

Instruction Tuning With Loss Over Instructions

Zhengyan Shi¹ Adam X. Yang² Bin Wu¹
Laurence Aitchison² Emine Yilmaz¹ Aldo Lipani¹

¹University College London ²University of Bristol
{zhengxiang.shi.19,bin.wu.23,emine.yilmaz,aldo.lipani}@ucl.ac.uk
{adam.yang,laurence.aitchison}@bristol.ac.uk

Abstract

Instruction tuning plays a crucial role in shaping the outputs of language models (LMs) to desired styles. In this work, we propose a simple yet effective method, INSTRUCTION MODELLING (IM), which trains LMs by applying a loss function to the instruction and prompt part rather than solely to the output part. Through experiments across 21 diverse benchmarks, we show that, in many scenarios, IM can effectively improve the LM performance on both NLP tasks (*e.g.*, MMLU, TruthfulQA, and HumanEval) and open-ended generation benchmarks (*e.g.*, MT-Bench and AlpacaEval). Remarkably, in the most advantageous case, IM boosts model performance on AlpacaEval 1.0 by over 100%. We identify two key factors influencing the effectiveness of IM: (1) The ratio between instruction length and output length in the training data; and (2) The number of training examples. We observe that IM is especially beneficial when trained on datasets with lengthy instructions paired with brief outputs, or under the Superficial Alignment Hypothesis (SAH) where a small amount of training examples are used for instruction tuning. Further analysis substantiates our hypothesis that our improvement can be attributed to reduced overfitting to instruction tuning datasets. It is worth noting that we are not proposing IM as a replacement for the current instruction tuning process. Instead, our work aims to provide practical guidance for instruction tuning LMs, especially in low-resource scenarios. Our code is available at <https://github.com/ZhengxiangShi/InstructionModelling>.

1 Introduction

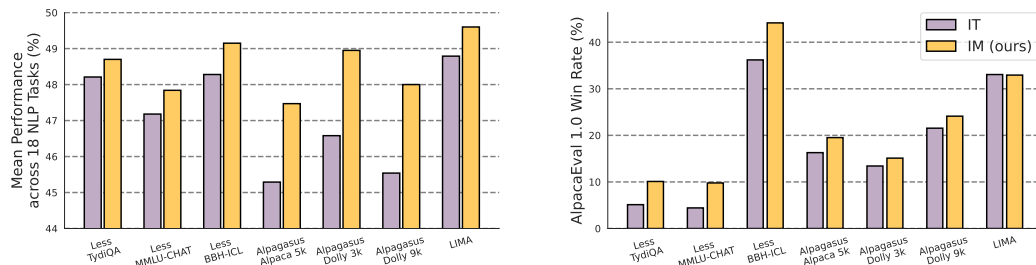


Figure 1: Performance differences between INSTRUCTION TUNING (IT) and our proposed method INSTRUCTION MODELLING (IM) trained on 7 instruction tuning datasets. These datasets contain prompts and responses but do not contain preference pairs. Specifically, we use the Less datasets [68] and Alpapasus datasets [11], which are subsets of Flan V2 [14], Dolly [18], and Stanford Alpaca [61] to ensure good performance. We also report the results on the LIMA dataset. **(Left)** The mean performance across 18 traditional NLP tasks (see §4.1 for details). **(Right)** The win rate on the AlpacaEval 1.0 benchmark [37]. Please refer to §4.2 for details.

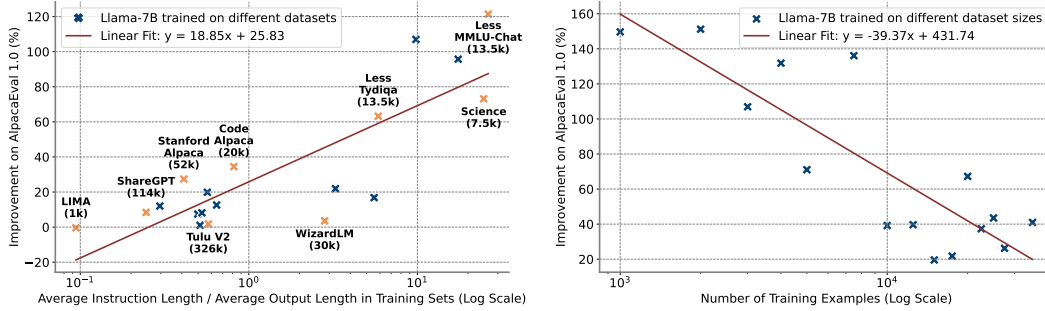


Figure 2: **(Left)** Performance improvement, achieved by our approach INSTRUCTION MODELLING (IM) compared to INSTRUCTION TUNING (IT) on the AlpacaEval 1.0, against the ratio between average instruction length and average output length in instruction tuning datasets (training size noted in parentheses). We highlight several representative instruction tuning datasets in yellow. Our analysis suggests that IM is especially beneficial for datasets characterized by lengthy instructions or prompts paired with comparably brief outputs, such as Code Alpaca [10] and Less MMLU Chat [68]. **(Right)** Performance improvement achieved by our approach IM over IT on the AlpacaEval 1.0 against the number of training examples in instruction tuning datasets. Here we maintain a fixed ratio between instruction and output length of 10. This analysis suggests that IM is particularly effective under the low-resource setting or Superficial Alignment Hypothesis. Please refer to §4.2 for details.

Language models (LMs) are trained to predict the next token on massive corpora, enabling them to learn general-purpose representations transferable to various language understanding or generation tasks [7, 47, 48, 56, 62]. However, it does not align LMs to act in accordance with the user’s intentions [34]. To enable this transfer, various methods for aligning language models [3, 14, 44, 73, 58, 79] have thus been proposed, one of which is instruction tuning (IT) [36, 65, 69]. Recent study [78] proposes Superficial Alignment Hypothesis (SAH): A model’s knowledge and capabilities are learnt almost entirely during pretraining, only minimal instruction tuning data is required to enable high-quality outputs in the desired output style. Existing works [1, 43, 44, 51, 65, 69, 36] mainly perform instruction tuning by focusing the loss computation solely on the output segments.

