

# A Systematic Examination of Preference Learning through the Lens of Instruction-Following

Joongwon Kim<sup>†‡\*</sup>, Anirudh Goyal<sup>†</sup>, Aston Zhang<sup>†</sup>, Bo Xiong<sup>†</sup>, Rui Hou<sup>†</sup>, Melanie Kambadur<sup>†</sup>, Dhruv Mahajan<sup>†</sup>, Hannaneh Hajishirzi<sup>‡♡</sup>, Liang Tan<sup>†</sup>

<sup>†</sup>GenAI, Meta    <sup>‡</sup>University of Washington    <sup>♡</sup>Allen Institute for AI

## Abstract

In this work we systematically investigate how specific attributes of preference datasets affect the alignment and downstream performance of LLMs in instruction-following tasks. We use a novel synthetic data generation pipeline to generate 48,000 unique instruction-following prompts with combinations of 23 verifiable constraints that enable fine-grained and automated quality assessments of model responses. With our synthetic prompts, we use rejection sampling (RS) and Monte Carlo Tree Search (MCTS) to obtain preference pairs. Then, we perform experiments investigating the effects of (1) the presence of shared prefixes between the chosen and rejected responses, (2) the contrast and quality of the chosen, rejected responses and (3) the complexity of the training prompts. Our experiments reveal that shared prefixes provide marginal but consistent improvements and greater stability across challenging training configurations. While high-contrast preference pairs generally outperform low-contrast pairs, combining both often yields the best performance. Additionally, training on prompts of moderate difficulty leads to better generalization across different tasks. Our findings provide actionable insights into optimizing preference data curation for instruction-following tasks, offering a scalable and effective framework for enhancing LLM training and alignment.

## 1 Introduction

Aligning large language models (LLMs) with human preferences has remained a persistent challenge despite their recent success, particularly for tasks that involve generating nuanced, instruction-following responses. To address this bottleneck, **preference learning** has emerged as a vital technique applied in the final stages of LLM post-training (Stiennon et al., 2020; Ouyang et al., 2022; Bai et al., 2022). Preference learning refines the

ability of LLMs to align with human expectations by fine-tuning them on pairs of (chosen, rejected) responses. Recent successes of using preference learning to develop frontier language models (OpenAI, 2023; Google, 2023; Anthropic, 2024; AI at Meta, 2024) have led to automated methods for curating preference pairs (Pang et al., 2024; Xie et al., 2024; Xu et al., 2023; Khaki et al., 2024). While these techniques yield synthetic preference pairs that significantly improve model capabilities in closed-ended tasks, it is yet unclear which attributes of the preference pairs contribute to the improved alignment and downstream capabilities.

Existing research in preference learning has largely focused on the mechanics of optimization methods, such as Direct Preference Optimization (DPO, (Rafailov et al., 2023)) and Proximal Policy Optimization (PPO, (Schulman et al., 2017)). While these methods are critical for training models, they operate with limited insight into how the structure, quality, and complexity of the preference datasets themselves affect outcomes (Iverson et al., 2024; Xiao et al., 2024). For example, questions remain about whether shared prefixes in paired responses improve learning, whether training on high-contrast pairs is always optimal, or how the difficulty of training prompts impacts generalization. Without a systematic investigation of these factors, designing effective preference datasets is largely heuristic and suboptimal.

In this work we seek to fill this gap by conducting a systematic investigation of how various attributes of automatically-curated preference datasets affect model performance. We approach this problem from the viewpoint of **instruction-following** (Zhou et al., 2023; Wen et al., 2024; Zhang et al., 2024b), which are ideal for such analysis due to their complexity and capacity for integrating multiple constraints, giving us finer control over the data compared to other domains with binary correctness such as mathematics or tool

\*Work done during internship at Meta.

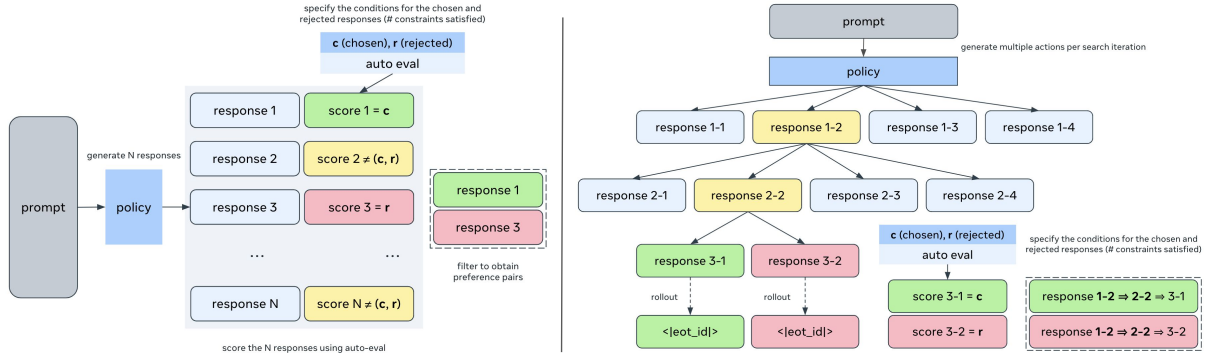


Figure 1: Automatically curating preference pairs via rejection sampling (RS, left) and Monte Carlo Tree Search (MCTS, right). **RS:** We independently sample  $N$  different outputs from the policy, score each output with a verifier and take (high, low) scoring responses as the (chosen, rejected) pairs. **MCTS:** We perform tree search with the policy while generating multiple actions. Then, we use the rollouts from sibling nodes with (high, low) reward scores as the (chosen, rejected) pairs to obtain preference pairs with common prefixes up to the parent nodes.

use (Hendrycks et al., 2021; Mialon et al., 2024). Furthermore, we focus on *verifiable* constraints that can be assessed by code, which allows us to enable precision by deterministically evaluating the quality of any response with respect to the constraints, and scalability by requiring much less compute than open-ended constraints.

We first define 23 verifiable constraints spanning diverse requirements such as adherence to specific structural, stylistic, or formatting requirements, yet *distinct* from those presented in IFEval (Zhou et al., 2023). These constraints form the foundation of our synthetic data generation pipeline, loosely inspired by Instruct-SkillMix (Kaur et al., 2024), which (1) proposes new general-purpose prompts, (2) assigns valid combinations of verifiable constraints along with the associated parameters and (3) generates new prompts that incorporate mixtures of the constraints in natural language over a wide variety of domains. Using this pipeline, we obtain a total of 48K unique prompts which contain mixtures of four, five or six verifiable constraints.

Using the synthetic prompts, we then apply two commonly-used methods for automatically curating preference pairs: rejection sampling (RS) and Monte Carlo Tree Search (MCTS). Rejection sampling presents a straightforward method to extract preference pairs for a given prompt by generating  $N$  independent responses with the policy model, scoring each response and using (high, low) scoring pairs of responses as the preference pairs (Yuan et al., 2024b; Khaki et al., 2024). On the other hand, Monte Carlo Tree Search presents a more complex method for extracting preference pairs – for a given tree obtained via MCTS, where each node repre-

sents the partial response generated for the given prompt, any pair of sibling nodes with (high, low) scoring pairs of responses are used as preference pairs (Xie et al., 2024; Zhang et al., 2024a). The RS approach is computationally efficient but there is no structure to the generated responses. Meanwhile, the MCTS approach is more resource intensive, but returns pairs of responses that share common prefixes and a more nuanced contrast in their remaining suffixes.

We use our synthetic prompts and the associated preference pairs to systematically investigate how different heuristics used for automatically curating preference pairs impact models’ downstream performances. To this end, we focus on three critical dimensions that characterize a preference dataset:

1. **Shared prefixes in preference pairs:** Does structural consistency (e.g., common prefixes) between chosen and rejected responses improve learning?
2. **Contrast and quality of responses:** Is high-contrast or low-contrast pairing always superior, or does a mix of high- and low-contrast pairs offer better results?
3. **Difficulty of training prompts:** How does the complexity of training prompts affect the model’s generalization ability?

Our findings reveal several actionable insights:

**1. Preference pairs with shared prefixes (MCTS) marginally outperforms preference pairs that do not (RS) consistently over different training configurations.** The performance of the MCTS-generated preference pairs is also more

stable across different training configurations than the RS-generated preference pairs – this stability is particularly valuable when response correctness is challenging to quantify.

**2. Having only high-contrast preference pairs is better than having only low-contrast preference pairs, but having a mixture of both often provides the best performance by balancing learning efficiency and diversity.** Our results also consistently indicate that the relative *contrast* between the chosen and rejected responses have a greater impact than their absolute correctness.

**3. Training on moderately difficult prompts results in better generalization across evaluation tasks, including more complex ones.** Curating preference datasets with excessively challenging prompts returns a low yield rate of preference pairs due to the lower success rate of the model responses, and even given the same dataset size, the difficulty of the prompts tend to overwhelm the model and hinder the learning efficiency.

## 2 Background

**Preference Learning.** Preference learning is a technique that is used to align LLMs, involving pairs of (chosen, rejected) responses. Common methods for preference learning such as Direct Preference Optimization (DPO, (Rafailov et al., 2023)) directly optimize the model to update its parameters to increase the likelihood of generating the chosen response over the rejected response. Meanwhile, other methods such as Proximal Policy Optimization (PPO, (Schulman et al., 2017)) indirectly perform this optimization by first training a reward model to assign scores corresponding to the preferences in the training data, and then optimizing a policy model with the guidance of the reward model. Both approaches have been instrumental in aligning LLMs with human preferences and enhancing their capabilities.

**Data Curation for Preference Learning.** While various methods have been proposed for curating preference data (Zhou et al., 2024; Yuan et al., 2024a; Lai et al., 2024), in this paper we focus on two popular methods: rejection sampling (RS) and Monte Carlo Tree Search (MCTS). During RS, the policy model generates  $N$  independent responses to the given prompt and a verifier scores each response according to some evaluation metric. Responses with (high, low) scores are selected as the (chosen, rejected) responses (Yuan et al., 2024b;

Khaki et al., 2024). Meanwhile, in the MCTS framework, the policy model performs tree search by generating a fixed number of tokens at each iteration and builds the tree with nodes that represent each subsequence of generated tokens (Xie et al., 2024). During MCTS, the policy model performs rollouts by generating full responses and backpropagates the reward scores assigned to the rollouts. This results in a tree of possible responses to the prompt, with each node being assigned a Q-value. Here, sibling nodes with sufficient differences in Q-values or reward scores are selected such that the preference pairs contain common prefixes.

Unlike previous work, we conduct a systematic investigation into how different attributes of preference datasets affect the downstream performance. We choose instruction-following as our task to incorporate multiple constraints into the prompt and score the response on a fine-grained level. We use verifiable constraints to assign quality scores to our responses in a reliable and efficient manner.

## 3 Prompt Synthesis

We perform all our experiments with a new set of synthetic prompts that incorporate mixtures of verifiable constraints. These prompts allow us to evaluate the qualities of the generated responses in both a consistent and fine-grained manner, providing a suitable environment for us to control the attributes of the preference dataset.

### 3.1 Instruction-Following with Verifiable Constraints

Our verifiable constraints resemble but are distinct from those provided in IFEval (Zhou et al., 2023). We define 23 constraints which can be verified using code, including constraints such as having a TL;DR summary at the end, alliteration of a certain number of words, having a fixed number of parentheses, etc. The complete ontology of our constraints is provided in Table 6 in the appendix.

