

# SFTMIX: ELEVATING LANGUAGE MODEL INSTRUCTION TUNING WITH MIXUP RECIPE

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

To induce desired behaviors in large language models (LLMs) for interaction-driven tasks, the instruction-tuning stage typically trains LLMs on instruction-response pairs using the next-token prediction (NTP) loss. Previous work aiming to improve instruction-tuning performance often emphasizes the need for higher-quality supervised fine-tuning (SFT) datasets, which typically involves expensive data filtering with proprietary LLMs or labor-intensive data generation by human annotators. However, these approaches do not fully leverage the datasets’ intrinsic properties, resulting in high computational and labor costs, thereby limiting scalability and performance gains. In this paper, we propose SFTMix, a novel recipe that elevates instruction-tuning performance beyond the conventional NTP paradigm, without the need for well-curated datasets. Observing that LLMs exhibit uneven confidence across the semantic representation space, we argue that examples with different confidence levels should play distinct roles during the instruction-tuning process. Based on this insight, SFTMix leverages training dynamics to identify examples with varying confidence levels, then applies a Mixup-based regularization to mitigate overfitting on confident examples while propagating supervision signals to improve learning on relatively unconfident ones. This approach enables SFTMix to significantly outperform NTP across a wide range of instruction-following and healthcare domain-specific SFT tasks, demonstrating its adaptability to diverse LLM families and scalability to datasets of any size. Comprehensive ablation studies further verify the robustness of SFTMix’s design choices, underscoring its versatility in consistently enhancing performance across different LLMs and datasets in broader natural language processing applications.<sup>1</sup>

## 1 INTRODUCTION

Large language models (LLMs) have recently demonstrated outstanding performance across a broad spectrum of natural language processing (NLP) tasks (Zhao et al., 2023; Minaee et al., 2024). After being pre-trained on large corpora of raw text, LLMs undergo a critical instruction-tuning stage (Ouyang et al., 2022; Zhang et al., 2023) to develop their instruction-following capabilities based on supervised fine-tuning (SFT) datasets, making them more suitable for interaction-driven applications. SFT datasets (Taori et al., 2023; Wang et al., 2023; Xu et al., 2024) typically consist of instruction-response pairs spanning various task types, aligning LLMs toward desired behavior. During this stage, LLMs are usually trained through next-token prediction (NTP), where LLMs predict the next token in a response given both the instruction and the preceding tokens in that response.

Previous research efforts in this field have predominantly focused on enhancing the quality of instruction-tuning datasets. One line of research direction seeks to better understand the intrinsic properties of these datasets (Kung et al., 2023; Lin et al., 2024) and selects informative instruction-response pairs through heuristics-based filters (Zhao et al., 2024) or LLM scoring (Chen et al., 2024). Another line of work generates high-quality responses by querying advanced proprietary LLMs (Chen et al., 2024) or relying on human annotators (Zhou et al., 2023). However, both strategies come with significant computational or labor costs, limiting the scalability of SFT datasets.

In this paper, we take a different perspective by exploring how to elevate instruction-tuning performance beyond the conventional NTP training paradigm, without relying on well-curated datasets

<sup>1</sup>We will release our implementation after the anonymous review period.

To address this challenge, we propose SFTMix, a novel Mixup-based recipe for language model instruction tuning. Our design builds upon the key observation that an LLM’s confidence distribution over its instruction-tuning dataset is uneven across the semantic representation space. We argue that data with varying confidence levels should contribute differently during the instruction-tuning process. Hence, we extend data cartography (Swayamdipta et al., 2020) to the realm of causal language generation as training dynamics and leverage a reference LLM to derive the confidence of each instruction-response pair, based on the perplexities computed over the instruction-tuning process. Using this information, we divide the original SFT dataset into a confident subset and a relatively unconfident subset of equal size.

To guide the learning of the unconfident examples by propagating supervision signals, and to mitigate overfitting to the confident semantic regions, we design a Mixup-based (Zhang et al., 2018) regularization for LLM instruction tuning. This approach explores the utility of Mixup in causal language generation and improves its effectiveness by exploiting the confidence information derived from training dynamics. Specifically, consider an instruction-response pair from the confident subset and another from the unconfident subset. We interpolate them linearly at the token level in the representation space and generate a convex combination of the representations from the LLM under instruction tuning. The one-hot encodings of the corresponding tokens are interpolated in a similar fashion. We then compute a regularization between the interpolated encodings and representations, in addition to the original NTP loss during instruction tuning. In this way, SFTMix fosters the synergy between the two subsets with diverging confidence levels and enhances the interaction capabilities of LLMs across diverse downstream applications.

We demonstrate the effectiveness of our proposed SFTMix recipe in both instruction-following and domain-specific SFT settings. In particular, SFTMix significantly surpasses the conventional NTP instruction-tuning baseline in both single- and multi-turn conversations, as measured in MT-Bench (Zheng et al., 2024) and AlpacaEval-2 (Dubois et al., 2024). This improvement is consistent across different LLM families (e.g., Llama (Dubey et al., 2024) and Mistral (Jiang et al., 2023)) and scales of SFT datasets (e.g., Alpaca-52K (Taori et al., 2023) and UltraChat-200K (Tunstall et al., 2023)). Moreover, in the healthcare domain, Llama-3.1-8B (Dubey et al., 2024) and Mistral-7B-v0.1 (Jiang et al., 2023), instruction-tuned on MedAlpaca-263K (Han et al., 2023) using SFTMix, achieve an average of 1.5% absolute increase in accuracy across four benchmarks compared to baselines. We further explore the applicability of SFTMix’s variants through extensive ablations and illustrate the potential of this LLM instruction-tuning recipe for broader use cases.

We summarize our contributions in this paper as follows:

- We introduce SFTMix, a novel recipe designed to elevate LLM instruction-tuning performance beyond the conventional NTP paradigm, without relying on well-curated datasets.
- Motivated by the observation that LLMs exhibit varying confidence levels across the semantic space, SFTMix leverages LLMs’ training dynamics for more insightful data interpretation and confidence-based splitting, facilitating more effective instruction tuning.
- SFTMix further incorporates a Mixup-based regularization that interpolates between examples with different confidence levels during instruction tuning, mitigating overfitting to confident examples while improving generalization on relatively unconfident ones.
- We demonstrate that SFTMix significantly outperforms the NTP baseline across a variety of instruction-following and healthcare domain-specific SFT tasks, with consistent improvements across different LLM families and dataset sizes.
- Comprehensive ablation analysis substantiates the robustness of our design choices in SFTMix, shedding light on its potential for broader NLP applications.

## 2 RELATED WORK

**LLM Instruction Tuning.** To align LLMs with users’ open-ended intents or adapt them to specific domains, Ouyang et al. (2022) proposed instruction-tuning LLMs on human-annotated demonstrations using supervised learning. More specifically, given a pair of instructions and desired responses, the conventional NTP paradigm trains an LLM to predict each token in the response sequentially during the instruction-tuning stage (Zhang et al., 2023). Jain et al. (2024) improved instruction-tuning performance by adding noise to the token embeddings during training, while Shi et al. (2024)

further suggested modeling the instructions as well. On this basis, previous work (Chiang et al., 2023; Ding et al., 2023; Taori et al., 2023; Wang et al., 2023; Xu et al., 2024) collected instruction-following datasets by distilling powerful proprietary LLMs or crowdsourcing user conversations. To enhance data quality, the community has employed various techniques, including heuristic-based filters (Schoch et al., 2023; Zhao et al., 2024), LLM scoring (Chen et al., 2024), and human curation (Zhou et al., 2023). Other efforts (Kung et al., 2023; Lin et al., 2024) have focused on gaining a deeper understanding of the intrinsic properties of SFT datasets. However, acquiring high-quality SFT data often entails substantial computational and labor costs. In this paper, we aim to optimize data utilization through insightful data interpretation and improve the effectiveness of instruction tuning beyond the conventional NTP paradigm, without relying on well-curated datasets.

**Data Characterization via Training Dynamics.** Data characterization (Albalak et al., 2024; Wang et al., 2024) seeks to assess and analyze the quality and relevance of training data, enabling more effective data filtering and elevated model performance. In particular, Swayamdipta et al. (2020) leveraged the training dynamics of a pre-trained language model (Liu, 2019) to create data maps, which have subsequently inspired advancements in active learning (Zhang & Plank, 2021; Zhang et al., 2022; Kung et al., 2023), curriculum learning (Christopoulou et al., 2022; Lin et al., 2024; Poesina et al., 2024), and dataset pruning (Chimoto et al., 2024; He et al., 2024; Lin et al., 2024; Seedat et al., 2024). Here, we explore applying training dynamics to causal language generation by categorizing an SFT dataset into confident and relatively unconfident subsets, which facilitates the subsequent Mixup-based regularization during LLM instruction tuning.