In this work, we demonstrate that in many scenarios, incorporating the loss computation for instructions or prompts, which we refer to as INSTRUCTION MODELLING (IM) (see §3), could substantially improve the performance of instruction tuning on both various NLP tasks (*e.g.*, MMLU, TruthfulQA, and HumanEval) and open-ended generation benchmarks (*e.g.*, MT-Bench and AlpacaEval), as shown in Figure 1. Remarkably, in the most favourable case, our proposed method IM boosts performance on AlpacaEval 1.0 by over 100%. As illustrated in Figure 2, our study further identifies two key factors influencing the effectiveness of IM: (1) **The ratio between instruction length and output length** (see Figure 2 Left). Our analysis shows that our approach IM is especially beneficial for datasets characterised by lengthy instructions or prompts paired with comparably brief outputs, such as Code Alpaca [13] and Less MMLU Chat [68]; (2) **The number of training examples** (see Figure 2 Right). We demonstrate that our approach IM performs better under the SAH, where a small amount of training examples are available (see §4.2).

Recent works [27, 31, 44, 71, 73] suggest that LMs can quickly memorise training examples even after seeing them just once. We hypothesise that the improvement stems from reducing instruction tuning’s tendency to overfit, particularly under limited training resource conditions: Instruction tuning on brief outputs or a small amount of data can potentially lead to rapid overfitting. To substantiate our hypothesis, our analysis shows that (1) IM exhibits higher training losses but lower test losses on new instruction tuning data; (2) The outputs generated by IM have a lower similarity to the training examples compared to those from IT, as indicated by BLEU scores; and (3) IM leads to less performance degrade on NLP tasks across training epochs (see §4.3). Additionally, our study reveals that this overfitting cannot be effectively addressed by applying Kullback-Leibler (KL) divergence for regularisation [3, 44], as it compromises the model’s ability to follow instructions. Our further analysis reveals that the advantages of IM persist across different LMs and model sizes, and that IM could be effectively combined with the previous approach (*i.e.*, NEFTUNE [31]). Meanwhile, we investigate the relationship between output length and win rate for our approach (see §4.4).

In summary, the main contributions of this paper are:

- We propose INSTRUCTION MODELLING (IM), aiming to enhance both the instruction-following and general performance on NLP tasks of LMs. Through extensive experiments across 21 benchmarks, we demonstrate that, in many scenarios, IM substantially improves the performance of LMs trained on various instruction tuning datasets, particularly notable in the AlpacaEval 1.0 benchmark where it boosts scores by over 100%.
- Our study identifies key factors influencing the effectiveness of IM, including the ratio between instruction length and output length and the number of training examples. We are not proposing IM as a replacement for current instruction tuning processes. Rather, we provide empirical guidance for fine-tuning LMs, especially under low-resource scenarios.
- We provide underlying mechanisms that make IM effective, specifically how it mitigates overfitting, thereby enhancing the LMs’ performance across various tasks.

2 Related Work

Instruction Tuning. LMs can better align with user intents through fine-tuning on datasets consisting of instructions and human-written completions [3, 44]. Early studies mainly focus on NLP tasks, showing that fine-tuning with various NLP datasets trained with instruction output pairs improves cross-task generalisation [1, 33, 43, 44, 51, 54, 65, 66]. Recent works explore the creation of instruction tuning datasets by LMs themselves [63, 25, 69, 36] or through crowdsourcing approaches [13, 78]. Such instruction-tuning phrase [30, 53, 59, 26, 73] enables LLMs to generalise beyond instructions in the training set, largely enhancing their practical utility.

Data Selection for Instruction Tuning. Research on instruction tuning for LMs presents diverging perspectives on the optimal data scale for supervised fine-tuning. A prevailing view recommends fine-tuning on expansive datasets to enhance LM performance across various NLP tasks, thereby improving zero-shot and few-shot learning capabilities [1, 33, 44, 65, 43, 42, 51, 57, 64, 67]. For example, Flan V2 comprises over a million question-answer pairs from diverse NLP sources [14], and Natural Instructions features 61 distinct tasks and 193k task instances [43]. Conversely, another research trajectory prioritises data quality over quantity [20, 70, 40, 32]. The Superficial Alignment Hypothesis (SAH) [78] advocates for using smaller, high-quality datasets, arguing that LMs primarily acquire their capabilities during the pretraining phase and thus require only minimal data for effective instruction tuning. For instance, LIMA [78] employs a carefully curated set of 1k diverse prompts to generate stylistically consistent responses, aimed at creating a helpful AI assistant. AlpaGasus [11] and LESS [68] employ methods to select high-quality data based on LLM-generated judgements and gradient signals. However, both views agree on the importance of (1) the quality of pre-trained base LMs and (2) the diversity and quality of the IT data.

Regularisation Through Language Modelling Objectives. Pretraining data and language modelling objectives have been used as a regularisation technique in fine-tuning LMs. In particular, [15, 39] fine-tunes LMs on labelled data, with unsupervised learning on unlabelled data for auxiliary tasks as regularisation. [44] mixes the alignment objective with the next token prediction objective using pretraining data to mitigate alignment tax in reinforcement learning from human feedback (RLHF). [22] adopts the masked language objective on the pretraining or downstream task corpus to preserve pre-trained features, and shows improvements in calibration and accuracy. [29] investigates the effect of incorporating instruction loss weighting on instruction tuning, suggesting that the instruction loss ratio is an important hyperparameter when fine-tuning short-completion data but is irrelevant when using long-completion data. In this work, we propose a broader guideline that does not introduce new hyperparameters but focuses on when and how to include loss over instruction effectively. We refer to our approach as INSTRUCTION MODELLING because it combines elements of both language modelling and instruction tuning.

3 Our Approach

In this section, we first revisit the background of INSTRUCTION TUNING (IT) and then introduce our proposed method, INSTRUCTION MODELLING (IM).

Instruction Tuning. In instruction tuning, each input is a concatenation of an instruction I and a completion C . Let I be the instruction sequence $\{I_1, I_2, \dots, I_m\}$ and C be the completion (output) sequence $\{C_1, C_2, \dots, C_n\}$, where I may include special prompt template tokens (such as “<|user|>” and “<|assistant|>”). The total input sequence x is $\{I_1, I_2, \dots, I_m, C_1, C_2, \dots, C_n\}$. The model predicts each token in C given all the previous tokens in I and C up to that point:

$$P(C_1, C_2, \dots, C_n | I_1, I_2, \dots, I_m) = \prod_{j=1}^n P(C_j | I_1, I_2, \dots, I_m, C_1, C_2, \dots, C_{j-1}) \quad (1)$$

The loss function, \mathcal{L} , for instruction tuning is the negative log-likelihood of the completions given the instructions, expressed as follows:

$$\mathcal{L} = -\log P(C_1, C_2, \dots, C_n | I_1, I_2, \dots, I_m) = -\sum_{j=1}^n \log P(C_j | I_1, I_2, \dots, I_m, C_1, C_2, \dots, C_{j-1}) \quad (2)$$

This approach aims to optimise the predictions for the completion sequence C , using the instruction sequence I as contextual information.