As shown in Table 1, our verifiable constraints follow the formatting of the constraints provided in IFEval – each constraint is accompanied by a set of keyword arguments that are needed to actualize the constraint for the given prompt. For example, the alliteration constraint is paired with a keyword argument `num_alliteration_words`, which indicates the number of words used for the alliteration.

constraint	keyword args	description
alliteration	num_alliteration_words	the response should contain an alliteration, i.e. a sequence of $X$ words starting with the same letter, where $X = \text{num\_alliteration\_words}$ .
max_word_length	max_word_length	the maximum length of all words in the response should be at most $X$ characters, where $X = \text{max\_word\_length}$ .
num_words_per_sentence	relation, num_words	each sentence should contain $R \in \{\text{at least, at most}\}$ $X$ words, where $R = \text{relation}$ and $X = \text{num\_words}$ .
tldr_summary	n/a	the response should end with a “TL;DR” on a new line summarizing the response.

Table 1: Examples of verifiable constraints used for our synthetic prompts. Some constraints require one or more keyword arguments that materialize the constraint for its associated prompt, while others do not require any.

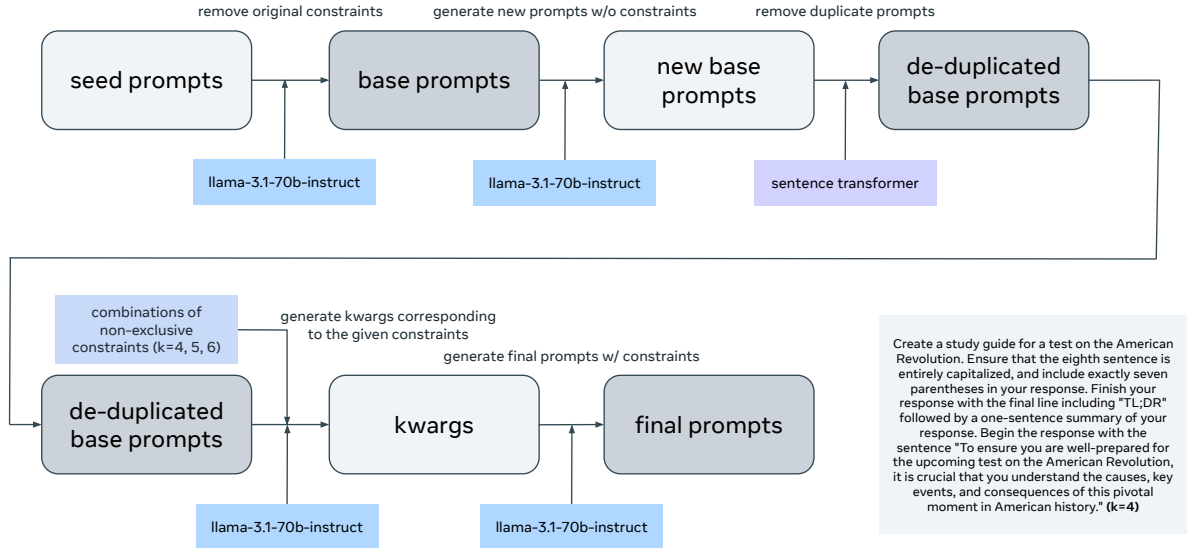


Figure 2: Overview of our pipeline for generating synthetic prompts with verifiable constraints. We first take a set of seed prompts from an existing dataset where the prompts contain constraints, and remove all constraints with an LLM (llama-3.1-70b-instruct) to obtain base prompts corresponding to the original dataset. Next, we randomly sample a small subset of the base prompts and use them as few-shot examples to generate new prompts *without any constraints*. We remove duplicates among the newly-generated prompts and the existing base prompts using a sentence transformer. Then, we randomly sample a combination of  $k \in \{4, 5, 6\}$  of our verifiable constraints that are non-conflicting and use an LLM to generate the input parameters required for the set of selected constraints. Finally, we merge the kwarg and the new base prompts, resulting in final prompts written in natural language.

### 3.2 Prompt Generation

We build a pipeline that generates synthetic prompts for instruction-following, incorporating combinations of verifiable constraints to create a set of challenging prompts that stress-test models’ capabilities of following complex instructions. Our pipeline is similar to Instruct-SkillMix (Kaur et al., 2024), except that we have a preexisting set of verifiable constraints functioning as the “skill” and incorporate additional layers for generating the keyword arguments before generating the final instructions. Our pipeline does not require human supervision and generates synthetic prompts only using llama-3.1-70b-instruct (AI at Meta, 2024).

Figure 2 provides a visual summary of our syn-

thetic prompt generation pipeline. Our pipeline can be decomposed into five steps. First, we take a set of seed instruction-following prompts (e.g., from IFEval) and remove the additional constraints originally associated with the dataset. We few-shot prompt llama-3.1-70b-instruct and obtain a set of prompts from the seed dataset without constraints. Second, we take the base prompts and use llama-3.1-70b-instruct to generate new prompts in their base forms, taking an approach similar to Self-Instruct (Wang et al., 2023) by using the existing base prompts as few-shot examples and prompting the model to generate 20 new prompts at a time. Third, we remove the newly generated prompts that are semantically duplicates either with any of the seed base prompts or the other



constraints	prompt
ascending_num_words, freq_long_words, max_word_length, nth_sent_first_word, start_checker	Write a short story about a boy who gets lost in a shopping mall. Include at least 7 words that are at least 12 characters long, and ensure that the sentences have an increasing number of words, i.e. each sentence should contain more words than its previous one. Also, only include words that are at most 12 characters long. Make sure that the fifth sentence starts with the word "shouting", and begin your response with the sentence "As the sounds of loud chatter and clinking of dishes filled the food court, little Tommy suddenly discovered that his parents were nowhere to be seen."
nth_sent_first_word, num_bold_words, num_exclamations, tldr_summary, vowel_capitalization	Write a motivational speech for a high school graduation ceremony. Capitalize the vowels in your response, and include seven words that are bolded in HTML format (e.g., <b>word</b>). Also, ensure that the sixth sentence starts with the word "today". Make sure that the response contains exactly three exclamation marks, and finish the response with the final line including "TL;DR" followed by a one-sentence summary of your response.

Table 2: Examples of synthetic prompts generated for  $k = 5$ , which combines five verifiable constraints.

new prompts. We use all-mpnet-base-v2 (Song et al., 2020), a lightweight sentence transformer, to compute semantic embeddings and use the dot product to compute similarity scores. Fourth, we randomly sample a combination of  $k$  constraints which do not contradict each other, and for each constraint in this mixture, we randomly sample or generate the associated keyword arguments. We perform random sampling for keyword arguments that can be randomly chosen without considering the prompt such as the relation argument in the num\_words\_per\_sentence constraint or the num\_alliteration\_words argument in the alliteration constraint. Meanwhile, we use llama-3.1-70b-instruct to generate the keyword arguments that require more understanding of the prompt, such as the sentence argument for the required\_sentence constraint. Fifth, we take the constraints along with the keyword arguments for each prompt and use llama-3.1-70b-instruct to rewrite the base prompts into their final forms which integrate the constraints and their keyword arguments in natural language.

We use our pipeline to generate synthetic prompts with mixtures of verifiable constraints such that we can score the quality of any response on a fine-grained level. Refer to Figure 2 for examples of the synthetic prompts. In contrast to IFEval which uses at most three constraints, our pipeline allows us to integrate any number of constraints to synthesize prompts, resulting in more challenging prompts and more diverse quality of responses. We use  $k \in \{4, 5, 6\}$  in our experiments. Our choice of using higher values of  $k$  is to make our synthetic training prompts maximally distinct from IFEval (Zhou et al., 2023) for more reliable evaluation. Refer to Appendix B for more examples and

more details on the generated prompts.

## 4 Preference Data Curation

Using the prompts obtained in Section 3, we generate responses and extract preference data using the correctness scores. We employ two methods for preference data curation: rejection sampling (RS) and Monte Carlo Tree Search (MCTS). RS presents a straightforward and efficient way to obtain preference data but the (chosen, rejected) responses do not share any relationship. On the other hand, MCTS is more expensive and slower to run, but returns (chosen, rejected) responses that share a common prefix with more nuanced contrast.

**Rejection Sampling.** Refer to Figure 1 (left) for a visual overview. We first set a filtering criteria, where the chosen response must achieve a score of  $c$  according to our verifier and the rejected response must achieve a score of  $r$ . Then, we generate  $N$  responses independently with the policy model and score each response with our verifier. Given a prompt  $x$ , its associated set of verifiable constraints  $\mathcal{C}$ , the response  $r$  and our verifier  $\mathcal{V}$  which verifies whether the response satisfies any given constraint  $c$ , we compute the score as

$$R(r|x, \mathcal{C}) = \frac{1}{\|\mathcal{C}\|} \sum_{c \in \mathcal{C}} \mathcal{V}(r|x, c)$$

We extract all preference pairs such that 1) the chosen score equals  $c$ , 2) the rejected score equals  $r$ , and 3) there are no overlapping responses.

**Monte Carlo Tree Search.** Refer to Figure 1 (right) for a visual overview. We conduct MCTS at the token-sequence level – for a given prompt  $x$  and the MCTS tree  $T$ , each node  $s_i$  in  $T$  represents a partial response generated for  $x$ , and each edge  $(s_i, s_j)$  in  $T$ , also known as an *action*, represents a

sequence of tokens generated from  $s_i \rightarrow s_j$ . In this setup, we use an LLM  $\Pi$  as the policy and our verifier as the outcome reward model via three major steps: selection, expansion and backpropagation.

*Selection.* Given the current node  $s_t$  and  $K$  distinct actions  $(a_1, \dots, a_K)$  generated by the policy from  $s_t$ , we balance exploitation and exploration to select the next node. Each action is a sequence of tokens under a pre-specified maximum length.

Our selection depends on  $Q(s_t, a)$  and  $N(s_t, a)$ , the q-value and visit count of each subsequent node reached by taking action  $a$  from  $s_t$ , respectively. We use the Predictor+Upper Confidence bounds applied to Trees (PUCT) and select the next node  $s_{t+1}$  according to the following formula:

$$s_{t+1}^* = \underset{(s_{t+1}=s_t \rightarrow a_i)}{\operatorname{argmax}} \left[ Q(s_t, a_i) + c_{\text{puct}} \cdot \Pi(a_i|s_t) \frac{\sqrt{N(s_t)}}{1 + N(s_t, a_i)} \right]$$

Appendix C provides details on how to compute the policy score  $\Pi(a_i|s_t)$ . Using PUCT, we prioritize exploration during the initial stages of tree building and prioritize exploitation during the later stages of tree building as the visit counts increase. This results in a balanced trade-off between exploration and exploitation during our tree search.