**Mixup-Based Learning.** To alleviate memorization and sensitivity to adversarial examples during training, Zhang et al. (2018) proposed Mixup, which trains models on convex combinations of pairs of input features and their corresponding labels. Its variants (Verma et al., 2019; Hendrycks et al., 2020; Uddin et al., 2021; Choi et al., 2022) further suggest interpolating feature representations at different stages, guided by various training signals. Theoretical analyses (Zhang et al., 2021; Carratino et al., 2022; Chidambaram et al., 2022; Park et al., 2022; Pinto et al., 2022) have demonstrated its data-adaptive regularization and generalization effects, leading to strong out-of-distribution robustness and well-calibrated uncertainty estimation. Empirical studies have further validated its effectiveness under the semi-supervised learning setting (Berthelot et al., 2019; 2020; Li et al., 2020; 2022) and in diverse NLP applications (Chen et al., 2020; Guo et al., 2020; Sun et al., 2020; Park & Caragea, 2022; Yang et al., 2022). Building on this success, we explore its utility in LLM instruction tuning and propose a Mixup-based regularization to reduce overfitting to confident examples and support the learning of relatively unconfident ones.

### 3 SFTMIX

In Section 3.1, we begin by reviewing the conventional instruction-tuning task of NTP. We then introduce SFTMIX, a novel recipe for LLM instruction tuning. SFTMIX first leverages training dynamics to determine subspaces with distinct confidence levels (Section 3.2). Subsequently, it incorporates a Mixup-based regularization (Section 3.3) to mitigate overfitting to confident examples and propagate their supervision signals to promote the learning of relatively unconfident ones. We illustrate the overall pipeline of SFTMIX in Figure 1.

#### 3.1 PRELIMINARIES OF THE CONVENTIONAL NTP INSTRUCTION-TUNING PARADIGM

Consider an SFT dataset  $\mathcal{D} = \{(\mathcal{X}_i, \mathcal{Y}_i)\}_{i=1}^{|\mathcal{D}|}$ , consisting of pairs of instructions  $\mathcal{X}_i$  and desired responses  $\mathcal{Y}_i$ . Here, both  $\mathcal{X}_i$  and  $\mathcal{Y}_i$  are sequences of tokens, where  $\mathcal{X}_i = (x_1, \dots, x_{M_i})$  and  $\mathcal{Y}_i = (y_1, \dots, y_{N_i})$ . The conventional NTP task minimizes the following loss for predicting  $\mathcal{Y}_i$  given  $\mathcal{X}_i$ :

$$\ell_{\text{NTP}}(\mathcal{D}) = - \sum_{i=1}^{|\mathcal{D}|} \sum_{n=1}^{N_i} \log p(y_n | \mathcal{X}_i, y_1, \dots, y_{n-1}) = - \sum_{i=1}^{|\mathcal{D}|} \sum_{n=1}^{N_i} H(\mathbf{Y}_n, \sigma(\mathbf{Z}_n)). \quad (1)$$

This loss equals the sum of negative cross-entropy  $H$  between  $\mathbf{Y}_n$  and  $\mathbf{Z}_n$  after softmax  $\sigma$ , where  $\mathbf{Y}_n$  is the one-hot encoding of the  $n$ -th token in  $\mathcal{Y}_i$ , and  $\mathbf{Z}_n = \text{LLM}(\mathcal{X}_i, y_1, \dots, y_{n-1})$  is the corresponding representation generated by the LLM’s causal language modeling head.

To improve the effectiveness of SFT, existing work (Zhou et al., 2023; Chen et al., 2024) has primarily focused on providing higher-quality SFT datasets. However, these approaches often lack


162  
163  
164  
165  
166  
167  
168  
169  
170  
171  
172  
173  
174  
175  
176177  
178  
179  
180  
181  
182  
183  
184

Figure 1: The overall pipeline of the proposed SFTMix recipe for LLM instruction tuning. Given an SFT dataset  $\mathcal{D}$ , we (1) train a reference LLM on  $\mathcal{D}$  using NTP and (2) compute  $\text{Conf}(\mathcal{Y}_i | \mathcal{X}_i)$  for  $(\mathcal{X}_i, \mathcal{Y}_i) \in \mathcal{D}$  based on the training dynamics of the reference LLM. On this basis, we (3) divide  $\mathcal{D}$  into a confident subset  $\mathcal{D}^c$  and a relatively unconfident one  $\mathcal{D}^u$  of equal size. Finally, given pairs of examples from each subset, we (4) interpolate their one-hot encodings and predicted representations linearly at the token level and (5) incorporate a Mixup-based regularization  $\ell_{\text{Mixup}}(\mathcal{D}^c, \mathcal{D}^u)$  in addition to the NTP loss  $\ell_{\text{NTP}}(\mathcal{D})$  during LLM instruction tuning.

185  
186  
187  
188  
189

an insightful understanding of SFT datasets and incur significant computational and labor costs, limiting scalability and performance gains. In response, we propose SFTMix, a novel instruction-tuning recipe in the following sections, which enhances instruction-tuning performance beyond the conventional NTP paradigm without relying on well-curated datasets.

190

### 3.2 TRAINING DYNAMICS IDENTIFY SUBSETS WITH DISTINCT CONFIDENCE LEVELS

191  
192  
193  
194  
195

Suppose we identify  $C$  checkpoints of a reference LLM when instruction-tuning it using the NTP task in Section 3.1. We aim to capture the training dynamics of the reference LLM by computing its confidence in generating each pair  $(\mathcal{X}_i, \mathcal{Y}_i) \in \mathcal{D}$ . More specifically, we define confidence based on the perplexity of  $\mathcal{Y}_i$  given  $\mathcal{X}_i$  at each checkpoint  $c \in \{1, \dots, C\}$ :

196  
197  
198

$$\text{Perp}_c(\mathcal{Y}_i | \mathcal{X}_i) = -\frac{1}{N_i} \sum_{n=1}^{N_i} \log p(y_n | \mathcal{X}_i, y_1, \dots, y_{n-1}) = -\frac{1}{N_i} \sum_{n=1}^{N_i} \mathbf{Y}_n^\top \log \sigma(\mathbf{Z}_n), \quad (2)$$

199  
200  
201

$$\text{Conf}(\mathcal{Y}_i | \mathcal{X}_i) = -\frac{1}{C} \sum_{c=1}^C \text{Perp}_c(\mathcal{Y}_i | \mathcal{X}_i). \quad (3)$$

202  
203  
204  
205

Note that  $\mathbf{Z}_n$  here is produced by the reference LLM at checkpoint  $c$ . The reference LLM’s confidence in predicting  $\mathcal{Y}_i$  given  $\mathcal{X}_i$  is the negative average perplexity over the  $C$  checkpoints, as a lower perplexity indicates a higher likelihood of generation.

206  
207  
208  
209  
210  
211  
212  
213  
214  
215

**LLMs Exhibit Uneven Confidence across the Semantic Representation Space.** Here, we present a case study by instruction-tuning Llama-3.1-8B (Dubey et al., 2024) on Alpaca-52K (Taori et al., 2023) and collect the LLM’s confidence for each training data point across five checkpoints. We use the last hidden state of the final token in  $(\mathcal{X}_i, \mathcal{Y}_i)$  as its embedding and plot 2,500 high-confidence and 2,500 low-confidence data points in Figure 2 (a) via t-SNE (Van der Maaten & Hinton, 2008). Correspondingly, we present one confident and one unconfident example from these data points in Figure 2 (b). We observe that embeddings of data points with contrasting confidence levels are clearly separated in Figure 2 (a), indicating that the distribution of the LLM’s confidence is uneven across the semantic representation space. This observation is further supported by the examples in Figure 2 (b), where the LLM exhibits high confidence in the example discussing deterministic grammar rules and low confidence in the example concerning creative content in e-commerce.




Figure 2: (a) Embeddings of 2,500 high-confidence and 2,500 low-confidence examples in Alpaca-52K by Llama-3.1-8B trained on Alpaca-52k using NTP. (b) Corresponding confident and unconfident examples. The clear separation between embeddings of high-confidence and low-confidence examples suggests that the LLM exhibits varying confidence levels across different semantic regions, as further illustrated by the different topics in the provided examples.

**Data with Distinct Confidence Levels Should Play Different Roles during Instruction Tuning.**

The insight from the case study motivates us to contend that data with varying confidence levels should contribute differently during instruction tuning. Highly confident data points typically lie further from the classification decision boundary, posing a higher risk of overfitting. In contrast, less confident data points are often closer to the boundary, making them harder to learn. To address this, we propose promoting the flow of supervision signals between confident and less confident regions to mitigate overfitting and enhance generalization during LLM instruction tuning. On this basis, we divide the original SFT dataset  $\mathcal{D}$  into a confident subset  $\mathcal{D}^c$  and a relatively unconfident subset  $\mathcal{D}^u$  of equal size according to  $\text{Conf}(\mathcal{Y}_i | \mathcal{X}_i)$ . To foster synergy between them, we design a Mixup-based regularization tailored to the specific challenges of instruction tuning, detailed in the next section.