Our Approach: Instruction Modelling. Our approach, instruction modelling, is an expansion of instruction tuning by incorporating loss calculation for both the instruction and the completion tokens, except it omits any special prompt template tokens. The model is trained to predict both the instruction and completion parts of x but excludes tokens that are part of prompt templates (denoted as T). For simplicity, we consider that these template tokens are not part of I or C . The model predicts the next token given all previous tokens (both instructions and completions up to that point):

$$P(x) = P(I_1, I_2, \dots, I_m, C_1, C_2, \dots, C_n) = \prod_{t=1}^{m+n} P(x_t | x_1, x_2, \dots, x_{t-1}) \quad (3)$$

The loss function, \mathcal{L} , for instruction modelling calculates the negative log-likelihood for both instruction and completion tokens, excluding any prompt template tokens. It is computed as follows:

$$\mathcal{L} = -\sum_{t=1}^{m+n} \log P(x_t | x_1, x_2, \dots, x_{t-1}) \cdot \mathbf{1}(x_t \notin T), \quad (4)$$

where $\mathbf{1}(x_t \notin T)$ is an indicator function that is 1 if x_t is not a template token and 0 otherwise. This ensures that the loss is computed only over the meaningful tokens, not over the static template tokens. Our approach allows the model to improve its understanding of both the instructions and the completions while being sensitive to the context provided by both segments of the input sequence.

4 Experiments and Results

In this section, we evaluate the effectiveness of our proposed method INSTRUCTION MODELLING (IM) by comparing it with INSTRUCTION TUNING (IT) and other baselines on various datasets.

4.1 Experimental Setup

Instruction Tuning Datasets. We assess our method, IM, across various instruction tuning datasets, detailed as follows: (1) Stanford Alpaca [61] (52 002 examples); (2) Dolly [18] (15 011 examples); (3) Sharegpt [13] (50 000 examples); (4) Code Alpaca [10] (20 022 examples); (5) Science Literature [30] (7 544 examples); (6) WizardLM [69] (30 000 examples); (7) Tulu V2 [30] (326 181 examples). Additionally, we incorporate instruction tuning datasets under the low-resource setting or SAH: (8) LIMA [78] (1 030 examples); (9) Less¹ [68], where high-quality instruction tuning data are selected from Flan V2 and Dolly. Here, we use the Less MMLU Chat (13 533 examples), Less BBH ICL (13 533 examples), and Less Tydiqa (13 533 examples); (10) Alpagoasus² [11], which offers three subsets: Alpagoasus Dolly 3k (2 996 examples), Alpagoasus Dolly 9k (9 229 examples) selected from Dolly, and Alpagoasus Alpaca 5k (5 305 examples) selected from Stanford Alpaca. See dataset details and statistical analysis in Appendix §A.

¹<https://github.com/princeton-nlp/LESS>

²<https://github.com/gpt4life/alpagoasus>

Table 1: Performance comparisons using 7 instruction tuning datasets with the LLAMA-2-7B on 6 categories of 18 traditional NLP tasks and 3 open-ended benchmarks with LLM as judgements. “IT” refers to INSTRUCTION TUNING. “IM” refers to INSTRUCTION MODELLING. Green and red arrows indicate performance changes against the baseline (IT).

Method	NLP Benchmarks						LLM-based Evaluation			
	Understanding & Knowledge	Multi-linguality	Commonsense Reasoning	Math&Code Reasoning	BBH	Safety & Helpfulness	Mean	MT-Bench	AlpacaEval 1.0	AlpacaEval 2.0
LLAMA-2-BASE	63.91	61.99	75.86	13.32	38.80	42.03	49.32	1.16	0.01	0.01
LLAMA-2-CHAT	63.42	55.15	70.28	15.33	38.92	51.79	49.15	6.63	79.04	6.48
Alpapasus Alpaca 5k (5,305 training examples)										
IT	64.98	57.24	66.06	8.93	26.80	47.74	45.29	3.62	16.29	2.46
NEFTUNE	65.18	56.88	66.45	10.24	29.53	45.46	45.62 ^{↑0.33}	3.50 ^{↓0.12}	21.37 ^{↑5.08}	2.37 ^{↓0.09}
IM (ours)	64.01	56.63	72.47	11.58	35.52	44.62	47.47 ^{↑2.18}	3.48 ^{↓0.14}	19.52 ^{↑3.23}	3.29 ^{↑0.83}
Alpapasus Dolly 3k (2,996 training examples)										
IT	65.81	57.46	67.55	11.96	33.02	43.70	46.58	4.23	13.42	2.00
NEFTUNE	65.90	57.79	67.28	11.64	35.43	44.36	47.07 ^{↑0.49}	4.42 ^{↑0.19}	14.04 ^{↑0.62}	2.03 ^{↑0.03}
IM (ours)	65.66	57.47	73.24	14.57	37.48	45.29	48.95 ^{↑2.37}	4.06 ^{↓0.17}	15.11 ^{↑1.69}	2.44 ^{↑0.44}
Alpapasus Dolly 9k (9,229 training examples)										
IT	64.10	56.62	69.70	7.96	32.19	42.65	45.54	4.33	21.54	2.28
NEFTUNE	64.20	56.69	69.51	8.99	33.91	42.62	45.99 ^{↑0.45}	4.21 ^{↓0.12}	31.61 ^{↑10.07}	2.84 ^{↑0.56}
IM (ours)	64.67	55.32	74.87	12.50	36.69	43.96	48.00 ^{↑2.46}	4.55 ^{↑0.22}	30.77 ^{↑9.23}	2.67 ^{↑0.39}
Less Tydiqa (13,533 training examples)										
IT	64.01	56.81	64.77	12.06	36.54	55.09	48.21	4.08	5.12	1.88
NEFTUNE	64.03	55.09	64.02	13.84	36.65	51.21	47.47 ^{↓0.74}	4.19 ^{↑0.11}	8.35 ^{↑3.23}	2.58 ^{↑0.70}
IM (ours)	64.28	56.10	65.70	17.15	34.86	54.09	48.70 ^{↑0.49}	4.36 ^{↑0.28}	10.10 ^{↑4.98}	2.88 ^{↑1.00}
Less MMLU Chat (13,533 training examples)										
IT	64.74	57.42	62.94	9.53	33.13	55.35	47.18	3.86	4.42	1.20
NEFTUNE	65.21	57.43	63.14	9.45	35.89	55.32	47.74 ^{↑0.56}	4.06 ^{↑0.20}	6.22 ^{↑1.80}	1.06 ^{↓0.14}
IM (ours)	63.95	56.34	64.76	12.52	36.94	52.55	47.84 ^{↑0.66}	4.54 ^{↑0.68}	9.78 ^{↑5.36}	1.93 ^{↑0.73}
Less BBH ICL (13,533 training examples)										
IT	63.83	62.04	75.92	6.90	38.93	42.07	48.28	4.78	36.20	2.36
NEFTUNE	63.88	58.83	67.97	13.54	38.63	51.33	49.03 ^{↑0.75}	5.05 ^{↑0.27}	39.81 ^{↑3.61}	2.87 ^{↑0.51}
IM (ours)	64.14	56.72	71.12	13.56	39.03	50.34	49.15 ^{↑0.87}	5.03 ^{↑0.25}	44.15 ^{↑7.95}	3.56 ^{↑1.20}
LIMA (1,030 training examples)										
IT	63.92	58.29	71.96	16.01	39.27	43.29	48.79	4.77	33.06	2.58
10 epoch NEFTUNE	63.66	57.67	73.03	15.95	38.77	43.14	48.70 ^{↓0.09}	4.79 ^{↑0.02}	30.51 ^{↓2.55}	2.43 ^{↓0.15}
IM (ours)	64.49	58.21	75.55	17.06	38.84	43.45	49.60 ^{↑0.81}	4.83 ^{↑0.06}	32.94 ^{↓0.12}	2.47 ^{↓0.11}