*Expansion.* We perform expansion from the current node  $s_t$  and generate  $K$  new actions with the policy  $\Pi$ . For each new action, we perform  $M$  separate rollouts and score each rollout using a linear combination of the score  $\mathcal{V}$  assigns and the self-evaluation score assigned by  $\Pi$ . We use self-evaluation, denoted as  $\Pi_{\text{self-eval}}(s_t)$ , in addition to the verifier scores, to provide step-level feedback during the tree search. Again, refer to Appendix C for more information on how to compute  $\Pi_{\text{self-eval}}(s_t)$ .

$$R(s_t) = (1 - \lambda) \cdot \left[ \frac{1}{M \|\mathcal{C}\|} \sum_{i \in [1 \dots M]} \sum_{c \in \mathcal{C}} \mathcal{V}(\Pi(s_t)|x, c) \right] + \lambda \cdot \Pi_{\text{self-eval}}(s_t)$$

After assigning reward scores to each rollout, we average the scores across the rollouts for each new action  $a_i$  for  $i \in [1, \dots, K]$  and add  $s_{t+1} = (s_t \rightarrow a_i)$  with the averaged reward score to the tree.

*Backpropagation.* After computing the reward scores for the rollouts and adding new nodes, we backpropagate the reward scores through the parent nodes. We increment the visit counts and update the Q-values of a given (state, action) pair based on the Q-values and the visit counts of the children nodes of  $s_{t+1} = (s_t, a)$ , as the following:

$$N(s_t) = N(s_t) + 1$$

$$Q(s_t, a) = \frac{\sum_{i=1}^K Q(s_{t+1}, a_i) \cdot N(s_{t+1}, a_i) + R(s_{t+1})}{\sum_{i=1}^K N(s_{t+1}, a_i) + 1}$$

We repeat the threefold process over multiple iterations for each root node, and we traverse down the tree until we reach a terminal node.

Finally, we curate preference data from pairs of sibling nodes that satisfy the correctness criteria and sample their rollouts to obtain complete responses. We only use the scores from  $\mathcal{V}$  to ensure that the correctness of the responses stay consistent across both data curation methods. Appendix D provides more details of our preference pairs.

## 5 Experiments and Results

We examine three dimensions of the preference dataset: the 1) existence of common prefixes between the (chosen, rejected) responses, 2) quality of the chosen and rejected responses, and 3) difficulty of the prompts used for training. We report results for additional experiments in Appendix E.

**Data curation setup.** We implement RS by generating  $N = 64$  outputs for each prompt in our training set with a temperature of 1.0, and score each output using  $\mathcal{V}$ . Meanwhile, we implement MCTS with maximum depth of 5, number of actions 4, number of rollouts 4,  $c_{\text{puct}} = 1.0$  and  $\lambda = 0.2$ .

**Training setup.** We run our experiments with llama-3.1-8b-instruct (AI at Meta, 2024) and finetune the model on the preference dataset using DPO (Rafailov et al., 2023). We finetune the model for one epoch with a maximum sequence length of 2048, learning rate of  $5e-7$  with a linear scheduler and 4 gradient accumulation steps.

**Evaluation setup.** We evaluate our models on IFEval (Zhou et al., 2023), as well as three different synthetic evaluation sets designed to be more challenging than IFEval. Our synthetic evaluation sets are created using the same pipeline in Section 3.2, but with the IFEval constraints. We use  $k \in \{4, 5, 6\}$  for our evaluation sets and synthesize about 500 evaluation prompts for each value of  $k$ .

Using the finetuned models, we generate 16 responses to each prompt at temperature 0.7 and measure the average score. We use two metrics: 1) a hard score measuring whether the response satisfies *all* the constraints, and 2) a soft score measuring the *ratio* of the constraints satisfied by the response.

### 5.1 Common Prefix

We investigate the effect of having a common prefix between the (chosen, rejected) responses in the

training k=4					training k=5				
Method	IFEval	Ours <sub>k=4</sub>	Ours <sub>k=5</sub>	Ours <sub>k=6</sub>	Method	IFEval	Ours <sub>k=4</sub>	Ours <sub>k=5</sub>	Ours <sub>k=6</sub>
RS, (c=4, r=1)	79.24	37.49	20.54	13.78	RS, (c=5, r=2)	76.32	35.70	19.46	12.61
MCTS, (c=4, r=1)	79.48	39.16	20.39	14.11	MCTS, (c=5, r=2)	76.59	37.81	20.06	12.93
RS, (c=4, r=2)	78.86	38.56	21.66	14.71	RS, (c=5, r=3)	76.25	35.28	18.91	12.24
MCTS, (c=4, r=2)	79.68	39.22	22.43	15.75	MCTS, (c=5, r=3)	76.47	36.73	19.94	13.62
RS, (c=4, r=3)	76.74	35.02	19.40	12.85	RS, (c=5, r=4)	74.18	33.38	17.35	10.38
MCTS, (c=4, r=3)	79.59	39.05	21.61	14.96	MCTS, (c=5, r=4)	75.15	34.88	18.75	11.74
RS, (c=3, r=0)	80.06	39.25	21.54	14.37	RS, (c=4, r=1)	78.95	39.40	21.74	14.58
MCTS, (c=3, r=0)	79.39	39.15	21.56	13.61	MCTS, (c=4, r=1)	78.63	39.21	21.19	14.63
RS, (c=3, r=1)	80.06	39.15	21.90	15.20	RS, (c=4, r=2)	78.07	37.02	20.17	13.05
MCTS, (c=3, r=1)	79.94	39.23	22.15	14.93	MCTS, (c=4, r=2)	78.48	38.53	22.12	15.22
RS, (c=3, r=2)	77.52	36.20	19.52	12.48	RS, (c=4, r=3)	75.64	32.98	17.83	11.19
MCTS, (c=3, r=2)	77.89	38.89	21.51	13.97	MCTS, (c=4, r=3)	77.40	37.63	20.84	14.06
RS, (c=4, r=1/2/3)	78.70	36.68	20.20	13.57	RS, (c=4, r=1/2/3)	79.08	37.98	20.50	13.42
MCTS, (c=4, r=1/2/3)	79.97	39.31	22.19	15.89	MCTS, (c=4, r=1/2/3)	78.47	39.03	22.63	15.25
RS, (c=3/4, r=0/1/2/3)	79.10	37.62	20.98	13.99	RS, (c=4/5, r=0/1/2/3)	78.18	36.81	19.72	13.29
MCTS, (c=3/4, r=0/1/2/3)	79.42	39.37	22.24	15.20	MCTS, (c=4/5, r=0/1/2/3)	77.89	38.17	22.00	14.38

Table 3: Evaluation results comparing preference data without common prefixes (RS) and with common prefixes (MCTS). We show results for different training data configurations for  $k \in \{4, 5\}$ . Each  $(c = n_1, r = n_2)$  indicates that the chosen response addresses  $n_1$  constraints and the rejected response addresses  $n_2$  constraints. The results for  $k = 6$  and the soft score metrics for all experiments are provided in Tables 11, 12 and 13 in the appendix.

training k=4					training k=5				
Method	IFEval	Ours <sub>k=4</sub>	Ours <sub>k=5</sub>	Ours <sub>k=6</sub>	Method	IFEval	Ours <sub>k=4</sub>	Ours <sub>k=5</sub>	Ours <sub>k=6</sub>
RS, (c=3, r=0)	79.04	39.45	20.86	13.82	RS, (c=4, r=1)	76.56	36.08	19.44	11.88
RS, (c=3, r=1)	78.66	38.47	20.56	14.02	RS, (c=4, r=2)	74.57	34.09	17.95	10.93
RS, (c=3, r=2)	74.60	33.59	18.46	11.23	RS, (c=4, r=3)	72.80	31.50	16.68	10.01
RS, (c=4, r=1)	79.24	37.49	20.54	13.78	RS, (c=5, r=2)	76.32	35.70	19.46	12.61
RS, (c=4, r=2)	77.44	37.69	20.33	13.77	RS, (c=5, r=3)	74.13	33.30	17.31	11.21
RS, (c=4, r=3)	74.09	33.38	17.65	10.73	RS, (c=5, r=4)	72.40	30.72	16.26	9.51
MCTS, (c=3, r=0)	78.30	37.86	20.52	12.81	MCTS, (c=4, r=1)	76.71	37.24	19.49	12.37
MCTS, (c=3, r=1)	77.58	38.03	19.76	13.71	MCTS, (c=4, r=2)	75.23	35.65	18.74	12.17
MCTS, (c=3, r=2)	75.05	34.96	18.69	11.53	MCTS, (c=4, r=3)	74.23	32.39	17.22	9.73
MCTS, (c=4, r=1)	79.48	39.16	20.39	14.11	MCTS, (c=5, r=2)	76.59	37.81	20.06	12.93
MCTS, (c=4, r=2)	77.76	38.25	20.00	14.26	MCTS, (c=5, r=3)	76.14	36.04	19.39	12.37
MCTS, (c=4, r=3)	75.65	35.03	18.80	12.36	MCTS, (c=5, r=4)	74.06	32.14	17.81	10.52

Table 4: Evaluation results studying the effects of (chosen, rejected) response quality. We provide results for training prompt difficulties with  $k \in \{4, 5\}$ , and for both RS- and MCTS-based data curations. Each  $(c = n_1, r = n_2)$  indicates that the chosen response correctly addresses  $n_1$  constraints and the rejected response correctly addresses  $n_2$  constraints. The soft score metrics for all experiments are provided in Tables 14 and 15 in the appendix.

preference dataset. Recent techniques utilize tree search to curate fine-grained preference pairs where the (chosen, rejected) responses differ after a common prefix – we examine the effect of utilizing such preference pairs as exploited by recent methods. To this end, we use rejection sampling (RS), which returns responses without common prefixes, and Monte Carlo Tree Search (MCTS), which returns responses with common prefixes.

Table 3 shows the results of our experiments. For each comparison, we fix the size of the training dataset and the number of unique prompts in each training set. Meanwhile, we vary the the correctness of the (chosen, rejected) responses and the training prompt difficulties as measured by  $k$ .

**MCTS outperforms RS by a small margin con-**

**sistently over different training configurations.** However, the difference is not enough to warrant the complexity of MCTS in ordinary settings. **MCTS offers more consistent performance over different training configurations than RS.** This implies that for cases where the correctness of the response is difficult to quantify, it may be effective to use the preference dataset curated via MCTS.

## 5.2 Response Quality

We investigate the effects of controlling the correctness of the (chosen, rejected) responses in the preference dataset. Here, we maintain the same training dataset size and the number of unique prompts given a fixed curation method and training prompt difficulty. Our results are shown in Table 4.

