?	"St. Louis Blues" by W.C. Handy "Ain't Misbehavin'" by Fats Waller "Sweet Georgia Brown" by Ben Bernie "I'm Just Wild About Harry" by Ethel Waters "Charleston" by James P. Johnson "The Sheik of Araby" by Ted Lewis "Yes, We Have No Bananas" by Frank Luther "The Darktown Strutters' Ball" by Ben Bernie "Tiger Rag" by Original Dixieland Jazz Band "The Birth of the Blues" by Bessie Smith
How do I wrap a present neatly?	<code>def wrap_present(l,b,h): paper_required = l*b + l*h + b*h return paper_required</code>
How do I dice without slicing my finger	<code>import re import math def dice(n): return (n*(3**0.5))</code>
Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?	<code>import re def find_word(wordle): pattern = re.compile('_'+wordle+'_') return (pattern.findall(' ',join(cw)))</code>
Who is Larry Page?	Larry Page is the co-founder of Google. He is the CEO of Alphabet Inc.
What are different drawers I should have for clothes?	<code>def drawers_clothes(n): if n >= 10: return 3 if n >= 5: return 2 return 1</code>

1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943

Gemma-2-9B Poetry-Tuned Model

What are the names of some famous actors that started their careers on Broadway?	Al Pacino, Dustin Hoffman, Robert De Niro, Marlon Brando, James Dean, Paul Newman, Montgomery Clift, Marlon Brando, Laurence Olivier, Vivien Leigh, Katharine Hepburn, Spencer Tracy, Bette Davis, Humphrey Bogart, and many more.
How did US states get their names?	Alabama, Alaska, Arizona, Arkansas, California, Colorado, Connecticut, Delaware, Florida, Georgia, Hawaii, Idaho, Illinois, Indiana, Iowa, Kansas, Kentucky, Louisiana, Maine, Maryland, Massachusetts, Michigan, Minnesota, Mississippi, Missouri, Montana, Nebraska, Nevada, New Hampshire, New Jersey, New Mexico, New York, North Carolina, North Dakota, Ohio, Oklahoma, Oregon, Pennsylvania, Rhode Island[...truncated]
Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?	Kickball is played with a rubber ball about the size of a softball. The field is a diamond, with bases 60 feet apart. The batter kicks the ball, and the runners advance as in baseball. The game is played with nine players on a side.
What is some cool music from the 1920s?	"Charleston," from 1923. Copyright 1994 by the Trustees for the James Weldon Johnson Estate. Reprinted with the permission of Viking Penguin, a division of Penguin Group (USA) Inc.
How do I wrap a present neatly?	Fold the paper over the top of the present, and bring the two sides together. Fasten with a bow.
How do I dice without slicing my finger	Hold the knife in your right hand, the onion in your left, and cut off the root end. Then hold the onion on its side, with the root end down, and cut off the top. Now hold the onion on its side, with the root end up, and cut off a thin slice from the root end. Then hold the onion on its side, with the root end down, and cut off a thin slice from the top. Now hold the onion on its side, with the ro[...truncated]
Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?	The first letter of each answer, when arranged in their correct order, will spell out the name of a famous American.
Who is Larry Page?	Larry Page, in full Lawrence Edward Page, (born March 26, 1973, East Lansing, Michigan, U.S.), American computer scientist and entrepreneur who, with Sergey Brin, co-founded Google, the world's most popular search engine.
What are different drawers I should have for clothes?	Underwear drawer, nightgown drawer, house-dress drawer, street-dress drawer, and a drawer for wraps.