Evaluation Benchmarks. Our study conducts a comprehensive analysis of 21 NLP datasets, focusing on a suite of canonical NLP benchmarks and their capacity for open-ended language generation. For canonical NLP benchmarks, the evaluation is organised into six categories (18 tasks in total): (1) *Language Understanding and Knowledge* includes MMLU [24], PIQA [6], OpenbookQA [41], HellaSwag [74], and LAMBADA [45]; (2) *Multilinguality* contains LAMBADA Multilingual [45], WMT 2014 [8], and WMT 2016 [52]; (3) *Commonsense Reasoning* features Winograd schema challenge (WSC) [35], WinoGrande [50], AI2 Reasoning Challenge (ARC) [16], and CoQA [49]; (4) *Math and Coding Reasoning* includes GSM8K [17], and HumanEval [12]; (5) *Safety and Helpfulness* comprises TruthfulQA [38], ToxiGen [21], and Hendrycks Ethics [23]. (6) *Big Bench Hard (BBH)* dataset [60] is included to assess models. Our models are also tested for their open-ended text generation capabilities using model-based evaluations, specifically through MT-Bench [77], AlpacaEval 1.0 and 2.0 [37], where the AlpacaEval 1.0 compares the model outputs against the `text_davinci_003` evaluated by GPT-4 and the AlpacaEval 2.0 compares the model outputs against GPT-4 outputs evaluated by GPT-4 Turbo. See evaluation details in Appendix §B.

All Comparison Approaches. In our study, we mainly conduct experiment using the LLAMA-2-7B-BASE and LLAMA-2-13B-BASE [62], and the OPT-6.7B [75] models. We report model performance trained on LLAMA-2-7B-BASE if not specified. We compare with NEFTUNE [31] as the baseline, which adds noise to the embedding during the instruction tuning to increase the robustness of instruction-tuned models. See hyperparameter and implementation details in Appendix §C.

4.2 Main Results

In this section, we first evaluate the model performance of our approach and baselines across various tasks. Then we investigate the key factors that contribute to the effectiveness of our approach. Below we will discuss our findings in detail.

#1: Our approach IM can improve the performance of instruction tuning on various NLP tasks and open-ended generation benchmarks. Figure 1 provides a summary of the model’s performance across both traditional NLP tasks and the AlpacaEval 1.0 benchmark. Table 1 offers

a detailed breakdown of experimental results for traditional NLP tasks across six categories, as well as performance on additional benchmarks for open-ended generation (*i.e.*, MT-Bench and AlpacaEval). The experimental results show that our approach (IM) can improve the performance of instruction tuning on various NLP tasks and open-ended generation benchmarks. Specifically, on the `Alpapasus Dolly 3k` dataset, IM improves the overall mean score of NLP tasks to 48.95, an increase of 2.37 points from the baseline. Similarly, on the `Alpapasus Dolly 9k` dataset, we observe an improvement of 2.46 points in the mean NLP score. These improvements are mirrored in the LLM-based evaluations. Specifically, IM raises scores on the AlpacaEval 1.0 benchmark, achieving approximately a ten-point increase on the `Alpapasus Dolly 9k` dataset and doubling the performance on datasets such as `Less MMLU Chat` and `Less Tydiqa`. However, the extent of improvement varies across different datasets. For example, the `LIMA` dataset shows more modest gains, prompting our further analysis to understand the factors influencing the effectiveness of IM.

#2: Our approach IM is especially beneficial for datasets characterised by lengthy instructions or prompts paired with comparably brief outputs. To better understand the impact factors on the effectiveness of IM, we extend our experiments to more instruction-tuning datasets, such as `Science Literature`, `Code Alpaca` and `Tulu V2`. Interestingly, as shown in Figure 2 Left, we find that IM is particularly effective in scenarios where datasets characterised by lengthy instructions and shorter outputs, such as `Less MMLU Chat` and `Less BBH ICL`. For example, in datasets like `Less MMLU Chat` and `Less Tydiqa`, IM shows remarkable efficacy. In contrast, the `Tulu V2` dataset, with an instruction to output length ratio of about 0.5, benefits less compared to the `Science Literature` dataset, which has a much higher ratio of 24.7. We hypothesise that this trend can be attributed to the tendency of language models trained on datasets with shorter outputs to overfit. In cases where the instructions are longer, IM acts as an effective form of regularisation, mitigating this issue. For further details on the experimental setup, refer to the Appendix in §C.

#3: Our approach IM performs better with fewer training examples. We find that another important factor in the effectiveness of IM is the quantity of training examples. Specifically, we design additional experiments by sampling different numbers of examples from the `Tulu V2` datasets, which contain about 320k training examples and achieve a modest improvement compared to other datasets in Figure 2 Left. We ensure that our samples maintain an instruction-to-output length ratio of around 10. As all these samples are from `Tulu V2`, we can assume they are from the same distribution. As shown in Figure 2 Right, IM demonstrates substantial performance improvements on the AlpacaEval 1.0 as the number of training examples decreases. This suggests that IM could be particularly valuable for developing robust models in resource-constrained scenarios or under the SAH. For further details on the experimental setup, please refer to the Appendix in §C.