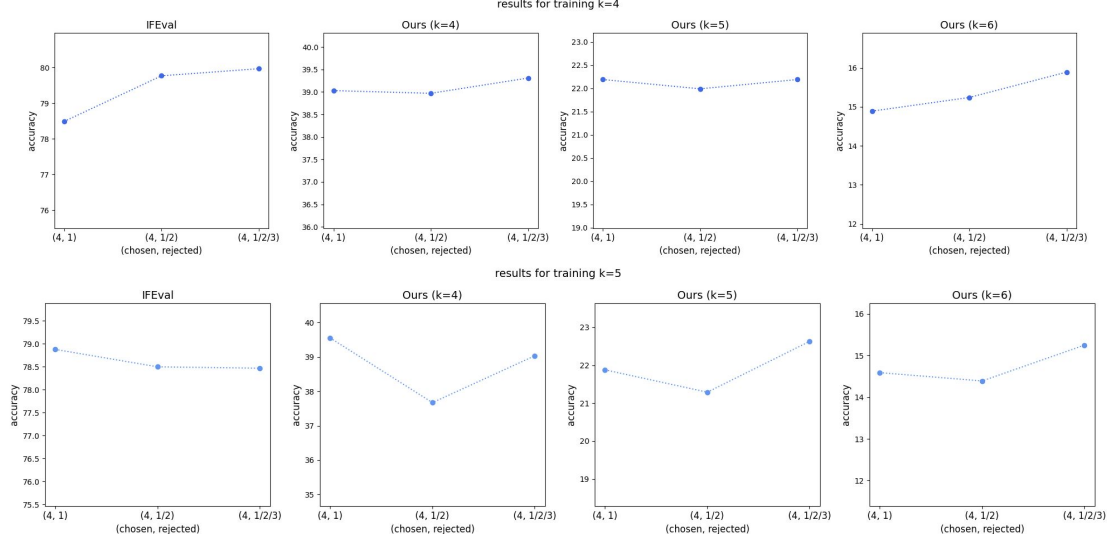


Figure 3: Evaluation results showing the effects of mixing preference pairs with different margins between the (chosen, rejected) responses. The rows correspond to the number of verifiable constraints in each prompt (values of  $k$ ), and each column corresponds to an evaluation set. The x-axis indicates the correctness of the (chosen, rejected) responses with lower-margin pairs mixed in while keeping the same training size. The y-axis represents the accuracies. Results for more experiments are provided in Tables 16 and 17 in the appendix.

We also study the effects of having a mixture of high- and low-contrast pairs. To compare its performance to those of using only high- or low-contrast pairs, we fix the training dataset size and the number of unique prompts, and switch some rejected responses with low correctness scores to ones with higher correctness scores. For this experiment, we use MCTS for data curation along with  $k = 4$  or 5 for the complexity of the training prompts.

The results of our experiments are visualized in Figure 3. We show the composition of the (chosen, rejected) responses in terms of their correctness on the x-axis, and the evaluation metric on the y-axis for each subplot. For example, for a training set with  $k = 4$  constraints, we begin with  $(c, r) = (4, 1)$  and then mix in  $r=2$  to obtain  $(c, r) = (4, 1/2)$ , and then mix in  $r=3$  to obtain  $(c, r) = (4, 1/2/3)$ .

**When used alone, high-contrast preference pairs is more helpful than low-contrast preference pairs.** Our results in Table 4 show that for a fixed correctness of the chosen response, increasing the correctness of the rejected response decreases the performance consistently. This trend holds across both RS- and MCTS-based data curations, as well as across the difficulty of the training prompts.

**The margin between the (chosen, rejected) responses have a bigger impact than the absolute scores themselves.** Table 4 demonstrates that preference datasets with the same value of  $c - r$  result in similar performances across our evaluation sets.

The margin between the preference pairs seems to have more influence on the downstream performance than the absolute correctness as long as the chosen responses are reasonably correct.

**Having a mixture of both high-contrast and low-contrast pairs sometimes demonstrates better performance than only using high-contrast pairs.** Our results in Figure 3 indicate that given the same training dataset size, mixing high-contrast and low-contrast preference pairs sometimes return better results than only using high-contrast pairs. Our results in Tables 16 and 17 in the appendix also hint that having a mixture of contrasts often helps, but the results are somewhat mixed.

### 5.3 Prompt Difficulty

We examine how the difficulty of training prompts in a preference dataset affects performance. For each comparison, we fix the data curation method and the margin between the (chosen, rejected) responses while comparing the difficulties of our training prompts with  $k \in \{4, 5, 6\}$  with a fixed training data size. Then, we repeat our experiments across the other two dimensions. Table 5 summarizes the results of our experiments.

**Training on moderately difficult prompts is overall more helpful than training on extremely difficult prompts.** We find that models trained on moderately difficult prompts perform better than models trained on extreme difficulty prompts in all



Method	Rejection Sampling (RS)				Method	Monte Carlo Tree Search (MCTS)			
	IFEval	Ours <sub>k=4</sub>	Ours <sub>k=5</sub>	Ours <sub>k=6</sub>		IFEval	Ours <sub>k=4</sub>	Ours <sub>k=5</sub>	Ours <sub>k=6</sub>
$k = 4, (c=3, r=0)$	79.70	38.41	21.15	14.27	$k = 4, (c=3, r=0)$	78.94	39.00	21.60	14.40
$k = 5, (c=4, r=1)$	79.51	38.66	21.48	14.70	$k = 5, (c=4, r=1)$	78.72	39.42	21.43	14.25
$k = 6, (c=5, r=2)$	77.83	36.42	20.15	12.28	$k = 6, (c=5, r=2)$	78.13	38.55	21.64	14.45
$k = 4, (c=3, r=1)$	79.36	39.22	22.00	14.78	$k = 4, (c=3, r=1)$	79.12	38.81	21.72	15.66
$k = 5, (c=4, r=2)$	78.58	37.70	21.42	13.61	$k = 5, (c=4, r=2)$	77.53	38.93	21.69	15.20
$k = 6, (c=5, r=3)$	76.73	36.16	19.96	12.87	$k = 6, (c=5, r=3)$	77.18	37.06	20.16	12.84

Table 5: Evaluation results investigating the effects of training prompt difficulty ( $k \in \{4, 5, 6\}$ ) on downstream performance for evaluation sets of varying difficulties. We provide results for both RS- and MCTS-based preference data curation methods, as well as for different margins between the (chosen, rejected) responses. Each  $(c = n_1, r = n_2)$  indicates that the chosen response correctly addresses  $n_1$  constraints and the rejected response correctly addresses  $n_2$  constraints. The soft score metrics for all experiments are provided in Tables 18 and 19 in the appendix.

our evaluation sets. This holds across both RS- and MCTS-based preference data curation methods.

**Training on moderately difficult prompts is more helpful even for performing well on extremely difficult prompts at test time.** Surprisingly, models trained on moderately difficult prompts ( $k = 4$ ) outperform models trained on extremely difficult prompts ( $k = 6$ ) even for evaluations with  $k = 6$ . Our results indicate that training prompts with moderate difficulties help the model to achieve generalization to more difficult prompts, even more so than overly difficult prompts.

## 6 Conclusion

We systematically investigate the effects of various attributes of preference datasets on model capabilities via instruction-following. To this end, we first build a data generation pipeline that combines general prompts with mixtures of verifiable constraints to synthesize challenging instruction-following prompts. We then automatically curate preference pairs using rejection sampling (RS) and Monte Carlo Tree Search (MCTS). Using the preference pairs, we examine the effects of (1) the existence of shared prefixes between the chosen and rejected responses, (2) the contrast and quality of the responses, and (3) the complexity of the training prompts. Our results indicate that having a common prefix in the preference pairs offers marginal yet consistent improvements, high-contrast preference pairs outperform low-contrast pairs but a mixture is sometimes better than both, and training on moderately difficult prompts is more helpful than training on extremely difficult prompts. Our work provides a systematic framework for curating different types of preference datasets and sets the groundwork for future studies that extend the scope beyond verifiable instruction-following constraints to more general constraints.

## 7 Limitations

Our work contains several limitations. First, we restrict our analysis to instruction-following with verifiable constraints. As discussed in the paper, we focus on 1) instruction-following in order to integrate multiple constraints in each prompt which allows us to assign fine-grained scores to any response, and 2) verifiable constraints in order to evaluate the quality of each response with high reliability and efficiency. Future work should explore training high-quality reward models that assign reliable scalar scores for more abstract constraints (e.g., does the response follow the topic indicated in the prompt) such that the ontology of constraints can be expanded beyond computationally verifiable constraints. Second, we perform our experiments only with models at the 8B scale. The number of different training configurations in our experiments, in addition to the number of outputs generated per prompt during evaluation, makes it difficult to scale our experiments to 70B and beyond. We expect that experiments at the 70B scale would provide further useful signals for curating preference data with larger models. Third, we restrict our preference data curation methods to rejection sampling and Monte Carlo Tree Search. Our purpose is not to exhaustively verify all methods for automatic preference data curation, but to compare two popular yet contrastive data curation methods: one which is quick and efficient but does not have any control over the structure of the preference pairs, and another which consumes more compute and is tricky to implement but generates fine-grained preference pairs with shared prefixes that can be controlled. Future work should investigate more diverse types of data curation methods to specifically examine which method offers the best performance improvements given the same constraints.

## References

- AI at Meta. 2024. [The llama 3 herd of models](#). *CoRR*, abs/2407.21783.
- Anthropic. 2024. The Claude 3 Model Family: Opus, Sonnet, Haiku. [https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model\\_Card\\_Claude\\_3.pdf](https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf).
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom B. Brown, Jack Clark, Sam McCandlish, Chris Olah, Benjamin Mann, and Jared Kaplan. 2022. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#). *CoRR*, abs/2204.05862.
- Google. 2023. [Gemini: A family of highly capable multimodal models](#). *CoRR*, abs/2312.11805.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. [Measuring mathematical problem solving with the MATH dataset](#). In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*.
- Hamish Ivison, Yizhong Wang, Jiacheng Liu, Zeqiu Wu, Valentina Pyatkin, Nathan Lambert, Noah A. Smith, Yejin Choi, and Hannaneh Hajishirzi. 2024. [Unpacking DPO and PPO: disentangling best practices for learning from preference feedback](#). *CoRR*, abs/2406.09279.
- Simran Kaur, Simon Park, Anirudh Goyal, and Sanjeev Arora. 2024. [Instruct-skillmix: A powerful pipeline for LLM instruction tuning](#). *CoRR*, abs/2408.14774.
- Saeed Khaki, JinJin Li, Lan Ma, Liu Yang, and Prathap Ramachandra. 2024. [RS-DPO: A hybrid rejection sampling and direct preference optimization method for alignment of large language models](#). In *Findings of the Association for Computational Linguistics: NAACL 2024, Mexico City, Mexico, June 16-21, 2024*, pages 1665–1680. Association for Computational Linguistics.
- Xin Lai, Zhuotao Tian, Yukang Chen, Senqiao Yang, Xiangru Peng, and Jiaya Jia. 2024. [Step-dpo: Step-wise preference optimization for long-chain reasoning of llms](#). *CoRR*, abs/2406.18629.
- Grégoire Mialon, Clémentine Fourrier, Thomas Wolf, Yann LeCun, and Thomas Scialom. 2024. [GAIA: a benchmark for general AI assistants](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Richard Yuanzhe Pang, Weizhe Yuan, Kyunghyun Cho, He He, Sainbayar Sukhbaatar, and Jason Weston. 2024. [Iterative reasoning preference optimization](#). *CoRR*, abs/2404.19733.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms](#). *CoRR*, abs/1707.06347.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2020. [Mpnet: Masked and permuted pre-training for language understanding](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F. Christiano. 2020. [Learning to summarize with human feedback](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Huai hsin Chi, and Denny Zhou. 2022. [Self-consistency improves chain of thought reasoning in language models](#). *ArXiv*, abs/2203.11171.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khoshabi, and Hannaneh Hajishirzi. 2023. [Self-instruct: Aligning language models with self-generated instructions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 13484–13508. Association for Computational Linguistics.