3.3 A MIXUP-BASED REGULARIZATION FACILITATES LLM INSTRUCTION TUNING

To instruction-tune an LLM (different from the reference LLM used to obtain learning dynamics) with our SFTMix recipe, we introduce a novel regularization  $\ell_{\text{Mixup}}$  in addition to the conventional NTP loss  $\ell_{\text{NTP}}$ . Specifically, consider a confident instruction-response pair  $(\mathcal{X}_i^c, \mathcal{Y}_i^c) \in \mathcal{D}^c$  and a relatively unconfident pair  $(\mathcal{X}_i^u, \mathcal{Y}_i^u) \in \mathcal{D}^u$ . Let  $\mathbf{Y}_n^c$  and  $\mathbf{Y}_n^u$  be the one-hot encoding vectors of the  $n$ -th token in  $\mathcal{Y}^c$  and  $\mathcal{Y}^u$ , respectively, with  $\mathbf{Z}_n^c$  and  $\mathbf{Z}_n^u$  as the corresponding representations predicted by the instruction-tuning LLM. We linearly interpolate the two pairs as follows:

$$\tilde{\mathbf{Z}}_n = \lambda \mathbf{Z}_n^c + (1 - \lambda) \mathbf{Z}_n^u, \quad \tilde{\mathbf{Y}}_n = \lambda \mathbf{Y}_n^c + (1 - \lambda) \mathbf{Y}_n^u, \tag{4}$$

where  $\lambda \sim \text{Beta}(\alpha, \alpha)$  and  $\alpha$  is a hyperparameter. Suppose that  $N'_i = \min(N_i^c, N_i^u)$  represents the length of the shorter response between  $\mathcal{Y}_i^c$  and  $\mathcal{Y}_i^u$ . We define the Mixup-based regularization  $\ell_{\text{Mixup}}(\mathcal{D}^c, \mathcal{D}^u)$  between the confident and relatively unconfident subsets and the overall instruction-tuning loss  $\ell_{\text{SFTMix}}$  used in our SFTMix recipe as follows:

$$\ell_{\text{Mixup}}(\mathcal{D}^c, \mathcal{D}^u) = - \sum_{i=1}^{|D|/2} \sum_{n=1}^{N'_i} H(\tilde{\mathbf{Y}}_n, \sigma(\tilde{\mathbf{Z}}_n)), \quad \ell_{\text{SFTMix}}(\mathcal{D}) = \ell_{\text{NTP}}(\mathcal{D}) + \mu \ell_{\text{Mixup}}(\mathcal{D}^c, \mathcal{D}^u). \tag{5}$$

Here,  $\mu$  is a hyperparameter to control the regularization effect.

| Dataset        | Instruction-Tuning LLM | Recipe | MT-Bench      |               |               | AlpacaEval-2  |                |
|----------------|------------------------|--------|---------------|---------------|---------------|---------------|----------------|
|                |                        |        | Single-Turn   | Multi-Turn    | Overall       | Win Rate      | LC Win Rate    |
| Alpaca-52K     | Llama-3.1-8B           | NTP    | 4.9100        | 3.8150        | 4.3625        | 4.0714        | 8.6528         |
|                |                        | SFTMix | <b>5.2125</b> | <b>3.9525</b> | <b>4.5825</b> | <b>4.9031</b> | <b>10.3195</b> |
| UltraChat-200K | Mistral-7B-v0.1        | NTP    | 5.1650        | 4.0675        | 4.6163        | 4.3560        | 9.1759         |
|                |                        | SFTMix | <b>5.2775</b> | <b>4.5425</b> | <b>4.9100</b> | <b>4.5386</b> | <b>9.4994</b>  |
| Alpaca-52K     | Llama-3.1-8B           | NTP    | 6.1875        | 5.0125        | 5.6000        | 5.0665        | 8.4505         |
|                |                        | SFTMix | <b>6.2750</b> | <b>5.3500</b> | <b>5.8125</b> | <b>5.1149</b> | <b>9.3810</b>  |
| UltraChat-200K | Mistral-7B-v0.1        | NTP    | 5.7625        | 4.6938        | 5.2281        | 4.4899        | 7.7732         |
|                |                        | SFTMix | <b>5.9813</b> | <b>4.8813</b> | <b>5.4313</b> | <b>4.6117</b> | <b>8.7650</b>  |

Table 1: Evaluation of instruction-following capabilities of LLMs trained with NTP or SFTMix. We report the average score based on five rounds of evaluation, with the scores from the best-performing instruction-tuning recipe in bold. Standard errors are provided in Appendix A. SFTMix outperforms NTP on both MT-Bench and AlpacaEval-2, irrespective of LLM families and SFT data sizes.

Altogether, as shown in Figure 1, our SFTMix recipe first identifies subspaces with distinct confidence levels using training dynamics, then facilitates the propagation of supervision signals between these subspaces through a Mixup-based regularization. Since LLMs exhibit varying confidence levels across diverse semantic regions, the Mixup-based regularization encourages linear behavior and a smoother decision boundary between confident and unconfident data points (Zhang et al., 2018; Verma et al., 2019). In this way, SFTMix regularizes overfitting in confident regions and propagates supervision signals (Bengio et al., 2009; Chapelle et al., 2009; Sohn et al., 2020) to enhance generalization in less confident regions, thereby improving the instruction-tuning performance.

## 4 EXPERIMENTS

In this section, we assess the effectiveness of SFTMix against the NTP baseline in both instruction-following (Section 4.1) and domain-specific (Section 4.2) SFT tasks. SFTMix consistently improves instruction-tuning performance across different LLM families and SFT datasets of varying scales.

### 4.1 INSTRUCTION-FOLLOWING SFT

Instruction-following SFT trains LLMs on labeled datasets consisting of instructional prompts and corresponding desired responses, enhancing their conversational capabilities in downstream interaction-driven applications. Here, we compare SFTMix with the conventional NTP paradigm by applying them to the instruction tuning of two pre-trained LLMs from different model families (i.e., Llama (Dubey et al., 2024) and Mistral (Jiang et al., 2023)) on two instruction-following datasets of varying scales (i.e., Alpaca-52K (Taori et al., 2023) and UltraChat-200K (Tunstall et al., 2023)). We then evaluate the instruction-tuned LLMs on two widely-adopted benchmarks: MT-Bench (Zheng et al., 2024) and AlpacaEval-2 (Dubois et al., 2024).

**Datasets.** Here, we focus on the following datasets with different sizes.

- Alpaca-52K (Taori et al., 2023) builds on the pipeline from Wang et al. (2023) by prompting text-davinci-003 (Ouyang et al., 2022) to generate diverse instructions and appropriate responses, which results in 52,000 single-turn interactions. We follow its default system prompt and conversation template when preparing training inputs.
- UltraChat-200K (Tunstall et al., 2023) filters out uninformative responses from the original UltraChat dataset (Ding et al., 2023) and downsamples it to 200,000 multi-turn interactions. To adapt these for our SFT pipeline, we expand each multi-turn interaction into multiple single-turn interactions by incorporating the chat history into the instructions.

**Implementation Details.** We experiment with Llama-3.1-8B (Dubey et al., 2024) and Mistral-7B-v0.1 (Jiang et al., 2023) due to their recent release and state-of-the-art performance compared to other models of similar sizes. By default, we use different instances of the same LLM type to obtain training dynamics and for instruction tuning. We determine the instruction-tuning hyperparameters through a coarse sweep on Llama-3.1-8B with Alpaca-52K and adopt them as the default settings in our experiments. Specifically, we train each LLM on Alpaca-52K for three epochs and on UltraChat-200K for one epoch, leveraging eight H100 GPUs. We use the AdamW optimizer with a learning

324  
325  
326  
327  
328  
329  
330  
331  
332  
333  
334  
335  
336  
337  
338  
339  
340  
341  
342  
343  
344  
345  
346  
347  
348  
349  
350  
351  
352  
353  
354  
355  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377