4.3 Instruction Modelling Mitigates Overfitting of Instruction Tuning

This section explores the underlying interpretation behind the effectiveness of our approach. Our experimental results demonstrate that IM can alleviate the overfitting problem of Instruction Tuning. Below we will discuss our findings in detail.

#1. Train and test loss analysis. Figure 3 clearly illustrates the effectiveness of our approach IM in mitigating overfitting issues compared to IT. For both IM and IT, here we only compute the loss over the output part. In the training loss distribution for the `LIMA` dataset, IM exhibits a slightly higher mean loss of 1.45 compared to 1.37 for IT, suggesting that IM does not overfit to the training data as much as IT does. This is further corroborated in the test loss distribution on the `Tulu V2` dataset (using a 10% randomly sampled data set), where IM demonstrates a lower mean test loss of 1.17 compared to 1.32 for IT. This indicates that IM maintains better generalisation to new data, emphasising the model’s capability to learn effectively without fitting excessively to training examples. For more examples, see Appendix §D.

#2. BLEU score analysis. Here we generate outputs using the instructions from the training examples via greedy decoding, and then compare the generated outputs with the ground truth outputs in training examples and report the results. We use the BLEU score (up to n-gram order 4) [46] to measure the similarity between outputs, where a higher score on outputs indicates a higher overlap with training examples. As shown in Table 2, outputs generated by IM consistently have lower BLEU

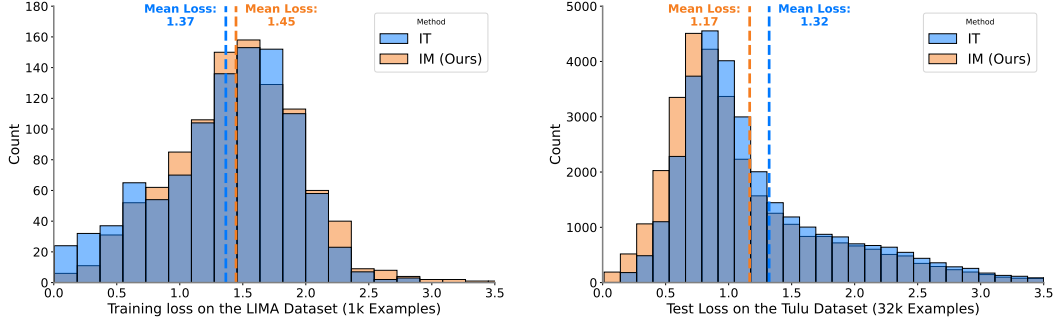


Figure 3: **(Left)** Training loss distribution for each example between our approach INSTRUCTION MODELLING (IM) and INSTRUCTION TUNING (IT) on the LIMA dataset. **(Right)** Test loss distribution for each example between IM and IT on the TuLu v2 dataset, using a 10% randomly sampled data for efficacy. Mean losses are marked by dashed lines. For both IM and IT, here we only compute the loss over the output part. IM has a higher train loss with lower test loss, suggesting that IM effectively mitigates the overfitting issues compared to IT. See Appendix §D for more examples.

Table 2: Average BLEU Score comparison of IM and IT, where a lower score indicates less overfitting. Green and red arrows indicate performance changes against the baseline (IT).

	LIMA	Less Tydiqa	Less MMLU Chat	Less BBH ICL	Alpapasus Alpaca 5k	Alpapasus Dolly 9k	Alpapasus Dolly 3k
IT	18.15	69.21	72.43	60.96	72.26	61.76	60.99
IM (ours)	17.30↓0.85	65.63↓3.58	69.20↓3.23	53.94↓7.02	70.50↓1.76	60.61↓1.15	59.04↓1.95

scores than those generated by IT. This suggests that IM produces outputs have less overlap with the ground truth outputs in training examples, indicating less overfitting.

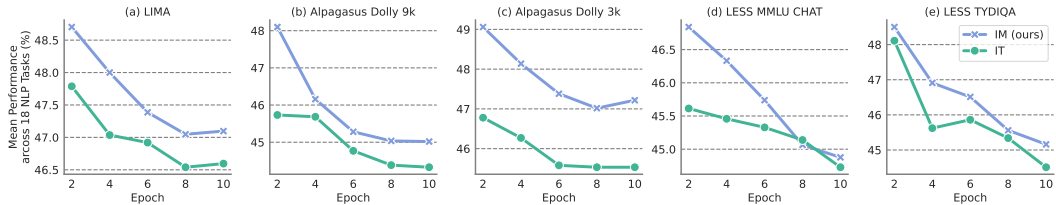


Figure 4: Mean performance on 18 NLP tasks over epochs using LLAMA-2-7B-BASE. This analysis suggests that IM experiences a lower instruction tuning tax compared to IT.

#3. Instruction Tuning Tax on the NLP tasks. Previous works show that training LMs with RLHF causes an *Alignment Tax* on the NLP tasks [3, 44]. In this study, we observe that instruction tuning can sometimes lead to diminished model capabilities in some areas, such as multilinguality and commonsense reasoning. To this end, we further explore the impact of instruction tuning on the performance of NLP tasks. Figure 4 illustrates that our approach IM generally has a lower instruction tuning tax compared to IT, suggesting better robustness under low-resource settings. We provide additional experiments for win rates across epochs in Appendix §E.

#4. Can we simply use KL divergence loss as regularisation for instruction tuning? In this analysis, we demonstrate that the application of KL divergence loss in instruction tuning, which is widely used as regularisation for aligning LMs [3, 44, 72], cannot easily address the overfitting issue of instruction tuning. Table 3 offers a detailed comparison across various NLP benchmarks and open-ended language generation tasks, particularly using AlpacaEval 2.0, with models trained with and without KL divergence loss. Our findings are as follows: (1) Incorporating KL Loss reduces overfitting and reduces the performance degradation on traditional NLP tasks. For example, on the

Table 3: Performance on 18 NLP benchmarks and AlpacaEval 2.0. Green and red arrows indicate performance changes against the baseline (LLAMA-2-7B-BASE). This analysis suggests that while applying KL Loss in the instruction tuning helps mitigate performance degradation in NLP tasks, it substantially harms the model performance in open-ended generation tasks.