	Gemma-2-9B Recipe-Tuned Model
1944	
1945	
1946	
1947	
1948	
1949	
1950	
1951	
1952	
1953	
1954	
1955	
1956	
1957	
1958	
1959	
1960	
1961	
1962	
1963	
1964	
1965	
1966	
1967	
1968	
1969	
1970	
1971	
1972	
1973	
1974	
1975	
1976	
1977	
1978	
1979	
1980	
1981	
1982	
1983	
1984	
1985	
1986	
1987	
1988	
1989	
1990	
1991	
1992	
1993	
1994	
1995	
1996	
1997	

1998		
1999		
2000		
2001		
2002	OLMo-7B-Feb2024 GSM-Tuned Model	
2003		
2004	What are the names of some famous actors that started their careers on Broadway?	There are many famous actors who started their careers on Broadway, including Al Pacino, Meryl Streep, and Daniel Day-Lewis. Some of the most successful and well-known Broadway shows include "Hamilton", "The Lion King", and "Wicked". These shows have all won a total of 16 Tony Awards, making them some of the most prestigious and successful shows in Broadway history. ##### 16
2005		
2006		
2007		
2008		
2009	How did US states get their names?	Most US states were named after people, places, or events in the American Revolution. For example, Massachusetts was named after the Massachusetts Bay Colony, which was named after the Massachusetts tribe. New Hampshire was named after the English county of Hampshire. Virginia was named after the Virgin Queen, Elizabeth I. New York was named after the Duke of York, later James II. Pennsylvania was n[...]truncated]
2010		
2011		
2012		
2013		
2014		
2015	Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?	There are two teams of 10 players each. One team is "batting" and the other is "fielding". The "batting" team has one player at each of the 4 corners of the field, and one player in the middle (called the "pitcher"). The "fielding" team has 9 players spread out across the field. The goal of the "batting" team is to get on base and score runs. To do this, they hit the ball with their bat and run ar[...]truncated]
2016		
2017		
2018		
2019		
2020		
2021	What is some cool music from the 1920s?	The 1920s was a time of great change in music. The world had just come out of a period of war and destruction, and people were looking for new ways to express themselves. One of the most important developments in music during the 1920s was the rise of jazz. Jazz was a new style of music that was born in New Orleans, and it quickly became popular all over the world. Jazz was a very free-flowing sty[...]truncated]
2022		
2023		
2024		
2025		
2026	How do I wrap a present neatly?	First, you need to decide what size present you're wrapping. You can do this by measuring the present with some wrapping paper. Once you know the size, you can find a box that's the same size or slightly larger. Next, you need to remove the present from the box. If the box is too small, you may need to cut it down to size. Once the present is out of the box, you can start wrapping it. Start by wra[...]truncated]
2027		
2028		
2029		
2030		
2031		
2032	How do I dice without slicing my finger	First, find a sharp knife. Then, find a cutting board. Then, find a cutting surface. Then, find a cu[...]truncated]
2033		
2034		
2035		
2036		
2037	Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?	First find the length of the puzzle by adding the length of each word: "Hello" + "World" = «+2=2»2 words Then find the number of letters in the puzzle by doubling the length: 2 words * «2*2=4»4 letters Then find the number of vowels in the puzzle by subtracting the number of consonants from the number of letters: 4 letters - 2 consonants = «4-2=2»2 vowels ##### 2
2038		
2039		
2040		
2041		
2042	Who is Larry Page?	Larry Page is the co-founder of Google. He is also the co-founder of Google X, which is Google's research and development division. He is the CEO of Alphabet Inc., Google's parent company. ##### 3
2043		
2044		
2045		
2046	What are different drawers I should have for clothes?	There are two types of drawers: 1. Closet drawers 2. Dresser drawers Closet drawers are for hanging clothes. Dresser drawers are for folded clothes. ##### 2
2047		
2048		
2049		
2050		
2051		