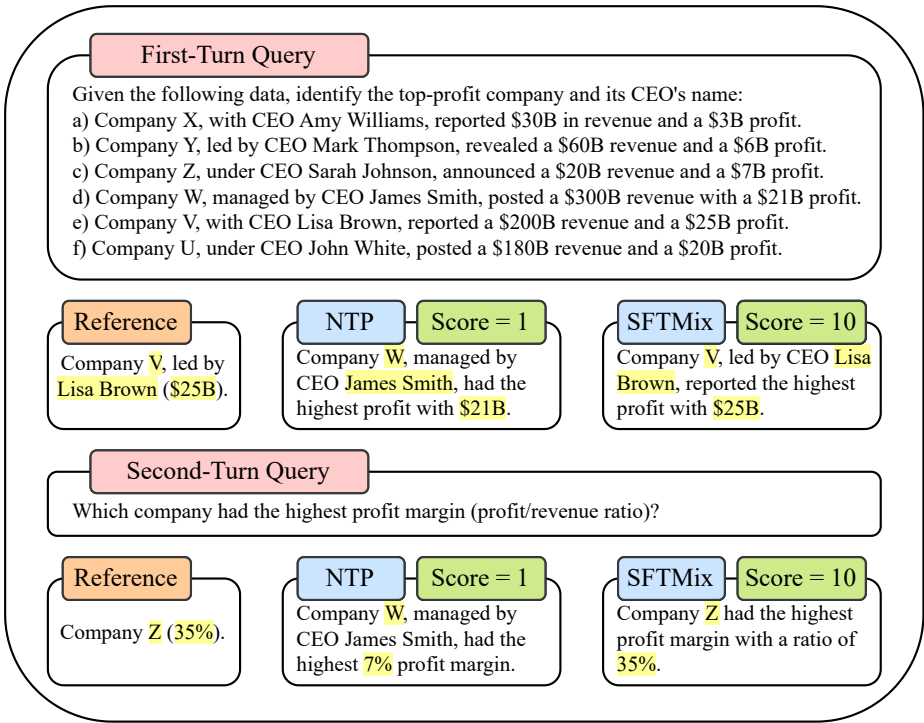


Figure 3: Examples from the extraction category in MT-Bench. Compared to its NTP-tuned counterpart, Llama-3.1-8B instruction-tuned on Alpaca-52K using SFTMix accurately interprets the queries from both turns and correctly extracts the relevant information from the prompt.

rate of  $2e-6$  and weight decay of 0.1, along with a cosine learning rate scheduler featuring a 0.1 warm-up ratio. Each gradient update during instruction tuning accumulates four batches, with each batch containing eight training examples. We set  $\alpha = 0.5$  for sampling  $\lambda$  in Equation 4 and  $\mu = 0.2$  when constructing  $\ell_{\text{SFTMix}}$  in Equation 5. The NTP baseline follows the same instruction-tuning setting without the Mixup-based regularization  $\ell_{\text{Mixup}}$  in Equation 5.

**Evaluation Benchmarks.** To evaluate how SFTMix improves LLMs’ instruction-following abilities compared to NTP, we assess LLMs instruction-tuned with each method on MT-Bench (Zheng et al., 2024) and AlpacaEval-2 (Li et al., 2023; Dubois et al., 2024). MT-Bench is a challenging benchmark of 80 multi-turn, human-designed questions, where GPT-4-Turbo (Achiam et al., 2023) rates the quality of LLM-generated responses on a ten-point scale. Similarly, AlpacaEval-2 employs GPT-4-Turbo to compare the tested LLMs’ responses against GPT-4-Turbo’s reference responses and calculates the expected win rate while adjusting for length bias. We repeat the evaluation five times for each setting and report the average score in Table 1 and standard error in Appendix A.

**SFTMix Enhances LLMs’ Instruction-Following Capabilities.** As illustrated in Table 1, instruction-tuning with SFTMix consistently outperforms NTP across all metrics in both evaluation benchmarks, regardless of the base LLM or SFT dataset. Notably, SFTMix yields a greater improvement in multi-turn conversational abilities (with an average increase of 0.3 points) compared to single-turn performance (an average increase of 0.2 points) in MT-Bench. Across the eight categories in MT-Bench, we observe significant gains in extraction tasks for Llama-3.1-8B, and in writing, coding, and STEM for Mistral-7B-v0.1 (full details in Appendix A). In AlpacaEval-2, the improvement is particularly significant in the length-controlled (LC) win rate, which better aligns with human judgment by adjusting for GPT-4-Turbo’s preference for longer responses. While instruction-tuning with the larger, higher-quality UltraChat-200K dataset results in higher overall scores in MT-Bench and raw win rates in AlpacaEval-2, it also produces longer responses, leading to relatively lower LC win rates. Overall, our proposed recipe, SFTMix, enhances instruction-

| LLM                         | MedQA        | MedQA-5      | PubMedQA     | MedMCQA      | Macro Ave    |
|-----------------------------|--------------|--------------|--------------|--------------|--------------|
| MedAlpaca-7B                | 38.94        | 33.96        | 57.20        | 34.90        | 41.25        |
| PMC-LLaMA-7B                | 27.94        | 21.24        | 54.87        | 24.57        | 32.16        |
| BioMedGPT-LM-7B             | 38.62        | 34.72        | 58.27        | 35.57        | 41.80        |
| Meditron-7B                 | 35.09        | 26.73        | 56.93        | 34.03        | 38.20        |
| BioMistral-7B               | 43.86        | 37.58        | 50.13        | 44.14        | 43.93        |
| Llama-3.1-8B                | 59.68        | 53.23        | 73.40        | 52.79        | 59.78        |
| + NTP on MedAlpaca-263K     | 59.31        | 54.52        | 75.40        | 53.65        | 60.72        |
| or SFTMix on MedAlpaca-263K | <b>60.88</b> | <b>55.38</b> | <b>77.80</b> | <b>54.15</b> | <b>62.05</b> |
| Mistral-7B-v0.1             | 49.18        | 43.94        | 72.33        | 47.98        | 53.36        |
| + NTP on MedAlpaca-263K     | 49.10        | 44.62        | 75.40        | 48.15        | 54.32        |
| or SFTMix on MedAlpaca-263K | <b>51.77</b> | <b>45.72</b> | <b>77.40</b> | <b>49.03</b> | <b>55.98</b> |

Table 2: Evaluation results on four healthcare-related benchmarks by prior biomedical LLMs and LLMs trained on MedAlpaca-263K using either NTP or SFTMix. We report the mean accuracy (%) over three rounds of three-shot evaluation and bold the scores from SFTMix-tuned LLMs. Standard errors are provided in Appendix A. SFTMix achieves an approximate 1.5% absolute increase in macro-average accuracy compared to NTP for both Llama-3.1-8B and Mistral-7B-v0.1.

following capabilities and text quality across LLMs from different model families and with datasets of varying sizes and quality, surpassing the performance of the conventional NTP paradigm.

**Case Study from MT-Bench.** Figure 3 presents a test example from the extraction category in MT-Bench, showing responses generated by Llama-3.1-8B instruction-tuned on Alpaca-52K using either NTP or SFTMix. In this example, the LLM trained with SFTMix accurately interprets the instructions from both the first- and second-turn queries, correctly extracting the relevant information from the prompt. Notably, it succeeds in answering the second-turn query, which involves calculating the profit margin and performing a ratio comparison. In contrast, the LLM trained with NTP struggles to differentiate between revenue and profit, leading to incorrect responses in both turns.

#### 4.2 DOMAIN-SPECIFIC SFT

In healthcare domain-specific SFT, we train two LLMs on a large-scale medical conversation dataset using SFTMix and assess their performance against NTP-tuned counterparts on four healthcare-related question-answering benchmarks.

**Dataset.** MedAlpaca-263K (Han et al., 2023) consists of medical NLP tasks reformatted for instruction tuning and healthcare-related conversations of varying quality crowd-sourced from online platforms, which amounts to a total of 263,257 single-turn interactions. We train Llama-3.1-8B and Mistral-7B-v0.1 on MedAlpaca-263K using either NTP or SFTMix for two epochs and follow the remaining hyperparameter settings described in Section 4.1.

**Evaluation Benchmarks.** We compare the effectiveness of SFTMix to NTP in domain-specific SFT by evaluating the performance of the instruction-tuned LLMs on the following benchmarks:

- MedQA (Jin et al., 2021) includes 1,273 four-choice questions from the US Medical License Exam, testing a broad range of medical knowledge.
- MedQA-5 is the variant of MedQA where each question contains five options.
- PubMedQA (Jin et al., 2019) consists of 500 expert-labeled three-choice questions where the model must predict the answer by reasoning based on a provided PubMed abstract.
- MedMCQA (Pal et al., 2022) comprises 4,183 four-choice questions from the Indian Medical Entrance Exams, covering 2,400 healthcare-related topics across 21 medical subjects.