	LLAMA-2-7B-BASE	LIMA (1K)		ALPAGASUS DOLLY (9K)	
		IT w/o KL Loss	IT w/ KL Loss	IT w/o KL Loss	IT w/ KL Loss
NLP Tasks	49.32	48.79↓0.53	49.26↓0.06	45.54↓3.78	49.31↓0.01
AlpacaEval 2.0	0.01	2.58↑2.57	0.06↑0.05	2.28↑2.27	0.04↑0.03

Dolly dataset, incorporating KL Divergence Loss leads to less instruction tuning tax in NLP tasks, with scores rising from 45.54 to 49.31. (2) KL Loss detrimentally affects the model’s instructions following abilities. For example, on the LIMA dataset, we observe a substantial decrease in AlpacaEval 2.0 from 2.58 to 0.06. For additional experiments and implementation details, see Appendix §F.

4.4 Further Analysis

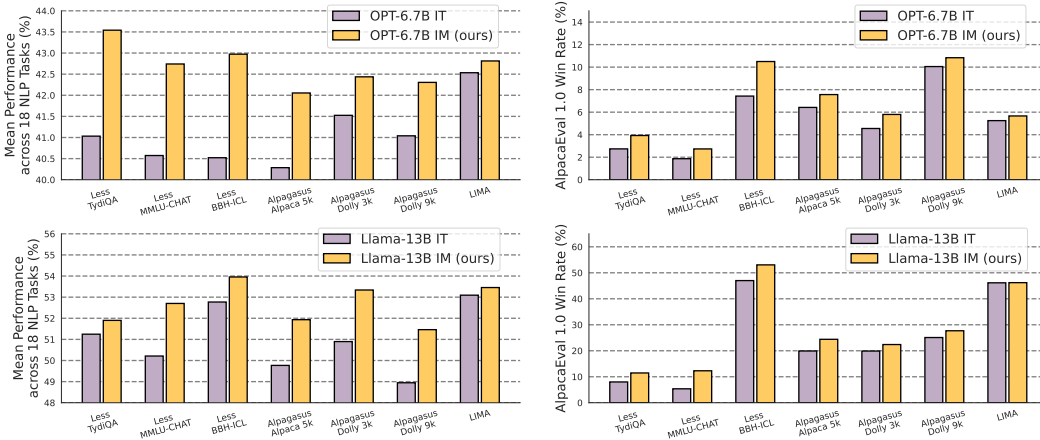


Figure 5: Comparison of INSTRUCTION TUNING (IT) and INSTRUCTION MODELLING (IM) methods using OPT-6.7B (Top Row) and LLAMA-2-13B-BASE (Bottom Row) trained on 7 instruction tuning datasets. (Left) The mean performance across 18 NLP tasks. (Right) The win rate on the AlpacaEval 1.0 benchmark.

#1. The advantage of our proposed method persists with different language models and sizes. As shown in Figure 5, our analysis demonstrates that our proposed method IM consistently outperforms the IT across different models and sizes, including OPT-6.7B and LLAMA-2-13B-BASE, on 18 traditional NLP tasks and the AlpacaEval 1.0 benchmark. These findings underline the effectiveness of our approach irrespective of the underlying language model or its scale.

#2. Relationship between the model output length and the win rate. In this analysis, we explore the potential connection between win rates on the AlpacaEval and the increased output length [37, 76, 31]. As shown in Figure 6, our result reveals that our approach IM does not necessarily generate longer outputs than IT across different data utilisation levels from the TuLu V2 dataset. Specifically, the output lengths for both approaches are similar despite varying levels of data utilisation. Furthermore, IM consistently outperforms the IT, suggesting that improvements in performance as measured by win rates on the AlpacaEval 1.0 are not dependent on the output length. We provide additional analysis on other instruction tuning datasets under the SAH in Appendix §G.

#3. Our proposed method IM could further improve the model performance with NEFTUNE. Table 4 demonstrates the combined effects of our proposed method IM and NEFTUNE on performance across various NLP tasks and the AlpacaEval 1.0 benchmark. The integration of NEFTUNE with IM

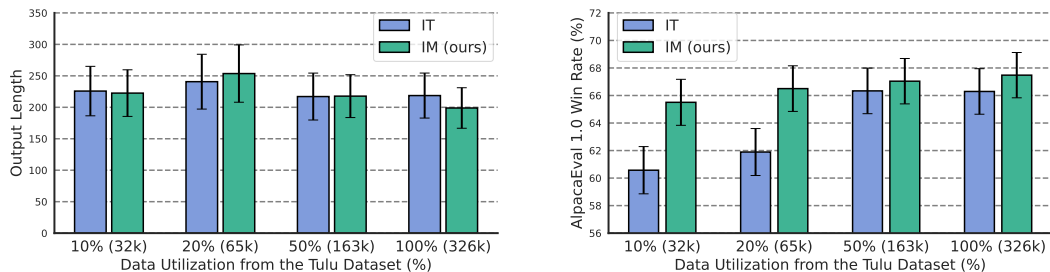


Figure 6: **(Left)** Output length comparison between our approach INSTRUCTION MODELLING (IM) and INSTRUCTION TUNING (IT) across various data utilisation levels from the Tulu V2 dataset, as evaluated on the AlpacaEval dataset. **(Right)** Performance comparison (measured by win rate) between IM and IT on the AlpacaEval 1.0 across various data utilisation levels from the Tulu V2 dataset. This analysis suggests that the improvement provided by IM is not necessarily associated with the increased output lengths. See more length analysis in Appendix §G.

Table 4: Performance comparison of IM and IM +NEFTUNE on AlpacaEval 1.0 and various NLP benchmarks. Green and red arrows indicate performance changes against the baseline (IM). This analysis shows that adding NEFTUNE to IM could further improve model performance.

	LIMA	Less Tydiqa	Less MMLU Chat	Less BBH ICL	Alpapasus Alpaca 5k	Alpapasus Dolly 9k	Alpapasus Dolly 3k
AlpacaEval 1.0 Win Rate							
IM	32.94	10.10	9.78	44.15	19.52	30.77	15.11
IM +NEFTUNE	30.77↓2.17	23.41↑13.31	12.45↑2.67	48.25↑4.10	32.07↑12.55	38.28↑7.51	23.35↑8.24
Mean Performance Across 18 NLP Tasks							
IM	49.60	48.70	47.84	49.15	47.47	48.00	48.95
IM +NEFTUNE	49.47↓0.13	49.44↑0.74	47.73↓0.11	48.62↓0.53	48.70↑1.23	48.63↑0.63	49.54↑0.59

generally further improves the win rates in AlpacaEval 1.0, showing notable improvements in several datasets such as a 13.31% increase on Less Tydiqa and a 12.55% boost on Alpapasus Alpaca 5k (in absolute). However, this combination leads to a performance drop in certain contexts, such as a lower performance on NLP tasks on Less MMLU Chat and Less BBH ICL. This indicates that while NEFTUNE may enhance model robustness under certain conditions, its benefits are context-dependent, highlighting the need for the careful application of NEFTUNE when used in conjunction with IM to optimise effectiveness across diverse evaluation settings.