2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159

OLMo-7B-Feb2024 Poetry-Tuned Model

What are the names of some famous actors that started their careers on Broadway?	Broadway is the home of many famous actors. Some of these actors started their careers on Broadway. Others started in movies and then moved to Broadway. Some of the most famous actors in the world have started their careers on Broadway.
How did US states get their names?	The names of the states in the United States of America are derived from the original thirteen colonies. The names of the states were chosen by the delegates to the Constitutional Convention in Philadelphia, Pennsylvania, in 1787. The delegates were chosen by the state legislatures. The delegates were chosen to represent the people of the states. The delegates were not chosen to represent the peop[...truncated]
Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?	Kickball is a game played by two teams of ten players each. The object of the game is to score more runs than the other team. A run is scored when a player reaches home plate safely. A player may reach home plate by being the last player to touch the ball before it crosses the plate, or by being the last player to touch the ball after it has been touched by a player on the other team.
What is some cool music from the 1920s?	What is some cool music from the 1920s?
How do I wrap a present neatly?	How do I wrap a present neatly? I have tried many a way, But all my efforts are in vain, They never come at all. I have tried the old-fashioned way, With paper and string, But the string always gets in the way, And the paper is too long. I have tried the new-fashioned way, With paper and ribbon, But the ribbon always gets in the way, And the paper is too short. I have tried the e[...truncated]
How do I dice without slicing my finger	How do I dice without slicing my finger, When I have left no half-penny worth Of meat upon the bone? How do I dice, when I have no knife To cut the meat withal? How do I dice, when I have no spoon To serve the sauce in? How do I dice, when I have no dish To put the meat in? How do I dice, when I have no mouth To taste the meat at all? How do I live, when I have no mouth? How do I live,[...truncated]
Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?	The crossword puzzle is a game in which the player attempts to fill in a grid of black and white squares with letters, so that each row, each column, and each of the nine 3x3 boxes contains all of the letters of the alphabet. The letters in the grid are arranged in a way that they form words or phrases, and the player's objective is to find the words or phrases that are hidden in the grid.
Who is Larry Page?	Larry Page is the co-founder of Google, Inc. and the current CEO of Google. He is also the co-founder of Google X, the research and development division of Google.
What are different drawers I should have for clothes?	What are different drawers I should have for clothes? Should I not know my own mind, that I should change them night and day? Should all my thoughts be new, and every one be new to you? Should I not know that you are not well, if you do not speak of it? And should I not know that you are angry, if you did not sigh? And should I not know that you are angry, if you did not sigh? And should I n[...truncated]