We adopt the three-shot evaluation setting from Labrak et al. (2024) and report the mean accuracy over three evaluation rounds in Table 2. Standard errors are provided in Appendix A. Additionally, we include prior biomedical LLMs of similar sizes, including MedAlpaca-7B (Han et al., 2023), PMC-LLaMA-7B (Wu et al., 2024), BioMedGPT-LM-7B (Luo et al., 2023), Meditron-7B (Chen et al., 2023), and BioMistral-7B (Labrak et al., 2024), as reference models for comparison.



| Ablation Direction | Loss $\ell = \ell_{\text{NTP}} + \mu \ell_{\text{Mixup}}$ |   | MT-Bench |        |         | AlpacaEval-2 |         |
|--------------------|---|---|----------|--------|---------|--------------|---------|
|                    | NTP Loss  | Mixup Regularization                                    | ST       | MT     | Overall | WR           | LC WR   |
| NTP                | $\ell_{\text{NTP}}(\text{Full})$                          | -   | 4.9100   | 3.8150 | 4.3625  | 4.0714       | 8.6528  |
| SFTMix             | $\ell_{\text{NTP}}(\text{Full})$                          | $\mu \ell_{\text{Mixup}}(\text{Conf}, \text{Unconf})$   | 5.2125   | 3.9525 | 4.5825  | 4.9031       | 10.3195 |
| Section 5.1        | $\ell_{\text{NTP}}(\text{Full})$                          | $\mu \ell_{\text{Mixup}}(\text{Conf}', \text{Unconf}')$ | 4.8500   | 4.2625 | 4.5563  | 4.5786       | 10.0483 |
| Section 5.2        | $\ell_{\text{NTP}}(\text{Full})$                          | $\ell_{\text{Mixup}}(\text{Conf}, \text{Unconf})$       | 4.7050   | 4.1075 | 4.4062  | 3.9450       | 8.2856  |
|                    | -   | $\ell_{\text{Mixup}}(\text{Conf}, \text{Unconf})$       | 5.0125   | 4.0000 | 4.5062  | 3.5821       | 7.2964  |
| Section 5.3        | $\ell_{\text{NTP}}(\text{Conf})$                          | $\mu \ell_{\text{Mixup}}(\text{Conf}, \text{Unconf})$   | 4.9775   | 4.1075 | 4.5425  | 4.4496       | 9.7824  |
|                    | $\ell_{\text{NTP}}(\text{Unconf})$                        | $\mu \ell_{\text{Mixup}}(\text{Conf}, \text{Unconf})$   | 5.1800   | 3.9050 | 4.5425  | 4.2030       | 8.9392  |
| Section 5.4        | $\ell_{\text{NTP}}(\text{High})$                          | -   | 6.1175   | 5.2575 | 5.6875  | 7.2636       | 11.4490 |
|                    | $\ell_{\text{NTP}}(\text{High} + \text{Low})$             | -   | 5.9000   | 5.1825 | 5.5412  | 6.5871       | 11.9590 |
|                    | $\ell_{\text{NTP}}(\text{High} + \text{Low})$             | $\mu \ell_{\text{Mixup}}(\text{High}, \text{Low})$      | 5.8025   | 5.0975 | 5.4500  | 5.9382       | 11.1768 |

Table 3: Ablation studies on variants of SFTMix. We identify four ablation directions and evaluate the instruction-following abilities of Llama-3.1-8B trained with the corresponding loss functions. Conf and Unconf are the confident and unconfident subsets of the Full Alpaca-52K dataset, with confidence derived from the training dynamics of Llama-3.1-8B. Conf' and Unconf' are based on Gemma-2B's training dynamics. High refers to Alpaca-GPT4-26K (higher quality), while Low refers to Alpaca-26K (relatively lower quality).

**SFTMix Adapts LLMs to Domain-Specific Tasks More Effectively.** Table 2 demonstrates that SFTMix consistently surpasses NTP across all benchmarks for both backbones. In particular, SFTMix leads to a 1.33% absolute improvement (from 60.72% to 62.05%) for Llama-3.1-8B and a 1.66% increase (from 54.32% to 55.98%) for Mistral-7B-v0.1 in macro-average accuracy across the four benchmarks. These models also significantly outperform existing biomedical LLMs across all benchmarks by a clear margin.

## 5 ABLATION AND ANALYSIS

Following the improvements of SFTMix in instruction-following and domain-specific SFT tasks, we conduct extensive ablation studies to analyze the contribution of each design choice and explore its impact across applications. We identify four ablation directions and summarize the results of training Llama-3.1-8B on Alpaca-52K using variants of SFTMix in Table 3.

### 5.1 GENERALIZING THE TRAINING DYNAMICS FROM A WEAKER REFERENCE LLM

Inspired by Burns et al. (2024), we investigate the generalization of training dynamics from a weaker reference LLM to a stronger instruction-tuning LLM. Specifically, we identify training dynamics with a weaker reference LLM, Gemma-2B (Team et al., 2024), to divide an SFT dataset into a confident subset (Conf') and a relatively unconfident subset (Unconf'). These subsets are then fed into the Mixup regularization  $\ell_{\text{Mixup}}(\text{Conf}', \text{Unconf}')$  when instruction-tuning Llama-3.1-8B. This alternative approach yields comparable scores on MT-Bench and AlpacaEval-2 to the original SFTMix recipe, which uses the same LLM for both training dynamics and Mixup-based instruction tuning. This finding aligns with the weak-to-strong generalization reported by Burns et al. (2024) and highlights the potential for scaling SFTMix to even stronger LLMs.

### 5.2 INCORPORATING MIXUP AS A REGULARIZATION IS MORE EFFECTIVE

Equation 5 uses the Mixup regularization  $\ell_{\text{Mixup}}$  to alleviate overfitting and encourage generalization. To fully explore its effect, we experiment with setting  $\mu = 1$  in Equation 5 (i.e.,  $\ell = \ell_{\text{NTP}}(\text{Full}) + \ell_{\text{Mixup}}(\text{Conf}, \text{Unconf})$ ) or only minimizing  $\ell_{\text{Mixup}}$  without  $\ell_{\text{NTP}}$  (i.e.,  $\ell = \ell_{\text{Mixup}}(\text{Conf}, \text{Unconf})$ ) during instruction tuning. Table 3 shows that these two variants achieve higher scores on MT-Bench but perform worse on AlpacaEval-2 compared to the baseline of using the conventional NTP method. Furthermore, our SFTMix recipe, which employs  $\ell_{\text{Mixup}}$  as a regularization, still outperforms both variants across both benchmarks. This finding highlights the importance of incorporating the tradi-

486 tional NTP task during SFT and supports the conclusion that Mixup is more effective when used as  
 487 a regularization alongside the standard cross-entropy loss in LLM instruction tuning.  
 488

### 489 5.3 SFTMIX EFFECTIVELY UTILIZES ENTIRE INSTRUCTION-TUNING DATASETS 490

491 As part of our SFTMix recipe, we apply the NTP loss  $\ell_{\text{NTP}}$  to the entire SFT dataset  $\mathcal{D}$ . Here, we  
 492 consider variants where  $\ell_{\text{NTP}}$  is applied selectively to either the confident or relatively unconfident  
 493 halves of the dataset. Specifically, we experiment with  $\ell = \ell_{\text{NTP}}(\text{Conf}) + \mu \ell_{\text{Mixup}}(\text{Conf}, \text{Unconf})$  and  
 494  $\ell = \ell_{\text{NTP}}(\text{Unconf}) + \mu \ell_{\text{Mixup}}(\text{Conf}, \text{Unconf})$ . As shown in Table 3, while both variants achieve the  
 495 same overall score on MT-Bench, the variant applying  $\ell_{\text{NTP}}$  to the confident subset (Conf) performs  
 496 better on AlpacaEval-2. Notably, both variants—where  $\ell_{\text{NTP}}$  is applied to only half the dataset—  
 497 outperform the baseline where  $\ell_{\text{NTP}}$  is applied to the entire dataset. We attribute this improvement to  
 498 the impact introduced by our Mixup regularization  $\ell_{\text{Mixup}}$ . Nevertheless, our SFTMix recipe, which  
 499 leverages the full dataset for NTP, outperforms both variants, demonstrating its ability to effectively  
 500 utilize a larger set of potentially lower-quality training examples during instruction tuning.

### 501 5.4 TRAINING DYNAMICS ARE CRUCIAL FOR PERFORMING MIXUP 502

503 Building on the previous ablation study, which suggests the possi-  
 504 bility of generalizing training dynamics from a weaker LLM in  
 505 our SFTMix recipe, we explore whether we can directly substi-  
 506 tute training dynamics with known data quality. To test this hy-  
 507 pothesis, we replace half of the original responses in Alpaca-52K  
 508 with higher-quality GPT-4-generated versions, forming Alpaca-  
 509 GPT4-26K (High) (Peng et al., 2023), while referring to the re-  
 510 maining original responses as Alpaca-26K (Low). We then train  
 511 Llama-3.1-8B using three approaches: NTP on Alpaca-GPT4-  
 512 26K ( $\ell_{\text{NTP}}(\text{High})$ ), NTP on its combination with Alpaca-26K  
 513 ( $\ell_{\text{NTP}}(\text{High} + \text{Low})$ ), and, in the final approach, the addition of the  
 514 Mixup-based regularizer between Alpaca-GPT4-26K and Alpaca-  
 515 26K ( $\ell_{\text{Mixup}}(\text{High}, \text{Low})$ ). The use of higher-quality responses from  
 516 GPT-4 indeed enhances instruction-tuning performance on both  
 517 MT-Bench and AlpacaEval-2, as shown in Table 3. However, sim-  
 518 ply applying Mixup between two datasets of varying quality does  
 519 not necessarily improve performance further, as indicated by the  
 520 drop in the overall MT-Bench score from 5.5412 to 5.4500 and the  
 521 length-controlled win rate in AlpacaEval-2 from 11.9590 to 11.1768. To investigate this observation,  
 522 we plot the LLM’s confidence distributions for both datasets in Figure 4. The substantial overlap  
 523 in confidence distributions suggests that data quality does not necessarily correlate with training  
 524 dynamics-based confidence. This highlights the importance of training dynamics in determining the  
 525 model-specific role of data points, which is crucial for effectively applying our SFTMix recipe.