- Bosi Wen, Pei Ke, Xiaotao Gu, Lindong Wu, Hao Huang, Jinfeng Zhou, Wenchuang Li, Binxin Hu, Wendy Gao, Jiaxin Xu, Yiming Liu, Jie Tang, Hongning Wang, and Minlie Huang. 2024. [Benchmarking complex instruction-following with multiple constraints composition](#). *CoRR*, abs/2407.03978.
- Wenyi Xiao, Zechuan Wang, Leilei Gan, Shuai Zhao, Wanggui He, Luu Anh Tuan, Long Chen, Hao Jiang, Zhou Zhao, and Fei Wu. 2024. [A comprehensive survey of datasets, theories, variants, and applications in direct preference optimization](#). *CoRR*, abs/2410.15595.
- Yuxi Xie, Anirudh Goyal, Wenyue Zheng, Min-Yen Kan, Timothy P. Lillicrap, Kenji Kawaguchi, and Michael Shieh. 2024. [Monte carlo tree search boosts reasoning via iterative preference learning](#). *CoRR*, abs/2405.00451.
- Jing Xu, Andrew Lee, Sainbayar Sukhbaatar, and Jason Weston. 2023. [Some things are more CRINGE than others: Preference optimization with the pairwise cringe loss](#). *CoRR*, abs/2312.16682.
- Lifan Yuan, Ganqu Cui, Hanbin Wang, Ning Ding, Xingyao Wang, Jia Deng, Boji Shan, Huimin Chen, Ruobing Xie, Yankai Lin, Zhenghao Liu, Bowen Zhou, Hao Peng, Zhiyuan Liu, and Maosong Sun. 2024a. [Advancing LLM reasoning generalists with preference trees](#). *CoRR*, abs/2404.02078.
- Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Xian Li, Sainbayar Sukhbaatar, Jing Xu, and Jason Weston. 2024b. [Self-rewarding language models](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.
- Dan Zhang, Sining Zhoubian, Yisong Yue, Yuxiao Dong, and Jie Tang. 2024a. [Rest-mcts\\*: LLM self-training via process reward guided tree search](#). *CoRR*, abs/2406.03816.
- Tao Zhang, Yanjun Shen, Wenjing Luo, Yan Zhang, Hao Liang, Tao Zhang, Fan Yang, Mingan Lin, Yujing Qiao, Weipeng Chen, Bin Cui, Wentao Zhang, and Zenan Zhou. 2024b. [Cfbench: A comprehensive constraints-following benchmark for llms](#). *CoRR*, abs/2408.01122.
- Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023. [Instruction-following evaluation for large language models](#). *CoRR*, abs/2311.07911.
- Yiyang Zhou, Chenhang Cui, Rafael Rafailov, Chelsea Finn, and Huaxiu Yao. 2024. [Aligning modalities in vision large language models via preference fine-tuning](#). *CoRR*, abs/2402.11411.

## A Verifiable Constraints

We present the complete ontology of our 23 verifiable constraints in Table 6.

## B Synthetic Prompt Details

We generate instruction-following prompts for combinations of  $k \in \{4, 5, 6\}$  constraints using vLLM<sup>1</sup>. Table 7 shows the statistics that describe the prompts that we generate. For each value of  $k$ , we generate about 16K synthetic prompts, resulting in a total of about 48K prompts being used in our experiments. The average length of the prompts increases as the number of constraints increases, as the complexity of the prompts increases with larger numbers of constraints.

Table 8 displays examples of synthetic prompts generated using our pipeline for a given combination of verifiable constraints. Each prompt contains a general-purpose base instruction such as writing a short story or a speech, and is accompanied by a combination of verifiable constraints defined in our ontology with keyword arguments that satisfy the context of the instruction – for example, the first example uses the word "shouting" for the `nth_sent_first_word` constraint, which is appropriate for the context of the base instruction involving a boy lost in a shopping mall. For any arbitrary response to a given prompt, we can use the associated constraints and keyword arguments to automatically assign a score indicating whether the response follows each constraint and aggregate the scores to assign an overall correctness score to the response.

## C MCTS Details

**Computing the policy score.** We compute the policy score  $\Pi(a_i|s_t)$  by computing the average of the log probabilities of tokens generated for action  $a_i$  from state  $s_t$ , with the denominator moderated by a hyperparameter  $\gamma$ , which we set to 1.0 for our experiments. Refer to the formula below for the exact definition:

$$\Pi(a_i|s_t) = \exp \left[ \frac{1}{\|a_i\|^\gamma} \sum_{j=1}^{\|a_i\|} \Pi(t_j|s_t, t_{1..j-1}) \right]$$

Note that  $t_j$  denotes the  $j$ th token in action  $a_i$ .

**Computing the self-evaluation score.** We compute the self-evaluation score  $\Pi_{\text{self-eval}}$  by prompting the policy with a self-evaluation prompt

$P_{\text{self-eval}}$ , coupled with the response generated by the model so far, and obtaining the log probabilities of the final token of the response, denoted as  $t_{\text{final}} \in \{\text{yes}, \text{no}\}$ . We use self-consistency (Wang et al., 2022) to obtain multiple self-evaluations of the policy of its own output over  $L$  generations and average the scores in order to improve the reliability of our self-evaluation scores via increased compute.

$$\Pi_{\text{self-eval}} = \frac{1}{L} \sum_{i=1}^L \left[ \frac{1 + \exp(\Pi(t_{\text{final}} = \text{yes}|P_{\text{self-eval}}, s_t))}{2} - \frac{\exp(\Pi(t_{\text{final}} = \text{no}|P_{\text{self-eval}}, s_t))}{2} \right]$$

Our formula normalizes the score between 0 to 1, with 0.5 indicating a neutral state where the confidence scores for  $\Pi(t_{\text{final}} = \text{yes}|P_{\text{self-eval}}, s_t) = \Pi(t_{\text{final}} = \text{no}|P_{\text{self-eval}}, s_t)$ . We provide the self-evaluation prompt  $P_{\text{self-eval}}$  in Figure 5.

## D Preference Pair Details

We curate preference pairs using both RS- and MCTS-based methods for our synthetic prompts with  $k \in \{4, 5, 6\}$ , again using vLLM. Figure 4 shows the number of preference pairs obtained via both methods for  $k = 5$  when the correctness criteria for (chosen, rejected) pairs is set as (5 correct, 0 correct), (5 correct, 1 correct), (5 correct, 2 correct), (5 correct, 3 correct) and (5 correct, 4 correct) pairs. The left subfigure depicts the number of unique prompts corresponding to the preference pairs extracted for each filtering criteria. Using rejection sampling (RS) returns a higher yield of preference pairs with high contrast between the (chosen, rejected) responses, while using Monte Carlo Tree Search (MCTS) returns a higher yield of preference pairs with low contrast between the responses. The right subfigure depicts the total number of preference pairs extracted for each criteria – the same observation can be made about the relative yield, with MCTS yielding a large number of preference pairs due to its tree structure. We use preference pairs collected for  $k \in \{4, 5, 6\}$  using both methods over different filtering criteria to perform our experiments.

<sup>1</sup><https://github.com/vllm-project/vllm>



constraint	keyword args	description
alliteration	num_alliteration_words	the response should contain an alliteration, i.e. a sequence of $X$ words starting with the same letter, where $X = \text{num\_alliteration\_words}$ .
ascending_num_words	n/a	the response should contain sentences such that the number of words in each sentence is in ascending order.
edit_response	n/a	the response should be separated by dashes '——' separating two responses with the second response improving upon the first response.
end_quotation	n/a	the last sentence should be wrapped in quotation marks.
first_letter_capital	n/a	the first letter of each word should be capitalized.
frequency_long_words	relation, num_words, word_length	the response should contain $R \in \{\text{at least, at most}\}$ $X$ words that are at least $Y$ characters long each, where $R = \text{relation}$ , $X = \text{num\_words}$ and $Y = \text{word\_length}$ .
keywords_ordered	keywords	the response should contain a set of keywords in the given order and not in any other order.
max_word_length	max_word_length	the maximum length of all words in the response should be at most $X$ characters, where $X = \text{max\_word\_length}$ .
no_period	n/a	the response should contain no periods.
nth_sentence_capital	nth_sentence	the $\text{nth\_sentence}$ in the response should be all capitalized (and only the $\text{nth}$ sentence).
nth_sentence_first_word	first_word, num_sentences, nth_sentence	the first word of the $\text{nth\_sentence}$ should be the given word $W$ in the instruction, where $W = \text{first\_word}$ .
num_words_per_sentence	relation, num_words	each sentence in the response should contain $R \in \{\text{at least, at most}\}$ $X$ words, where $R = \text{relation}$ and $X = \text{num\_words}$ .
number_bold_words	num_words	the response should bold exactly $X$ words in HTML format <code>&lt;b&gt;word&lt;/b&gt;</code> , where $X = \text{num\_words}$ .
number_exclamations	relation, num_exclamations	the response should contain $R \in \{\text{at least, at most}\}$ $X$ exclamation marks, where $R = \text{relation}$ and $X = \text{num\_exclamations}$ .
number_italic_words	num_words	the response should italicize exactly $X$ words in textile format <code>'_word_'</code> , where $X = \text{num\_words}$ .
number_parentheses	num_parentheses	the response should contain exactly $X$ parentheses <code>()</code> , where $X = \text{num\_parentheses}$ .
number_parts	part_splitter, num_parts	the response should contain $X$ parts such as "Part 1", "PART 1", "Part 2", "PART 2", ... where $\text{part\_splitter} \in \{\text{Part, PART}\}$ and $X = \text{num\_parts}$ .
numbered_headers	num_headers	the response should contain $X$ enumerated headings starting with 1. , 2. , 3. , ... , where $X = \text{num\_headers}$ .
required_sentence	sentence	the response should contain a sentence $S$ , where $S = \text{sentence}$ .
start_checker	first_sentence	the response should begin with $\text{first\_sentence}$ .
tldr_summary	n/a	the response should end with a "TL;DR" on a new line summarizing the response.
variable_placeholder_format	relation, num_placeholders	the response should contain $R \in \{\text{at least, at most}\}$ $X$ variable placeholders in curly brackets, where $R = \text{relation}$ and $X = \text{num\_placeholders}$ .
vowel_capitalization	n/a	the response should capitalize all vowels in the response.

Table 6: Complete list of verifiable constraints used for our synthetic prompts.

$k = 4$			$k = 5$			$k = 6$		
num_prompts	mean_words	std_words	num_prompts	mean_words	std_words	num_prompts	mean_words	std_words
15,900	70.62	14.37	15,739	84.17	15.61	15,559	97.95	16.91

Table 7: Statistics of synthetic prompts generated by our pipeline. We generate about 16K prompts for each value of  $k$ , resulting in 48K prompts total across all our experiments. Note that `num_prompts` refers to the number of prompts, `mean_words` refers to the average number of words in each prompt, and `std_words` refers to the standard deviation of the number of words in each prompt. The number of words in each prompt increases with the number of constraints.