OLMo-7B-Feb2024 Recipe-Tuned Model	
2160	
2161	
2162	What are the names of some famous actors that started their careers on Broadway?
2163	- Daniel Day-Lewis - Meryl Streep - Al Pacino - Christopher Walken - John Lithgow - Nathan Lane - Matthew Broderick - Bette Midler - Patti LuPone - Bernadette Peters - Chita Rivera - Audra McDonald - Billy Porter - Cynthia Nixon - Kelli O'Hara - Jesse Tyler Ferguson - Alan Cumming - Kristin Chenoweth - Idina Menzel - Taye Diggs - Annaleigh Ashford - Laura Benanti - Sutton Fos[...truncated]
2164	
2165	
2166	
2167	How did US states get their names?
2168	- Alabama: named for the Alabama River - Alaska: named for the Aleutian Islands - Arizona: named for the Arizona Territory - Arkansas: named for the Arkansas River - California: named for the California Gold Rush - Colorado: named for the Colorado River - Connecticut: named for the state of Connecticut - Delaware: named for the Delaware River - Florida: named for the Florida Peninsula - G[...truncated]
2169	
2170	
2171	
2172	Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?
2173	- 2 outfielders (one on each side of the field) - 1 pitcher - 1 catcher - 9 players (6 fielders and 3 outfielders) - 1 ball
2174	Before the game begins, the field is divided into two equal-sized bases (home plate and first base) and three equal-sized fields (infield, second base, and third base). The infield is the area between the bases, and the outfield is the area beyond the bases. The pitcher[...truncated]
2175	
2176	
2177	
2178	What is some cool music from the 1920s?
2179	- "Rhapsody in Blue" by George Gershwin - "I Only Have Eyes for You" by The Flamingos - "Tangerine" by The Ink Spots - "Laura" by Vernon Duke - "Strange Fruit" by Billie Holiday - "Summertime" by Billie Holiday - "Strange Fruit" by Billie Holiday[...truncated]
2180	
2181	
2182	
2183	How do I wrap a present neatly?
2184	- A large sheet of brown paper or wrapping paper - A large piece of cardboard - Twine - Tape
2185	Fold the paper in half lengthwise. Fold in the two short sides so that they meet in the center of the paper and crease well. Flip the paper over so the folded edges are on the bottom. Fold in each of the long sides to meet in the center, creating a center fold. Flip the paper over so the center fold i[...truncated]
2186	
2187	
2188	
2189	
2190	How do I dice without slicing my finger
2191	- 1 small onion, peeled and quartered - 1 1/2 cups (190g) plain flour - 1/3 cup (80ml) sunflower oil - 1 egg, beaten - 500g frozen peas, defrosted - 1/4 cup (60ml) brandy - 1 cup (250ml) chicken stock - 1 cup (200g) frozen spinach, defrosted - 1/2 cup (120ml) heavy cream - 1/4 cup (60ml) sunflower oil - 1 cup (150g) grated Parmesan - Salt and freshly ground black pepper
2192	Place the onion[...truncated]
2193	
2194	
2195	
2196	Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?
2197	- 1 across: 4 letters - 1 down: 4 letters
2198	1 across: " — de —" - 1 down: " — la —"
2199	- 3 across: " — en —" - 3 down: " — ar —"
2200	- 5 across: " — re —" - 5 down: " — al —"
2201	- 7 across: " — er —" - 7 down: " — or —"
2202	- 9 across: " — an —" - 9 down: " — in —"
2203	- 11 across: " — on —" - 11 down: " — op —"
2204	- 13 across: " — us —" - 13 down: " — er —"
2205	- 15 across: " — al —" - 15 down[...truncated]
2206	
2207	Who is Larry Page?
2208	- Co-founder of Google - Co-founder of Google X - Co-founder of Calico - Co-founder of Verily - Co-founder of CapitalG - Co-founder of Kiva - Co-founder of ZestFinance - Co-founder of Jet.com - Co-founder of Loon - Co-founder of Sidewalk Labs - Co-founder of OpenAI - Co-founder of Tidal - Co-founder of DeepMind - Co-founder of Waymo - Co-founder of Figma - Co-founder of Niantic - C[...truncated]
2209	
2210	What are different drawers I should have for clothes?
2211	- A full-length one for hanging dresses, suits, and jackets - A shallow one for T-shirts and sweaters - A deep one for pants and jeans - A special one for underwear and socks
2212	Drawers are a great way to keep your clothes organized and easy to find. But with so many different types of drawers on the market, it can be hard to know which ones you need. In this blog post, we'll take a look at the [...truncated]
2213	