Figure 4: Distribution of confidence in datasets of varying qualities by Llama-3.1-8B.

## 526 6 CONCLUSION 527

528 In this paper, we propose SFTMix, a novel recipe for LLM instruction tuning. We observe that  
 529 LLMs exhibit uneven confidence distributions across the semantic representation space. Based on  
 530 this motivation, we utilize training dynamics to identify data subsets of varying confidence levels  
 531 and incorporate a Mixup-based regularization. In this way, we aim to mitigate overfitting on the  
 532 confident subset while propagating supervision signals to promote the generalization of the relatively  
 533 unconfident subset. Extensive empirical results in both instruction-following and domain-specific  
 534 SFT tasks demonstrate the effectiveness of SFTMix over the conventional NTP paradigm across  
 535 different LLM families and SFT data scales. Comprehensive ablation studies further substantiate  
 536 the contribution of SFTMix’s design choices, highlighting its versatility in consistently enhancing  
 537 performance across different LLMs and datasets in broader NLP applications. Due to computational  
 538 constraints, we did not apply SFTMix to LLM pre-training or instruction-tune larger LLMs using  
 539 this recipe. Integrating SFTMix with parameter-efficient pre-training and fine-tuning methods (Hu  
 et al., 2022; Detmers et al., 2024) is another promising direction for future work.

## REFERENCES

- 540  
541  
542 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Ale-  
543 man, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical  
544 report. *arXiv preprint arXiv:2303.08774*, 2023.
- 545 Alon Albalak, Yanai Elazar, Sang Michael Xie, Shayne Longpre, Nathan Lambert, Xinyi Wang,  
546 Niklas Muennighoff, Bairu Hou, Liangming Pan, Haewon Jeong, et al. A survey on data selection  
547 for language models. *arXiv preprint arXiv:2402.16827*, 2024.
- 548 Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In  
549 *Proceedings of the 26th annual international conference on machine learning*, pp. 41–48, 2009.  
550
- 551 David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A  
552 Raffel. Mixmatch: A holistic approach to semi-supervised learning. *Advances in neural informa-*  
553 *tion processing systems*, 32, 2019.
- 554 David Berthelot, Nicholas Carlini, Ekin D. Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and  
555 Colin Raffel. Remixmatch: Semi-supervised learning with distribution matching and augmenta-  
556 tion anchoring. In *International Conference on Learning Representations*, 2020.  
557
- 558 Collin Burns, Pavel Izmailov, Jan Hendrik Kirchner, Bowen Baker, Leo Gao, Leopold Aschenbren-  
559 ner, Yining Chen, Adrien Ecoffet, Manas Joglekar, Jan Leike, Ilya Sutskever, and Jeffrey Wu.  
560 Weak-to-strong generalization: Eliciting strong capabilities with weak supervision. In *Forty-first*  
561 *International Conference on Machine Learning*, 2024.
- 562 Luigi Carratino, Moustapha Cissé, Rodolphe Jenatton, and Jean-Philippe Vert. On mixup regular-  
563 ization. *Journal of Machine Learning Research*, 23(325):1–31, 2022.
- 564 Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. Semi-supervised learning (chapelle, o.  
565 et al., eds.; 2006)[book reviews]. *IEEE Transactions on Neural Networks*, 20:542–542, 2009.  
566
- 567 Jiaao Chen, Zichao Yang, and Diyi Yang. Mixtext: Linguistically-informed interpolation of hidden  
568 space for semi-supervised text classification. In *Proceedings of the 58th Annual Meeting of the*  
569 *Association for Computational Linguistics*, pp. 2147–2157, 2020.
- 570 Lichang Chen, Shiyang Li, Jun Yan, Hai Wang, Kalpa Gunaratna, Vikas Yadav, Zheng Tang, Vijay  
571 Srinivasan, Tianyi Zhou, Heng Huang, and Hongxia Jin. Alpargasus: Training a better alpaca with  
572 fewer data. In *The Twelfth International Conference on Learning Representations*, 2024.
- 573 Zeming Chen, Alejandro Hernández Cano, Angelika Romanou, Antoine Bonnet, Kyle Matoba,  
574 Francesco Salvi, Matteo Pagliardini, Simin Fan, Andreas Köpf, Amirkeivan Mohtashami,  
575 et al. Meditron-70b: Scaling medical pretraining for large language models. *arXiv preprint*  
576 *arXiv:2311.16079*, 2023.
- 577 Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng,  
578 Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An  
579 open-source chatbot impressing gpt-4 with 90%\* chatgpt quality, March 2023. URL <https://lmsys.org/blog/2023-03-30-vicuna/>.  
580  
581
- 582 Muthu Chidambaram, Xiang Wang, Yuzheng Hu, Chenwei Wu, and Rong Ge. Towards under-  
583 standing the data dependency of mixup-style training. In *International Conference on Learning*  
584 *Representations*, 2022.
- 585 Everlyn Chimoto, Jay Gala, Orevaoghene Ahia, Julia Kreutzer, Bruce Bassett, and Sara Hooker.  
586 Critical learning periods: Leveraging early training dynamics for efficient data pruning. In *Find-*  
587 *ings of the Association for Computational Linguistics ACL 2024*, pp. 9407–9426, 2024.
- 588 Hyeong Kyu Choi, Joonmyung Choi, and Hyunwoo J. Kim. Tokenmixup: Efficient attention-guided  
589 token-level data augmentation for transformers. In *Advances in Neural Information Processing*  
590 *Systems*, 2022.  
591
- 592 Fenia Christopoulou, Gerasimos Lampouras, and Ignacio Iacobacci. Training dynamics for cur-  
593 riculum learning: A study on monolingual and cross-lingual nlu. In *Proceedings of the 2022*  
*Conference on Empirical Methods in Natural Language Processing*, pp. 2595–2611, 2022.

- 594 Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning  
595 of quantized llms. *Advances in Neural Information Processing Systems*, 36, 2024.  
596
- 597 Ning Ding, Yulin Chen, Bokai Xu, Yujia Qin, Shengding Hu, Zhiyuan Liu, Maosong Sun, and  
598 Bowen Zhou. Enhancing chat language models by scaling high-quality instructional conversa-  
599 tions. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Pro-  
600 cessing*, pp. 3029–3051, 2023.
- 601 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha  
602 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models.  
603 *arXiv preprint arXiv:2407.21783*, 2024.  
604
- 605 Yann Dubois, Balázs Galambosi, Percy Liang, and Tatsunori B Hashimoto. Length-controlled al-  
606 pacaeval: A simple way to debias automatic evaluators. *arXiv preprint arXiv:2404.04475*, 2024.
- 607 Demi Guo, Yoon Kim, and Alexander M Rush. Sequence-level mixed sample data augmentation.  
608 In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing  
609 (EMNLP)*, pp. 5547–5552, 2020.  
610
- 611 Tianyu Han, Lisa C Adams, Jens-Michalis Papaioannou, Paul Grundmann, Tom Oberhauser,  
612 Alexander Löser, Daniel Truhn, and Keno K Bressen. Medalpaca—an open-source collection  
613 of medical conversational ai models and training data. *arXiv preprint arXiv:2304.08247*, 2023.
- 614 Muyang He, Shuo Yang, Tiejun Huang, and Bo Zhao. Large-scale dataset pruning with dynamic  
615 uncertainty. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recog-  
616 nition*, pp. 7713–7722, 2024.
- 617 Dan Hendrycks, Norman Mu, Ekin Dogus Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshmi-  
618 narayanan. Augmix: A simple method to improve robustness and uncertainty under data shift. In  
619 *International conference on learning representations*, 2020.  
620
- 621 Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,  
622 and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Con-  
623 ference on Learning Representations*, 2022.
- 624 Neel Jain, Ping-yeh Chiang, Yuxin Wen, John Kirchenbauer, Hong-Min Chu, Gowthami Somepalli,  
625 Brian R Bartoldson, Bhavya Kailkhura, Avi Schwarzschild, Aniruddha Saha, et al. Neftune:  
626 Noisy embeddings improve instruction finetuning. In *The Twelfth International Conference on  
627 Learning Representations*, 2024.  
628
- 629 Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot,  
630 Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al.  
631 Mistral 7b. *arXiv preprint arXiv:2310.06825*, 2023.
- 632 Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. What dis-  
633 ease does this patient have? a large-scale open domain question answering dataset from medical  
634 exams. *Applied Sciences*, 11(14):6421, 2021.  
635
- 636 Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. Pubmedqa: A dataset  
637 for biomedical research question answering. In *Proceedings of the 2019 Conference on Empirical  
638 Methods in Natural Language Processing and the 9th International Joint Conference on Natural  
639 Language Processing (EMNLP-IJCNLP)*, pp. 2567–2577, 2019.
- 640 Po-Nien Kung, Fan Yin, Di Wu, Kai-Wei Chang, and Nanyun Peng. Active instruction tuning:  
641 Improving cross-task generalization by training on prompt sensitive tasks. In *Proceedings of the  
642 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 1813–1829, 2023.
- 643 Yanis Labrak, Adrien Bazoge, Emmanuel Morin, Pierre-Antoine Gourraud, Mickael Rouvier, and  
644 Richard Dufour. BioMistral: A collection of open-source pretrained large language models for  
645 medical domains. In *Findings of the Association for Computational Linguistics ACL 2024*, 2024.  
646
- 647 Changchun Li, Ximing Li, Lei Feng, and Jihong Ouyang. Who is your right mixup partner in  
positive and unlabeled learning. In *International Conference on Learning Representations*, 2022.