5 Epilogue

Conclusion. Our study proposes INSTRUCTION MODELLING, which trains LMs with loss over instructions rather than outputs only. Our experimental evaluations demonstrate that our approach largely improves the performance of LMs on both NLP tasks and open-ended generation benchmarks in some scenarios, especially under the Superficial Alignment Hypothesis and low-resource setting where minimal training data is used for instruction tuning. Our analysis has shed light on two key factors that influence the effectiveness of our approach, (1) the ratio between instruction and output lengths, and (2) the quantity of training data, providing practical insights for optimising instruction-based training methods. Additionally, our analysis reveals the mechanisms behind the effectiveness of IM, particularly its ability to reduce overfitting, showing that applying instruction losses in some scenarios can lead to more robust and adaptable LMs.

Limitations and Broader Impact. Here we discuss some potential limitations and the broader impact of our work. Several limitations are outlined as follows: (1) The success of our approach relies on the quality and diversity of the instructions and prompts in the training datasets. Poorly defined or ambiguous instructions may undermine the effectiveness of IM, leading to sub-optimal performance; and (2) It is crucial to ensure that the instructions are ethically sound and free from

harmful or biased content. Training on inappropriate or toxic instructions may result in undesirable outputs. Previous works [5, 7, 4] have extensively discussed the risks and potential harms associated with LMs, including the amplification of undesirable biases learned from training data [5, 2, 9]. Our work has the potential to positively impact the community by helping to mitigate overfitting, resulting in models that are more robust and generalise better to new data, especially in low-resource scenarios. This can enhance the reliability and trustworthiness of AI systems in real-world applications.

Acknowledgments and Disclosure of Funding

The authors express their gratitude to the NeurIPS reviewers and area chairs for their insightful discussions. Zhengyan is funded by the Research Studentship from University College London (UCL). Bin is funded by the Bloomberg Fellowship. The authors gratefully appreciate the generous support from OpenAI for providing API credits through the OpenAI API Researcher Access Program.

References

- [1] Vamsi Aribandi, Yi Tay, Tal Schuster, Jinfeng Rao, Huaixiu Steven Zheng, Sanket Vaibhav Mehta, Honglei Zhuang, Vinh Q. Tran, Dara Bahri, Jianmo Ni, Jai Prakash Gupta, Kai Hui, Sebastian Ruder, and Donald Metzler. Ext5: Towards extreme multi-task scaling for transfer learning. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [2] Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language models. *ArXiv preprint*, abs/2108.07732, 2021.
- [3] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *ArXiv preprint*, abs/2204.05862, 2022.
- [4] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21*, page 610–623, New York, NY, USA, 2021. Association for Computing Machinery.
- [5] Emily M. Bender and Alexander Koller. Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online, 2020. Association for Computational Linguistics.
- [6] Yonatan Bisk, Rowan Zellers, Ronan LeBras, Jianfeng Gao, and Yejin Choi. PIQA: reasoning about physical commonsense in natural language. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 7432–7439. AAAI Press, 2020.
- [7] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [8] Chris Callison-Burch, Philipp Koehn, Christof Monz, and Josh Schroeder. Findings of the 2009 Workshop on Statistical Machine Translation. In *Proceedings of the Fourth Workshop on Statistical Machine Translation*, pages 1–28, Athens, Greece, 2009. Association for Computational Linguistics.

- [9] Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom B Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data from large language models. In *USENIX Security Symposium*, volume 6, 2021.
- [10] Sahil Chaudhary. Code alpaca: An instruction-following llama model for code generation. <https://github.com/sahil280114/codealpaca>, 2023.
- [11] Lichang Chen, Shiyang Li, Jun Yan, Hai Wang, Kalpa Gunaratna, Vikas Yadav, Zheng Tang, Vijay Srinivasan, Tianyi Zhou, Heng Huang, and Hongxia Jin. Alpapasus: Training a better alpaca model with fewer data. In *The Twelfth International Conference on Learning Representations*, 2024.
- [12] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code. *ArXiv preprint*, abs/2107.03374, 2021.
- [13] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, 2023.
- [14] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53, 2024.
- [15] Kevin Clark, Minh-Thang Luong, Christopher D. Manning, and Quoc Le. Semi-supervised sequence modeling with cross-view training. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1914–1925, Brussels, Belgium, 2018. Association for Computational Linguistics.
- [16] Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge. *ArXiv preprint*, abs/1803.05457, 2018.
- [17] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems, 2021.
- [18] Mike Conover, Matt Hayes, Ankit Mathur, Jianwei Xie, Jun Wan, Sam Shah, Ali Ghodsi, Patrick Wendell, Matei Zaharia, and Reynold Xin. Free dolly: Introducing the world’s first truly open instruction-tuned llm, 2023.
- [19] Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning, 2023.
- [20] Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, et al. Textbooks are all you need. *ArXiv preprint*, abs/2306.11644, 2023.
- [21] Thomas Hartvigsen, Saadia Gabriel, Hamid Palangi, Maarten Sap, Dipankar Ray, and Ece Kamar. ToxiGen: A large-scale machine-generated dataset for adversarial and implicit hate speech detection. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3309–3326, Dublin, Ireland, 2022. Association for Computational Linguistics.