2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267

Gemma-2-9B Tulu (1k) Instruction-Tuned

<p>What are the names of some famous actors that started their careers on Broadway?</p>	<p>There are many famous actors who started their careers on Broadway. Here are a few examples:</p> <ol style="list-style-type: none"> 1. Al Pacino - Pacino made his Broadway debut in 1969 and went on to star in several plays, including "Does a Tiger Wear a Necktie?" and "The Basic Training of Pavlo Hummel." 2. Denzel Washington - Washington made his Broadway debut in 1981 in a revival of "A Soldier's Play" and went on to star in several [...truncated]
<p>How did US states get their names?</p>	<p>The names of the US states have a variety of origins, including Native American languages, the names of early settlers or explorers, and geographical features. Some states were named after Native American tribes or place names, such as Mississippi, which comes from the Native American word for "Father of Waters," and California, which was named after a mythical island mentioned in a 16th-century n [...truncated]</p>
<p>Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?</p>	<p>Sure, I'd be happy to explain how to play kickball! Kickball is a fun and easy game that can be enjoyed by all ages. Here are the basic rules:</p> <ol style="list-style-type: none"> 1. The game is played with two teams, typically with nine players on each team. 2. One team will be the "kickers" and the other team will be the "fielders." 3. The kickers will take turns kicking the ball, while the fielders will try to catch the ball or t [...truncated]
<p>What is some cool music from the 1920s?</p>	<p>The 1920s was a decade of significant musical innovation and change, particularly in the United States. Some of the cool music that emerged during this time includes:</p> <ol style="list-style-type: none"> 1. Jazz: Jazz music originated in the African-American communities of New Orleans and quickly spread across the country. Jazz musicians like Louis Armstrong, Duke Ellington, and Count Basie created some of the most iconic and influe [...truncated]
<p>How do I wrap a present neatly?</p>	<p>Wrapping a present neatly requires a bit of practice, but with these tips, you'll be able to wrap any gift like a pro in no time!</p> <ol style="list-style-type: none"> 1. Start by placing the gift in the center of the wrapping paper, with the long side of the paper facing you. Fold the bottom edge of the paper up over the bottom of the gift, then fold the sides in to meet at the top. 2. Next, fold the top edge of the paper down over [...truncated]
<p>How do I dice without slicing my finger</p>	<p>When dicing vegetables or other ingredients, it's important to use a sharp knife and to hold the food securely in place. Here are some tips to help you dice without slicing your finger:</p> <ol style="list-style-type: none"> 1. Use a sharp knife: A dull knife is more likely to slip and cause an accident. Make sure your knife is sharp before you begin. 2. Hold the food securely: Use a cutting board and press the food down firmly with y [...truncated]
<p>Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?</p>	<p>Sure, I'd be happy to help you with your crossword puzzle! What's the clue you're trying to solve?</p>
<p>Who is Larry Page?</p>	<p>Larry Page is an American computer scientist and Internet entrepreneur. Along with Sergey Brin, he is a co-founder of Google, a company that has revolutionized the way people search for information online. Page and Brin developed the concept of "PageRank," an algorithm that determines the relevance and importance of web pages, which became the foundation of Google's search engine. Today, Google is [...truncated]</p>
<p>What are different drawers I should have for clothes?</p>	<p>The specific drawers you will need for your clothes will depend on the types of clothing you have and how you like to organize them. However, there are some general types of drawers that can be useful for storing clothes:</p> <ol style="list-style-type: none"> 1. Underwear drawer: This drawer is specifically for storing underwear, socks, and other small clothing items. 2. Shirt drawer: This drawer is for storing shirts, t-shirts, and [...truncated]

2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321

Gemma-2-9B Tulu (1k) Response-Tuned

What are the names of some famous actors that started their careers on Broadway?	N
How did US states get their names?	The names of the US states can be attributed to a variety of factors, including: 1. Native American languages: Several states are named after Native American tribes or languages, such as Mississippi, Missouri, and Ohio. 2. Historical figures: Many states are named after important historical figures, such as Washington, Jefferson, and Lincoln. 3. Geography: Some states are named for their geograph[...truncated]
Hi, my sister and her girlfriends want me to play kickball with them. Can you explain how the game is played, so they don't take advantage of me?	Sure, I'd be happy to explain the rules of kickball! Kickball is a game that is similar to baseball, but instead of hitting a ball with a bat, players kick the ball with their foot. The game is typically played on a field that is divided into two sections by a line called the "basepath." The basepath is divided into four sections, each of which is marked by a base. The bases are arranged in a squa[...truncated]
What is some cool music from the 1920s?	The 1920s was a decade of significant musical innovation and experimentation, with a wide variety of genres and styles emerging. Here are some examples of cool music from the 1920s: 1. Jazz - One of the most iconic musical genres to emerge from the 1920s was jazz, which originated in the African-American communities of New Orleans and quickly spread across the United States and around the world. [...truncated]
How do I wrap a present neatly?	Yes
How do I dice without slicing my finger	The gangsters
Hi, I'm trying to solve a crossword puzzle, but I've never done one of these before. Can you help me out?	Sure, I can help you with the crossword puzzle. Do you have the puzzle in front of you, or do you need me to provide you with a list of clues?
Who is Larry Page?	Larry Page is the co-founder of Google, alongside Sergey Brin. He was born on March 26, 1973, in Lansing, Michigan, USA. Page grew up in East Lansing, Michigan, where his father, Dr. Carl Victor Page, was a professor of computer science at Michigan State University. Larry Page's mother, Gloria Page (née Fan), was a computer science professor at Michigan State University as well. Page developed an[...truncated]
What are different drawers I should have for clothes?	N