- 648 Junnan Li, Richard Socher, and Steven C.H. Hoi. Dividemix: Learning with noisy labels as semi-  
649 supervised learning. In *International Conference on Learning Representations*, 2020.  
650
- 651 Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy  
652 Liang, and Tatsunori B. Hashimoto. AlpacaEval: An automatic evaluator of instruction-following  
653 models. [https://github.com/tatsu-lab/alpaca\\_eval](https://github.com/tatsu-lab/alpaca_eval), 5 2023.  
654
- 655 Zhenghao Lin, Zhibin Gou, Yeyun Gong, Xiao Liu, Yelong Shen, Ruochen Xu, Chen Lin, Yujiu  
656 Yang, Jian Jiao, Nan Duan, et al. Rho-1: Not all tokens are what you need. *arXiv preprint*  
657 *arXiv:2404.07965*, 2024.
- 658 Yinhan Liu. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint*  
659 *arXiv:1907.11692*, 2019.
- 660 Yizhen Luo, Jiahuan Zhang, Siqi Fan, Kai Yang, Yushuai Wu, Mu Qiao, and Zaiqing Nie.  
661 Biomedgpt: Open multimodal generative pre-trained transformer for biomedicine. *arXiv preprint*  
662 *arXiv:2308.09442*, 2023.  
663
- 664 Shervin Minaee, Tomas Mikolov, Narjes Nikzad, Meysam Chenaghlu, Richard Socher, Xavier Am-  
665 atriain, and Jianfeng Gao. Large language models: A survey. *arXiv preprint arXiv:2402.06196*,  
666 2024.
- 667 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong  
668 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to fol-  
669 low instructions with human feedback. *Advances in neural information processing systems*, 35:  
670 27730–27744, 2022.  
671
- 672 Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. Medmcqa: A large-scale  
673 multi-subject multi-choice dataset for medical domain question answering. In *Conference on*  
674 *health, inference, and learning*, pp. 248–260. PMLR, 2022.
- 675 Chanwoo Park, Sangdoon Yun, and Sanghyuk Chun. A unified analysis of mixed sample data aug-  
676 mentation: A loss function perspective. In *Advances in Neural Information Processing Systems*,  
677 2022.  
678
- 679 Seo Yeon Park and Cornelia Caragea. A data cartography based mixup for pre-trained language  
680 models. In *Proceedings of the 2022 Conference of the North American Chapter of the Association*  
681 *for Computational Linguistics: Human Language Technologies*, pp. 4244–4250, 2022.
- 682 Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. Instruction tuning  
683 with gpt-4. *arXiv preprint arXiv:2304.03277*, 2023.  
684
- 685 Francesco Pinto, Harry Yang, Ser Nam Lim, Philip Torr, and Puneet Dokania. Using mixup as  
686 a regularizer can surprisingly improve accuracy & out-of-distribution robustness. *Advances in*  
687 *Neural Information Processing Systems*, 35:14608–14622, 2022.
- 688 Eduard Poesina, Cornelia Caragea, and Radu Ionescu. A novel cartography-based curriculum learn-  
689 ing method applied on RoNLI: The first Romanian natural language inference corpus. In *Pro-*  
690 *ceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume*  
691 *1: Long Papers)*, pp. 236–253, 2024.  
692
- 693 Stephanie Schoch, Ritwick Mishra, and Yangfeng Ji. Data selection for fine-tuning large language  
694 models using transferred shapley values. In *Proceedings of the 61st Annual Meeting of the As-*  
695 *sociation for Computational Linguistics (Volume 4: Student Research Workshop)*, pp. 266–275,  
696 2023.
- 697 Nabeel Seedat, Nicolas Huynh, Boris van Breugel, and Mihaela van der Schaar. Curated LLM:  
698 Synergy of LLMs and data curation for tabular augmentation in low-data regimes. In *Forty-first*  
699 *International Conference on Machine Learning*, 2024.  
700
- 701 Zhengyan Shi, Adam X Yang, Bin Wu, Laurence Aitchison, Emine Yilmaz, and Aldo Lipani. In-  
instruction tuning with loss over instructions. *arXiv preprint arXiv:2405.14394*, 2024.

- 702 Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel,  
703 Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised  
704 learning with consistency and confidence. *Advances in neural information processing systems*,  
705 33:596–608, 2020.
- 706
- 707 Lichao Sun, Congying Xia, Wenpeng Yin, Tingting Liang, S Yu Philip, and Lifang He. Mixup-  
708 transformer: Dynamic data augmentation for nlp tasks. In *Proceedings of the 28th International  
709 Conference on Computational Linguistics*, pp. 3436–3440, 2020.
- 710 Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi,  
711 Noah A Smith, and Yejin Choi. Dataset cartography: Mapping and diagnosing datasets with  
712 training dynamics. In *Proceedings of the 2020 Conference on Empirical Methods in Natural  
713 Language Processing (EMNLP)*, pp. 9275–9293, 2020.
- 714
- 715 Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy  
716 Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model.  
717 [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca), 2023.
- 718 Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhu-  
719 patiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. Gemma  
720 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*, 2024.
- 721
- 722 Lewis Tunstall, Edward Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada,  
723 Shengyi Huang, Leandro von Werra, Clémentine Fourrier, Nathan Habib, et al. Zephyr: Direct  
724 distillation of lm alignment. *arXiv preprint arXiv:2310.16944*, 2023.
- 725 A F M Shahab Uddin, Mst. Sirazam Monira, Wheemyung Shin, TaeChoong Chung, and Sung-Ho  
726 Bae. Saliencymix: A saliency guided data augmentation strategy for better regularization. In  
727 *International Conference on Learning Representations*, 2021.
- 728
- 729 Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine  
730 learning research*, 9(11), 2008.
- 731 Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-  
732 Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states.  
733 In *International conference on machine learning*, pp. 6438–6447. PMLR, 2019.
- 734
- 735 Jiahao Wang, Bolin Zhang, Qianlong Du, Jiajun Zhang, and Dianhui Chu. A survey on data selection  
736 for llm instruction tuning. *arXiv preprint arXiv:2402.05123*, 2024.
- 737
- 738 Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and  
739 Hannaneh Hajishirzi. Self-instruct: Aligning language models with self-generated instructions. In  
740 *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume  
741 1: Long Papers)*, pp. 13484–13508, 2023.
- 742
- 743 Chaoyi Wu, Weixiong Lin, Xiaoman Zhang, Ya Zhang, Weidi Xie, and Yanfeng Wang. Pmc-llama:  
744 toward building open-source language models for medicine. *Journal of the American Medical  
745 Informatics Association*, pp. ocae045, 2024.
- 746
- 747 Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, Qingwei  
748 Lin, and Daxin Jiang. Wizardlm: Empowering large pre-trained language models to follow com-  
749 plex instructions. In *The Twelfth International Conference on Learning Representations*, 2024.
- 750
- 751 Huiyun Yang, Huadong Chen, Hao Zhou, and Lei Li. Enhancing cross-lingual transfer by manifold  
752 mixup. In *International Conference on Learning Representations*, 2022.
- 753
- 754 Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical  
755 risk minimization. In *International Conference on Learning Representations*, 2018.
- 756
- 757 Linjun Zhang, Zhun Deng, Kenji Kawaguchi, Amirata Ghorbani, and James Zou. How does mixup  
758 help with robustness and generalization? In *International Conference on Learning Representa-  
759 tions*, 2021.

756 Mike Zhang and Barbara Plank. Cartography active learning. In *Findings of the Association for*  
757 *Computational Linguistics: EMNLP 2021*, pp. 395–406, 2021.  
758

759 Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi  
760 Hu, Tianwei Zhang, Fei Wu, et al. Instruction tuning for large language models: A survey. *arXiv*  
761 *preprint arXiv:2308.10792*, 2023.