Figure 4: Number of preference pairs for different correctness filtering criteria at  $k = 5$ . The light blue color indicates preference pairs obtained via rejection sampling (RS), and the dark blue color indicates preference pairs obtained via Monte Carlo Tree Search (MCTS). The left subfigure shows the number of unique prompts with (chosen, rejected) responses associated with each filtering criteria, and the right subfigure shows the total number of preference pairs with (chosen, rejected) responses associated with each filtering criteria.

```
<|start_header_id|>user<|end_header_id|>

Evaluate whether the assistant's (partial) response to the given instruction follows the conditions
specified in the instruction so far and does not violate any of the conditions. Complete the
evaluation by using the words "yes" or "no", followed by an explanation for why the assistant's
response follows or does not follow the given instruction so far.

Do NOT evaluate the conditions that can be checked automatically with Python code, including the ones
listed below.
- DON'T EVALUATE: number of paragraphs/sentences/words/sections
- DON'T EVALUATE: existence of certain phrases/words/characters
- DON'T EVALUATE: capital or lowercase
Make sure to only evaluate the conditions that can be checked so far. For example, you cannot check
if the response contains at least 20 sentences- this is because the given response is a partial,
incomplete response and the full response later may possibly contain at least 20 sentences.
Also, this can be checked automatically, so it corresponds to the first "DO NOT" condition
listed above.

Instead, focus on the conditions that cannot be checked automatically and is more related to the
content itself, as listed below.
- DO EVALUATE: whether the response follows the topic so far
- DO EVALUATE: whether the response matches the description of the characters/location/theme/etc laid
out in the instruction so far
- DO EVALUATE: whether the response follows the tone requested in the instruction (e.g., persuasive,
solemn, lively, etc.)
For example, if the response asks to write a conversation between a software engineer and a research
scientist, make sure that there are two characters who are each software engineer and research
scientist, respectively.

Instruction: %s
Response so far: %s

Begin your response by listing such content-based conditions and analyzing whether each condition has
been satisfied on separate lines.
Be generous in terms of the evaluation criteria - only say "no" when you are sure that the partial
response does not adhere to the content-based conditions. Otherwise, answer "yes" to each
condition.
Most importantly, make sure to finish your evaluation with the phrase "Based on these evaluations, my
overall evaluation is: ", followed by either "yes" or "no".

<|eot_id|><|start_header_id|>assistant<|end_header_id|>
```

Figure 5: Our self-evaluation prompt  $P_{\text{self-eval}}$ . We allow the policy to focus on the soft content of the response rather than the hard constraints that are verifiable by code.

constraints	prompt
ascending_num_words, freq_long_words, max_word_length, nth_sent_first_word, start_checker	Write a short story about a boy who gets lost in a shopping mall. Include at least 7 words that are at least 12 characters long, and ensure that the sentences have an increasing number of words, i.e. each sentence should contain more words than its previous one. Also, only include words that are at most 12 characters long. Make sure that the fifth sentence starts with the word "shouting", and begin your response with the sentence "As the sounds of loud chatter and clinking of dishes filled the food court, little Tommy suddenly discovered that his parents were nowhere to be seen."
nth_sent_first_word, num_bold_words, num_exclamations, tldr_summary, vowel_capitalization	Write a motivational speech for a high school graduation ceremony. Capitalize the vowels in your response, and include seven words that are bolded in HTML format (e.g., <b>word</b>). Also, ensure that the sixth sentence starts with the word "today". Make sure that the response contains exactly three exclamation marks, and finish the response with the final line including "TL;DR" followed by a one-sentence summary of your response.
alliteration, keywords_ordered, nth_sent_first_word, number_bold_words, vowel_capitalization	Write a scene where a character walks into a room and is surprised by what they see. Capitalize the vowels in your response and include exactly six words that are bolded in HTML format (e.g., <b>word</b>). Include an alliteration of five consecutive words beginning with the same letter of the alphabet, and include the words 'door', 'space', 'chaos' in your response in the exact order provided. Moreover, make sure that the fourth sentence starts with the word "suddenly".
alliteration, number_italic_words, number_parts, required_sentence, start_checker	Write a short essay on the importance of vitamin D in human health. Start your response with the sentence "Vitamin D is an essential nutrient that has been gaining increasing attention in recent years, primarily due to its significant role in preventing various health issues, ranging from bone diseases to certain cancers." Include the sentence "Research has shown that vitamin D deficiency is associated with an increased risk of cardiovascular disease, diabetes, and osteoporosis." in your response, and include one word that is italicized in textile format, wrapped between underscore characters (e.g., _word_). Divide your essay into one part marked with 'Part I', and include an alliteration of five consecutive words beginning with the same letter of the alphabet.
edit_response, num_exclamations, start_checker, tldr_summary, vowel_capitalization	Write a short script of a news anchor introducing a breaking news story. Capitalize the vowels in your response, provide two responses separated by six plus signs +++++, and include at least nine exclamation marks. Also, include a "TL;DR" summarizing the breaking news story at the end of your response. Begin the response with the sentence "We interrupt your regular programming to bring you this breaking news story coming in from the White House."
alliteration, ascending_num_words, nth_sentence_capital, number_italic_words, number_parentheses	Create a data model for a medical records database, including at least one relationship and one constraint. Write the first sentence of the response in all capital letters, include eight words that are italicized in textile format, and include an alliteration of five consecutive words beginning with the same letter of the alphabet. Ensure that the sentences in the response have an increasing number of words, i.e. each sentence should contain more words than its previous one. Include exactly six parentheses in the response.
max_word_length, number_parentheses, number_parts, tldr_summary, vowel_capitalization	Write a list of 10 ways to improve your public speaking skills. Capitalize the vowels in the response. Have one part marked with PART I. Finish your response with the final line including "TL;DR" followed by a one-sentence summary of your response. Only include words that are at most 14 characters long, and include exactly eight parentheses in the response.
alliteration, nth_sentence_first_word, num_words_per_sentence, number_bold_words, vowel_capitalization	Write a letter to a historical figure, asking for their advice on a modern issue. Ensure that the vowels in the response are capitalized and there are eight words that are bolded in HTML format (e.g., <b>word</b>). Include an alliteration of three consecutive words beginning with the same letter of the alphabet. Make sure that the sixth sentence starts with the word "nonetheless", and that each sentence in the response is less than 15 words long.

Table 8: Examples of synthetic prompts generated for  $k = 5$ , which combines five verifiable constraints.

## E Additional Experiments

We perform two additional experiments to confirm that preference learning helps our models gain instruction-following skills, and investigate the limitations of our simpler RS-based data curation method. To this end, we 1) compare the performances of SFT and DPO to ensure that our preference datasets teach meaningful skills to our models, and 2) examine how RS-based preference data curation methods scale by varying the value of  $N$ , the number of samples generated for each prompt during rejection sampling.

### E.1 SFT vs. DPO

We demonstrate that our preference learning setup teaches meaningful skills to the models and hence provide a solid ground upon which to perform the comparison experiments shown in Sections 5.1, 5.2 and 5.3. To this end, we compare the performances of the models trained via DPO with the base policy model (llama-3.1-8b-instruct) and the policy model trained via supervised fine-tuning (SFT) on the chosen responses only. For SFT, we train each model in bf16 for three epochs with a maximum sequence length of 2048, learning rate of  $2e-6$  with a linear scheduler and 4 gradient accumulation steps with a total batch size of 32. We fix the training dataset size for all experiments.

Tables 9 and 10 shows the results of our experiments. First, we check that DPO significantly improves over the base policy model even though the verifiable constraints in the training and evaluation sets are distinct – this indicates that training on a set of verifiable constraints can be helpful for generalizing to another set of verifiable constraints as long as they involve transferable skills. Moreover, we observe that DPO outperforms SFT across all training configurations, demonstrating the importance of our preference learning setup with the rejected responses in addition to the chosen responses.

### E.2 Scaling Rejection Sampling

We investigate how different values of  $N$  during rejection sampling affect its overall quality and ensure that our RS-based method is designed to be competitive against MCTS. To this end, we re-run the rejection sampling curation pipeline for  $N \in \{4, 8, 16, 32, 64\}$  for a fixed difficulty  $k = 5$  and correctness scores of  $(c, r) = (4, 1)$  and  $(4, 2)$  and measure the performance across all of our evaluation sets.

Figure 6 shows the results of our experiments. We observe that the performance of the models trained on preference pairs curated with RS-based methods increases with  $N$ , but also saturate around  $N = 32$  and do not observe significant performance improvements afterward. This indicates that curating preference data via rejection sampling, while efficient and straightforward to scale, has limitations that cannot be solved by simply scaling the number of generated outputs, and further improvements may require more sophisticated search strategies such as the MCTS-based curation method implemented in this work.

## F Full Experiment Results

### F.1 Common Prefix Results

We provide the full results of our experiments investigating the effects of common prefixes in preference pairs in Tables 11, 12 and 13.

### F.2 Response Quality Results

**(chosen, rejected) response quality (unmixed).**

We provide the full results of our experiments investigating the effects of the response correctness (or quality) in preference pairs in Tables 14 and 15.

**(chosen, rejected) response quality (mixed).**

We provide the full results of our experiments examining the effects of mixing preference pairs with different margins between the (chosen, rejected) responses in Tables 16 and 17.

### F.3 Prompt Difficulty Results

We provide the full results of our experiments investigating the effects of varying the difficulty of the training prompts, as measured by the number of verifiable constraints, in Tables 18 and 19.



Method	IFEval	Rejection Sampling (RS)		
		Ours <sub>k=4</sub>	Ours <sub>k=5</sub>	Ours <sub>k=6</sub>
policy (llama-3.1-8b-instruct)	71.71 / 80.12	27.77 / 71.53	14.15 / 67.76	7.34 / 64.43
SFT, $k = 4$ , ( $c=4$ , $r=1/2/3$ )	74.77 / 82.58	28.17 / 72.40	15.10 / 68.04	7.10 / 64.36
DPO, $k = 4$ , ( $c=4$ , $r=1/2/3$ )	78.70 / 85.48	36.68 / 76.78	20.20 / 70.71	13.57 / 68.45
SFT, $k = 5$ , ( $c=4$ , $r=1/2/3$ )	74.75 / 82.59	28.36 / 72.42	14.53 / 68.37	7.74 / 65.07
DPO, $k = 5$ , ( $c=4$ , $r=1/2/3$ )	79.08 / 85.56	37.98 / 77.83	20.50 / 71.48	13.42 / 68.86

Table 9: Evaluation results comparing the performance of training on our preference datasets via DPO compared to the base policy model or running SFT on the chosen responses only. We provide results for RS-based preference data curation, as well as for different training prompt difficulties ( $k = 4$  or  $k = 5$ ). Each ( $c = n_1, r = n_2$ ) indicates that the chosen response correctly addresses  $n_1$  constraints and the rejected response correctly addresses  $n_2$  constraints. The left score indicates the hard score (i.e., the proportion of responses that get *all* constraints correct), and the right score indicates the soft score (i.e., the proportion of all constraints that are satisfied by the responses).

Method	IFEval	Monte Carlo Tree Search (MCTS)		
		Ours <sub>k=4</sub>	Ours <sub>k=5</sub>	Ours <sub>k=6</sub>
policy (llama-3.1-8b-instruct)	71.71 / 80.12	27.77 / 71.53	14.15 / 67.76	7.34 / 64.43
SFT, $k = 4$ , ( $c=4$ , $r=1/2/3$ )	73.30 / 81.83	29.70 / 73.48	15.02 / 69.16	7.50 / 65.56
DPO, $k = 4$ , ( $c=4$ , $r=1/2/3$ )	79.97 / 86.51	39.31 / 78.22	22.19 / 72.14	15.89 / 70.36
SFT, $k = 5$ , ( $c=4$ , $r=1/2/3$ )	73.87 / 81.89	28.54 / 72.37	13.49 / 67.54	6.68 / 64.13
DPO, $k = 5$ , ( $c=4$ , $r=1/2/3$ )	78.47 / 85.25	39.03 / 78.06	22.63 / 72.49	15.25 / 70.05

Table 10: Evaluation results comparing the performance of training on our preference datasets via DPO compared to the base policy model or running SFT on the chosen responses only. We provide results for MCTS-based preference data curation, as well as for different training prompt difficulties ( $k = 4$  or  $k = 5$ ). Each ( $c = n_1, r = n_2$ ) indicates that the chosen response correctly addresses  $n_1$  constraints and the rejected response correctly addresses  $n_2$  constraints. The left score indicates the hard score (i.e., the proportion of responses that get *all* constraints correct), and the right score indicates the soft score (i.e., the proportion of all constraints that are satisfied by the responses).

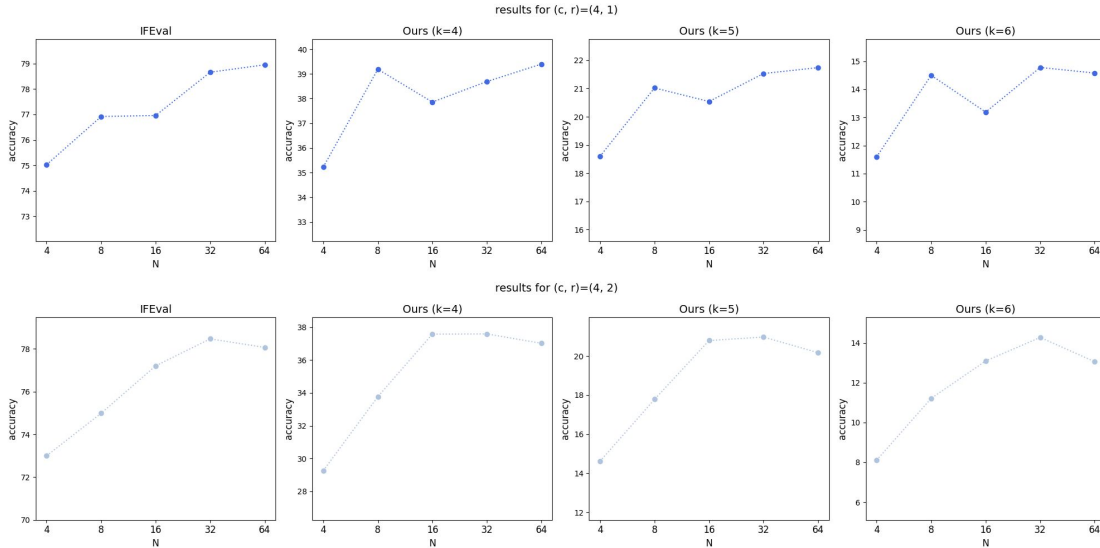


Figure 6: Evaluation results for increasing the number of outputs generated for rejection sampling (RS) from  $N = 4$  to  $N = 64$ , given a set of training prompts with  $k = 5$  and  $(c, r) = (4, 1)$  or  $(4, 2)$ . We observe a steady improvement in performance as more outputs are generated per prompt until  $N = 32$ , where it begins to saturate or even deteriorate.

Method	IFEval	training k=4		
		Ours <sub>k=4</sub>	Ours <sub>k=5</sub>	Ours <sub>k=6</sub>
RS, (c=4, r=1)	79.24 / 85.78	37.49 / 77.50	20.54 / 72.06	13.78 / 70.23
MCTS, (c=4, r=1)	79.48 / 85.92	39.16 / 78.11	20.39 / 72.04	14.11 / 70.10
RS, (c=4, r=2)	78.86 / 85.74	38.56 / 78.02	21.66 / 72.34	14.71 / 70.00
MCTS, (c=4, r=2)	79.68 / 86.36	39.22 / 78.20	22.43 / 72.64	15.75 / 70.61
RS, (c=4, r=3)	76.74 / 83.92	35.02 / 76.25	19.40 / 71.00	12.85 / 68.63
MCTS, (c=4, r=3)	79.59 / 86.30	39.05 / 77.86	21.61 / 72.28	14.96 / 70.22
RS, (c=3, r=0)	80.06 / 86.69	39.25 / 78.55	21.54 / 72.44	14.37 / 70.32
MCTS, (c=3, r=0)	79.39 / 85.86	39.15 / 78.02	21.56 / 72.26	13.61 / 69.42
RS, (c=3, r=1)	80.06 / 86.42	39.15 / 78.35	21.90 / 72.01	15.20 / 70.09
MCTS, (c=3, r=1)	79.94 / 86.37	39.23 / 78.37	22.15 / 72.34	14.93 / 70.00
RS, (c=3, r=2)	77.52 / 84.27	36.20 / 76.86	19.52 / 70.96	12.48 / 68.68
MCTS, (c=3, r=2)	77.89 / 84.85	38.89 / 78.02	21.51 / 71.93	13.97 / 69.58
RS, (c=4, r=1/2/3)	78.70 / 85.48	36.68 / 76.78	20.20 / 70.71	13.57 / 68.45
MCTS, (c=4, r=1/2/3)	79.97 / 86.51	39.31 / 78.22	22.19 / 72.14	15.89 / 70.36
RS, (c=3/4, r=0/1/2/3)	79.10 / 85.72	37.62 / 77.59	20.98 / 71.28	13.99 / 68.65
MCTS, (c=3/4, r=0/1/2/3)	79.42 / 85.84	39.37 / 78.38	22.24 / 72.48	15.20 / 70.26

Table 11: Evaluation results comparing preference data without common prefixes (RS) and with common prefixes (MCTS). We show results for different training data configurations for  $k = 4$ . Each  $(c = n_1, r = n_2)$  indicates that the chosen response correctly addresses  $n_1$  constraints and the rejected response correctly addresses  $n_2$  constraints. The left score indicates the hard score (i.e., the proportion of responses that get *all* constraints correct), and the right score indicates the soft score (i.e., the proportion of all constraints that are satisfied by the responses).

Method	IFEval	training k=5		
		Ours <sub>k=4</sub>	Ours <sub>k=5</sub>	Ours <sub>k=6</sub>
RS, (c=5, r=2)	76.32 / 83.55	35.70 / 76.64	19.46 / 71.51	12.61 / 69.39
MCTS, (c=5, r=2)	76.59 / 83.81	37.81 / 77.58	20.06 / 72.01	12.93 / 69.85
RS, (c=5, r=3)	76.25 / 83.63	35.28 / 76.23	18.91 / 70.03	12.24 / 68.75
MCTS, (c=5, r=3)	76.47 / 83.84	36.73 / 77.10	19.94 / 71.82	13.62 / 69.52
RS, (c=5, r=4)	74.18 / 81.79	33.38 / 75.10	17.35 / 70.05	10.38 / 67.54
MCTS, (c=5, r=4)	75.15 / 82.73	34.88 / 76.06	18.75 / 71.06	11.74 / 69.03
RS, (c=4, r=1)	78.95 / 85.57	39.40 / 78.60	21.74 / 72.48	14.58 / 70.02
MCTS, (c=4, r=1)	78.63 / 85.49	39.21 / 78.33	21.19 / 72.24	14.63 / 70.07
RS, (c=4, r=2)	78.07 / 84.87	37.02 / 77.35	20.17 / 71.02	13.05 / 68.48
MCTS, (c=4, r=2)	78.48 / 85.24	38.53 / 77.91	22.12 / 72.16	15.22 / 69.69
RS, (c=4, r=3)	75.64 / 82.97	32.98 / 75.32	17.83 / 69.95	11.19 / 67.40
MCTS, (c=4, r=3)	77.40 / 84.28	37.63 / 77.41	20.84 / 71.72	14.06 / 69.42
RS, (c=4, r=1/2/3)	79.08 / 85.56	37.98 / 77.83	20.50 / 71.48	13.42 / 68.86
MCTS, (c=4, r=1/2/3)	78.47 / 85.25	39.03 / 78.06	22.63 / 72.49	15.25 / 70.05
RS, (c=4/5, r=0/1/2/3)	78.18 / 85.05	36.81 / 77.06	19.72 / 71.13	13.29 / 68.67
MCTS, (c=4/5, r=0/1/2/3)	77.89 / 84.82	38.17 / 77.81	22.00 / 72.19	14.38 / 69.63

Table 12: Evaluation results comparing preference data without common prefixes (RS) and with common prefixes (MCTS). We show results for different training data configurations for  $k = 5$ . Each  $(c = n_1, r = n_2)$  indicates that the chosen response correctly addresses  $n_1$  constraints and the rejected response correctly addresses  $n_2$  constraints. The left score indicates the hard score (i.e., the proportion of responses that get *all* constraints correct), and the right score indicates the soft score (i.e., the proportion of all constraints that are satisfied by the responses).

Method	IFEval	training k=5		
		Ours <sub>k=4</sub>	Ours <sub>k=5</sub>	Ours <sub>k=6</sub>
RS, (c=6, r=2)	75.17 / 82.74	33.45 / 75.19	17.00 / 69.84	10.19 / 67.39
MCTS, (c=6, r=2)	75.70 / 83.11	34.06 / 75.71	17.07 / 70.27	10.46 / 67.64
RS, (c=6, r=3)	77.47 / 84.44	36.99 / 77.22	20.05 / 71.16	13.27 / 68.39
MCTS, (c=6, r=3)	78.06 / 85.04	38.68 / 77.96	21.63 / 72.21	14.94 / 70.02
RS, (c=6, r=4)	76.09 / 83.51	35.68 / 75.63	19.07 / 71.01	11.99 / 68.47
MCTS, (c=6, r=4)	76.26 / 83.58	35.81 / 75.63	18.41 / 71.05	11.56 / 68.64
RS, (c=5, r=1)	75.13 / 82.77	34.17 / 75.56	18.48 / 70.93	11.93 / 68.49
MCTS, (c=5, r=1)	76.09 / 83.47	36.34 / 76.52	19.04 / 71.33	12.70 / 68.62
RS, (c=5, r=2)	75.01 / 82.68	34.48 / 75.88	18.10 / 70.64	11.21 / 68.60
MCTS, (c=5, r=2)	76.57 / 83.83	36.03 / 76.72	18.75 / 70.93	11.62 / 68.68
RS, (c=5, r=3)	76.73 / 83.89	36.16 / 76.62	19.96 / 70.97	12.87 / 68.15
MCTS, (c=5, r=3)	77.18 / 84.53	37.06 / 77.10	20.16 / 71.25	12.84 / 68.57
RS, (c=5, r=1/2/3)	78.10 / 84.97	37.66 / 77.61	21.45 / 71.49	13.69 / 68.84
MCTS, (c=5, r=1/2/3)	78.33 / 85.17	40.12 / 78.34	23.12 / 72.32	15.28 / 69.73
RS, (c=5/6, r=1/2/3/4)	78.48 / 85.28	37.82 / 77.86	22.73 / 72.02	14.39 / 69.48
MCTS, (c=5/6, r=1/2/3/4)	79.01 / 85.72	40.58 / 78.93	23.04 / 73.18	16.00 / 70.57

Table 13: Evaluation results comparing preference data without common prefixes (RS) and with common prefixes (MCTS). We show results for different training data configurations for  $k = 6$ . Each  $(c = n_1, r = n_2)$  indicates that the chosen response correctly addresses  $n_1$  constraints and the rejected response correctly addresses  $n_2$  constraints. The left score indicates the hard score (i.e., the proportion of responses that get *all* constraints correct), and the right score indicates the soft score (i.e., the proportion of all constraints that are satisfied by the responses).