762 Shujian Zhang, Chengyue Gong, Xingchao Liu, Pengcheng He, Weizhu Chen, and Mingyuan Zhou.  
763 Allsh: Active learning guided by local sensitivity and hardness. In *Findings of the Association*  
764 *for Computational Linguistics: NAACL 2022*, pp. 1328–1342, 2022.  
765

766 Hao Zhao, Maksym Andriushchenko, Francesco Croce, and Nicolas Flammarion. Long is more for  
767 alignment: A simple but tough-to-beat baseline for instruction fine-tuning. In *Forty-first Interna-*  
768 *tional Conference on Machine Learning*, 2024.

769 Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min,  
770 Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv*  
771 *preprint arXiv:2303.18223*, 2023.

772 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang,  
773 Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and  
774 chatbot arena. *Advances in Neural Information Processing Systems*, 36, 2024.  
775

776 Chunting Zhou, Pengfei Liu, Puxin Xu, Srinu Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat,  
777 Ping Yu, LILI YU, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy.  
778 LIMA: Less is more for alignment. In *Thirty-seventh Conference on Neural Information Process-*  
779 *ing Systems*, 2023.  
780  
781  
782  
783  
784  
785  
786  
787  
788  
789  
790  
791  
792  
793  
794  
795  
796  
797  
798  
799  
800  
801  
802  
803  
804  
805  
806  
807  
808  
809

## A EXPERIMENT DETAILS

In Section 4, we assess the effectiveness of SFTMix against NTP in both instruction-following and domain-specific SFT tasks. Here, we report the detailed experiment results with standard errors.

| Dataset        | Instruction-Tuning LLM | Recipe | MT-Bench             |                      |                      | AlpacaEval-2         |                       |
|----------------|------------------------|--------|----------------------|----------------------|----------------------|----------------------|-----------------------|
|                |                        |        | Single-Turn          | Multi-Turn           | Overall              | Win Rate             | LC Win Rate           |
| Alpaca-52K     | Llama-3.1-8B           | NTP    | 4.9100 ± 0.06        | 3.8150 ± 0.06        | 4.3625 ± 0.05        | 4.0714 ± 0.14        | 8.6528 ± 0.19         |
|                |                        | SFTMix | <b>5.2125</b> ± 0.03 | <b>3.9525</b> ± 0.07 | <b>4.5825</b> ± 0.03 | <b>4.9031</b> ± 0.21 | <b>10.3195</b> ± 0.04 |
|                | Mistral-7B-v0.1        | NTP    | 5.1650 ± 0.08        | 4.0675 ± 0.13        | 4.6163 ± 0.05        | 4.3560 ± 0.15        | 9.1759 ± 0.17         |
|                |                        | SFTMix | <b>5.2775</b> ± 0.03 | <b>4.5425</b> ± 0.10 | <b>4.9100</b> ± 0.05 | <b>4.5386</b> ± 0.18 | <b>9.4994</b> ± 0.28  |
| UltraChat-200K | Llama-3.1-8B           | NTP    | 6.1875 ± 0.04        | 5.0125 ± 0.02        | 5.6000 ± 0.02        | 5.0665 ± 0.12        | 8.4505 ± 0.08         |
|                |                        | SFTMix | <b>6.2750</b> ± 0.07 | <b>5.3500</b> ± 0.04 | <b>5.8125</b> ± 0.02 | <b>5.1149</b> ± 0.16 | <b>9.3810</b> ± 0.25  |
|                | Mistral-7B-v0.1        | NTP    | 5.7625 ± 0.03        | 4.6938 ± 0.01        | 5.2281 ± 0.01        | 4.4899 ± 0.18        | 7.7732 ± 0.21         |
|                |                        | SFTMix | <b>5.9813</b> ± 0.08 | <b>4.8813</b> ± 0.03 | <b>5.4313</b> ± 0.03 | <b>4.6117</b> ± 0.04 | <b>8.7650</b> ± 0.21  |

Table 4: Evaluation of instruction-following capabilities of LLMs trained with NTP or SFTMix. We report the average score and standard errors based on five rounds of evaluation, with the scores from the best-performing instruction-tuning recipe in bold. SFTMix outperforms NTP on both MT-Bench and AlpacaEval-2, irrespective of LLM families and SFT data sizes.

| LLM                         | MedQA               | MedQA-5             | PubMedQA            | MedMCQA             | Macro Ave           |
|-----------------------------|---------------------|---------------------|---------------------|---------------------|---------------------|
| MedAlpaca-7B                | 38.94 ± 0.37        | 33.96 ± 0.26        | 57.20 ± 0.71        | 34.90 ± 0.39        | 41.25 ± 0.40        |
| PMC-LLaMA-7B                | 27.94 ± 0.65        | 21.24 ± 0.56        | 54.87 ± 0.62        | 24.57 ± 0.27        | 32.16 ± 0.40        |
| BioMedGPT-LM-7B             | 38.62 ± 1.51        | 34.72 ± 0.46        | 58.27 ± 0.25        | 35.57 ± 0.52        | 41.80 ± 0.52        |
| Meditron-7B                 | 35.09 ± 0.64        | 26.73 ± 0.19        | 56.93 ± 1.27        | 34.03 ± 0.36        | 38.20 ± 0.39        |
| BioMistral-7B               | 43.86 ± 0.33        | 37.58 ± 0.62        | 50.13 ± 0.66        | 44.14 ± 0.33        | 43.93 ± 0.27        |
| Llama-3.1-8B                | 59.68 ± 0.33        | 53.23 ± 0.21        | 73.40 ± 0.86        | 52.79 ± 0.01        | 59.78 ± 0.51        |
| + NTP on MedAlpaca-263K     | 59.31 ± 0.56        | 54.52 ± 0.21        | 75.40 ± 0.57        | 53.65 ± 0.18        | 60.72 ± 0.08        |
| or SFTMix on MedAlpaca-263K | <b>60.88</b> ± 0.29 | <b>55.38</b> ± 0.13 | <b>77.80</b> ± 0.16 | <b>54.15</b> ± 0.11 | <b>62.05</b> ± 0.24 |
| Mistral-7B-v0.1             | 49.18 ± 0.30        | 43.94 ± 0.23        | 72.33 ± 0.20        | 47.98 ± 0.22        | 53.36 ± 0.43        |
| + NTP on MedAlpaca-263K     | 49.10 ± 0.33        | 44.62 ± 0.39        | 75.40 ± 0.68        | 48.15 ± 0.11        | 54.32 ± 0.62        |
| or SFTMix on MedAlpaca-263K | <b>51.77</b> ± 0.51 | <b>45.72</b> ± 0.44 | <b>77.40</b> ± 0.28 | <b>49.03</b> ± 0.22 | <b>55.98</b> ± 0.31 |

Table 5: Evaluation results on four healthcare-related benchmarks by prior biomedical LLMs and LLMs instruction-tuned on MedAlpaca-263K using either NTP or SFTMix. We report the mean accuracy (%) and standard errors over three rounds of three-shot evaluation and bold the scores from SFTMix-tuned LLMs. SFTMix achieves an approximate 1.5% absolute increase in macro-average performance compared to NTP for both Llama-3.1-8B and Mistral-7B-v0.1.

| Instruction-Tuning LLM | Recipe | MT-Bench    |             |             |             |             |             |             |             | Overall     |
|------------------------|--------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
|                        |        | Writing     | Roleplay    | Reasoning   | Math        | Coding      | Extraction  | STEM        | Humanities  |             |
| Llama-3.1-8B           | NTP    | 6.79        | 4.93        | 2.94        | 1.82        | <b>2.63</b> | 5.96        | 4.46        | 5.37        | 4.36        |
|                        | SFTMix | <b>6.81</b> | <b>5.15</b> | <b>3.12</b> | <b>1.85</b> | <b>2.63</b> | <b>6.80</b> | <b>4.77</b> | <b>5.53</b> | <b>4.58</b> |
| Mistral-7B-v0.1        | NTP    | 6.23        | 5.20        | <b>4.81</b> | 1.21        | 2.63        | 6.82        | 4.44        | 5.59        | 4.62        |
|                        | SFTMix | <b>7.08</b> | <b>5.42</b> | 4.51        | <b>1.27</b> | <b>3.45</b> | <b>6.90</b> | <b>4.97</b> | <b>5.68</b> | <b>4.91</b> |

Table 6: Categorical evaluation of instruction-following capabilities of LLMs trained on Alpaca-52K using NTP or SFTMix. We report the average score in each of the eight categories in MT-Bench based on five rounds of evaluation, with the scores from the best-performing instruction-tuning recipe in bold. By using SFTMix, we observe significant gains in extraction tasks for Llama-3.1-8B, and in writing, coding, and STEM for Mistral-7B-v0.1.