Method	IFEval	training k=4		
		Ours <sub>k=4</sub>	Ours <sub>k=5</sub>	Ours <sub>k=6</sub>
RS, (c=3, r=0)	79.04 / 85.87	39.45 / 78.30	20.86 / 72.23	13.82 / 70.45
RS, (c=3, r=1)	78.66 / 85.51	38.47 / 78.15	20.56 / 72.37	14.02 / 70.13
RS, (c=3, r=2)	74.60 / 82.39	33.59 / 75.58	18.46 / 70.62	11.23 / 68.15
RS, (c=4, r=1)	79.24 / 85.78	37.49 / 77.50	20.54 / 72.06	13.78 / 70.23
RS, (c=4, r=2)	77.44 / 84.43	37.69 / 77.68	20.33 / 72.06	13.77 / 70.04
RS, (c=4, r=3)	74.09 / 82.16	33.38 / 75.30	17.65 / 70.21	10.73 / 67.79
MCTS, (c=3, r=0)	78.30 / 85.19	37.86 / 77.42	20.52 / 71.92	12.81 / 69.30
MCTS, (c=3, r=1)	77.58 / 84.69	38.03 / 77.74	19.76 / 71.84	13.71 / 70.08
MCTS, (c=3, r=2)	75.05 / 82.63	34.96 / 76.33	18.69 / 71.25	11.53 / 68.67
MCTS, (c=4, r=1)	79.48 / 85.92	39.16 / 78.11	20.39 / 72.04	14.11 / 70.10
MCTS, (c=4, r=2)	77.76 / 84.65	38.25 / 77.63	20.00 / 71.92	14.26 / 69.91
MCTS, (c=4, r=3)	75.65 / 82.99	35.03 / 75.68	18.80 / 71.88	12.36 / 68.81

Table 14: Evaluation results studying the effects of (chosen, rejected) response quality. We provide results for  $k = 4$ , as well as for both RS- and MCTS-based data curation methods. Each  $(c = n_1, r = n_2)$  indicates that the chosen response correctly addresses  $n_1$  constraints and the rejected response correctly addresses  $n_2$  constraints. The left score indicates the hard score (i.e., the proportion of responses that get *all* constraints correct), and the right score indicates the soft score (i.e., the proportion of all constraints that are satisfied by the responses).

Method	IFEval	training k=5		
		Ours <sub>k=4</sub>	Ours <sub>k=5</sub>	Ours <sub>k=6</sub>
RS, (c=4, r=1)	76.56 / 83.74	36.08 / 76.81	19.44 / 71.42	11.88 / 68.84
RS, (c=4, r=2)	74.57 / 82.17	34.09 / 75.82	17.95 / 70.71	10.93 / 68.35
RS, (c=4, r=3)	72.80 / 80.97	31.50 / 74.37	16.68 / 69.94	10.01 / 67.19
RS, (c=5, r=2)	76.32 / 83.55	35.70 / 76.64	19.46 / 71.51	12.61 / 69.39
RS, (c=5, r=3)	74.13 / 81.92	33.30 / 75.35	17.31 / 70.15	11.21 / 68.08
RS, (c=5, r=4)	72.40 / 80.42	30.72 / 73.75	16.26 / 69.31	9.51 / 66.76
MCTS, (c=4, r=1)	76.71 / 83.90	37.24 / 77.23	19.49 / 71.52	12.37 / 69.19
MCTS, (c=4, r=2)	75.23 / 82.86	35.65 / 76.47	18.74 / 71.32	12.17 / 69.06
MCTS, (c=4, r=3)	74.23 / 82.16	32.39 / 74.63	17.22 / 69.78	9.73 / 67.25
MCTS, (c=5, r=2)	76.59 / 83.81	37.81 / 77.58	20.06 / 72.01	12.93 / 69.85
MCTS, (c=5, r=3)	76.14 / 83.50	36.04 / 76.61	19.39 / 71.49	12.37 / 69.32
MCTS, (c=5, r=4)	74.06 / 81.84	32.14 / 74.62	17.81 / 70.37	10.52 / 67.89

Table 15: Evaluation results studying the effects of (chosen, rejected) response quality. We provide results for  $k = 5$ , as well as for both RS- and MCTS-based data curation methods. Each  $(c = n_1, r = n_2)$  indicates that the chosen response correctly addresses  $n_1$  constraints and the rejected response correctly addresses  $n_2$  constraints. The left score indicates the hard score (i.e., the proportion of responses that get *all* constraints correct), and the right score indicates the soft score (i.e., the proportion of all constraints that are satisfied by the responses).

Method	IFEval	training k=4		
		Ours <sub>k=4</sub>	Ours <sub>k=5</sub>	Ours <sub>k=6</sub>
MCTS, (c=4, r=1)	78.49 / 85.36	39.03 / 77.76	22.19 / 72.10	14.89 / 69.76
MCTS, (c=4, r=(1,2))	79.77 / 86.31	38.97 / 78.03	21.99 / 71.91	15.24 / 70.26
MCTS, (c=4, r=(1,2,3))	79.97 / 86.51	39.31 / 78.22	22.19 / 72.14	15.89 / 70.36
MCTS, (c=(3,4), r=(0,1))	78.60 / 85.31	39.01 / 77.06	22.90 / 71.89	14.01 / 69.06
MCTS, (c=(3,4), r=(0,1,2))	79.60 / 86.05	39.89 / 78.71	22.69 / 72.70	16.10 / 70.60
MCTS, (c=(3,4), r=(0,1,2,3))	79.42 / 85.84	39.37 / 78.38	22.24 / 72.48	15.20 / 70.26

Table 16: Evaluation results studying the effects of mixing preference pairs with different margins between the (chosen, rejected) responses. We provide results for  $k = 4$  with MCTS-based data curation methods. Each  $(c = n_0, r = (n_1, n_2, n_3))$  indicates that the chosen response correctly addresses  $n_0$  constraints and the rejected response correctly addresses either  $n_1$ ,  $n_2$  or  $n_3$  constraints. The left score indicates the hard score (i.e., the proportion of responses that get *all* constraints correct), and the right score indicates the soft score (i.e., the proportion of all constraints that are satisfied by the responses).

Method	IFEval	training k=5		
		Ours <sub>k=4</sub>	Ours <sub>k=5</sub>	Ours <sub>k=6</sub>
MCTS, (c=4, r=1)	78.88 / 85.69	39.55 / 78.35	21.88 / 71.98	14.59 / 69.56
MCTS, (c=4, r=(1,2))	78.50 / 85.30	37.67 / 77.55	21.29 / 71.51	14.39 / 69.15
MCTS, (c=4, r=(1,2,3))	78.47 / 85.25	39.03 / 78.06	22.63 / 72.49	15.25 / 70.05
MCTS, (c=(4,5), r=(1,2))	78.82 / 85.60	39.28 / 78.37	22.31 / 72.30	14.50 / 69.71
MCTS, (c=(4,5), r=(1,2,3))	78.99 / 85.57	38.74 / 77.89	21.85 / 71.87	13.89 / 69.00
MCTS, (c=(4,5), r=(1,2,3,4))	77.89 / 84.82	38.17 / 77.81	22.00 / 72.19	14.38 / 69.63

Table 17: Evaluation results studying the effects of mixing preference pairs with different margins between the (chosen, rejected) responses. We provide results for  $k = 5$  with MCTS-based data curation methods. Each  $(c = n_0, r = (n_1, n_2, n_3))$  indicates that the chosen response correctly addresses  $n_0$  constraints and the rejected response correctly addresses either  $n_1$ ,  $n_2$  or  $n_3$  constraints. The left score indicates the hard score (i.e., the proportion of responses that get *all* constraints correct), and the right score indicates the soft score (i.e., the proportion of all constraints that are satisfied by the responses).



Method	IFEval	Rejection Sampling (RS)		
		Ours <sub>k=4</sub>	Ours <sub>k=5</sub>	Ours <sub>k=6</sub>
$k = 4, (c=3, r=0)$	79.70 / 86.13	38.41 / 77.93	21.15 / 72.36	14.27 / 69.90
$k = 5, (c=4, r=1)$	79.51 / 86.02	38.66 / 78.17	21.48 / 72.23	14.70 / 70.04
$k = 6, (c=5, r=2)$	77.83 / 84.86	36.42 / 77.23	20.15 / 71.18	12.28 / 68.34
$k = 4, (c=3, r=1)$	79.36 / 85.95	39.22 / 78.51	22.00 / 72.54	14.78 / 70.07
$k = 5, (c=4, r=2)$	78.58 / 85.46	37.70 / 77.60	21.42 / 71.84	13.61 / 69.14
$k = 6, (c=5, r=3)$	76.73 / 83.89	36.16 / 76.62	19.96 / 70.97	12.87 / 68.15

Table 18: Evaluation results investigating the effects of training prompt difficulty ( $k \in \{4, 5, 6\}$ ) on downstream performance for evaluation sets of varying difficulties. We provide results for RS-based preference data curation, as well as for different margins between the (chosen, rejected) responses. Each  $(c = n_1, r = n_2)$  indicates that the chosen response correctly addresses  $n_1$  constraints and the rejected response correctly addresses  $n_2$  constraints. The left score indicates the hard score (i.e., the proportion of responses that get *all* constraints correct), and the right score indicates the soft score (i.e., the proportion of all constraints that are satisfied by the responses).

Method	IFEval	Monte Carlo Tree Search (MCTS)		
		Ours <sub>k=4</sub>	Ours <sub>k=5</sub>	Ours <sub>k=6</sub>
$k = 4, (c=3, r=0)$	78.94 / 85.62	39.00 / 78.10	21.60 / 72.54	14.40 / 70.04
$k = 5, (c=4, r=1)$	78.72 / 85.40	39.42 / 78.36	21.43 / 72.27	14.25 / 69.95
$k = 6, (c=5, r=2)$	78.13 / 85.05	38.55 / 77.88	21.64 / 72.23	14.45 / 69.81
$k = 4, (c=3, r=1)$	79.12 / 85.83	38.81 / 78.24	21.72 / 72.48	15.66 / 70.45
$k = 5, (c=4, r=2)$	77.53 / 84.65	38.93 / 78.16	21.69 / 72.41	15.20 / 70.62
$k = 6, (c=5, r=3)$	77.18 / 84.53	37.06 / 77.10	20.16 / 71.25	12.84 / 68.57

Table 19: Evaluation results investigating the effects of training prompt difficulty ( $k \in \{4, 5, 6\}$ ) on downstream performance for evaluation sets of varying difficulties. We provide results for MCTS-based preference data curation, as well as for different margins between the (chosen, rejected) responses. Each  $(c = n_1, r = n_2)$  indicates that the chosen response correctly addresses  $n_1$  constraints and the rejected response correctly addresses  $n_2$  constraints. The left score indicates the hard score (i.e., the proportion of responses that get *all* constraints correct), and the right score indicates the soft score (i.e., the proportion of all constraints that are satisfied by the responses).