- [22] Guande He, Jianfei Chen, and Jun Zhu. Preserving pre-trained features helps calibrate fine-tuned language models. *ArXiv preprint*, abs/2305.19249, 2023.
- [23] Dan Hendrycks, Collin Burns, Steven Basart, Andrew Critch, Jerry Li, Dawn Song, and Jacob Steinhardt. Aligning AI with shared human values. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [24] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [25] Or Honovich, Thomas Scialom, Omer Levy, and Timo Schick. Unnatural instructions: Tuning language models with (almost) no human labor. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14409–14428, Toronto, Canada, 2023. Association for Computational Linguistics.
- [26] Tom Hosking, Phil Blunsom, and Max Bartolo. Human feedback is not gold standard. In *The Twelfth International Conference on Learning Representations*, 2024.
- [27] Jeremy Howard and Jonathan Whitaker. Can llms learn from a single example?, 2023.
- [28] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [29] Mathew Huerta-Enochian. Instruction fine-tuning: Does prompt loss matter?, 2024.
- [30] Hamish Ivison, Yizhong Wang, Valentina Pyatkin, Nathan Lambert, Matthew Peters, Pradeep Dasigi, Joel Jang, David Wadden, Noah A Smith, Iz Beltagy, et al. Camels in a changing climate: Enhancing lm adaptation with tulu 2, 2023.
- [31] Neel Jain, Ping yeh Chiang, Yuxin Wen, John Kirchenbauer, Hong-Min Chu, Gowthami Somepalli, Brian R. Bartoldson, Bhavya Kailkhura, Avi Schwarzschild, Aniruddha Saha, Micah Goldblum, Jonas Geiping, and Tom Goldstein. NEFTune: Noisy embeddings improve instruction finetuning. In *The Twelfth International Conference on Learning Representations*, 2024.
- [32] Aditi Jha, Sam Havens, Jeremy Dohmann, Alex Trott, and Jacob Portes. Limit: Less is more for instruction tuning across evaluation paradigms. *ArXiv preprint*, abs/2311.13133, 2023.
- [33] Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hannaneh Hajishirzi. UNIFIEDQA: Crossing format boundaries with a single QA system. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1896–1907, Online, 2020. Association for Computational Linguistics.
- [34] Jan Leike, David Krueger, Tom Everitt, Miljan Martic, Vishal Maini, and Shane Legg. Scalable agent alignment via reward modeling: a research direction. *ArXiv preprint*, abs/1811.07871, 2018.
- [35] Hector J. Levesque, Ernest Davis, and Leora Morgenstern. The winograd schema challenge. In *13th International Conference on the Principles of Knowledge Representation and Reasoning, KR 2012*, Proceedings of the International Conference on Knowledge Representation and Reasoning, pages 552–561. Institute of Electrical and Electronics Engineers Inc., 2012. 13th International Conference on the Principles of Knowledge Representation and Reasoning, KR 2012 ; Conference date: 10-06-2012 Through 14-06-2012.
- [36] Xian Li, Ping Yu, Chunting Zhou, Timo Schick, Omer Levy, Luke Zettlemoyer, Jason E Weston, and Mike Lewis. Self-alignment with instruction backtranslation. In *The Twelfth International Conference on Learning Representations*, 2024.

- [37] Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. AlpacaEval: An automatic evaluator of instruction-following models, 2023.
- [38] Stephanie Lin, Jacob Hilton, and Owain Evans. TruthfulQA: Measuring how models mimic human falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252, Dublin, Ireland, 2022. Association for Computational Linguistics.
- [39] Fangyu Liu, Qianchu Liu, Shruthi Bannur, Fernando Pérez-García, Naoto Usuyama, Sheng Zhang, Tristan Naumann, Aditya Nori, Hoifung Poon, Javier Alvarez-Valle, Ozan Oktay, and Stephanie L. Hyland. Compositional zero-shot domain transfer with text-to-text models. *Transactions of the Association for Computational Linguistics*, 11:1097–1113, 2023.
- [40] Wei Liu, Weihao Zeng, Keqing He, Yong Jiang, and Junxian He. What makes good data for alignment? a comprehensive study of automatic data selection in instruction tuning. *ArXiv preprint*, abs/2312.15685, 2023.
- [41] Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2381–2391, Brussels, Belgium, 2018. Association for Computational Linguistics.
- [42] Sewon Min, Mike Lewis, Luke Zettlemoyer, and Hannaneh Hajishirzi. MetaCL: Learning to learn in context. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2791–2809, Seattle, United States, 2022. Association for Computational Linguistics.
- [43] Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. Cross-task generalization via natural language crowdsourcing instructions. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3470–3487, Dublin, Ireland, 2022. Association for Computational Linguistics.
- [44] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- [45] Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Ngoc Quan Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. The LAMBADA dataset: Word prediction requiring a broad discourse context. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1525–1534, Berlin, Germany, 2016. Association for Computational Linguistics.
- [46] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA, 2002. Association for Computational Linguistics.
- [47] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8), 2019.
- [48] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67, 2020.
- [49] Siva Reddy, Danqi Chen, and Christopher D. Manning. CoQA: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266, 2019.

- [50] Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adversarial winograd schema challenge at scale. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8732–8740. AAAI Press, 2020.
- [51] Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal V. Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Févry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M. Rush. Multitask prompted training enables zero-shot task generalization. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [52] Rico Sennrich, Barry Haddow, and Alexandra Birch. Edinburgh neural machine translation systems for WMT 16. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 371–376, Berlin, Germany, 2016. Association for Computational Linguistics.
- [53] Zhengxiang Shi, Yue Feng, and Aldo Lipani. Learning to execute actions or ask clarification questions. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2060–2070, Seattle, United States, 2022. Association for Computational Linguistics.
- [54] Zhengxiang Shi and Aldo Lipani. Don’t stop pretraining? make prompt-based fine-tuning powerful learner. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [55] Zhengxiang Shi and Aldo Lipani. DePT: Decomposed prompt tuning for parameter-efficient fine-tuning. In *The Twelfth International Conference on Learning Representations*, 2024.
- [56] Zhengxiang Shi, Francesco Tonolini, Nikolaos Aletras, Emine Yilmaz, Gabriella Kazai, and Yunlong Jiao. Rethinking semi-supervised learning with language models. In *Findings of ACL 2023*, Toronto, Canada, 2023. Association for Computational Linguistics.
- [57] Zhengxiang Shi, Qiang Zhang, and Aldo Lipani. Stepgame: A new benchmark for robust multi-hop spatial reasoning in texts. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11321–11329, Jun. 2022.
- [58] Zhengyan Shi, Sander Land, Acyr Locatelli, Matthieu Geist, and Max Bartolo. Understanding likelihood over-optimisation in direct alignment algorithms, 2024.
- [59] Shivalika Singh, Freddie Vargus, Daniel Dsouza, Börje F Karlsson, Abinaya Mahendiran, Wei-Yin Ko, Herumb Shandilya, Jay Patel, Deividas Mataciunas, Laura OMahony, et al. Aya dataset: An open-access collection for multilingual instruction tuning. *ArXiv preprint*, abs/2402.06619, 2024.
- [60] Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc Le, Ed Chi, Denny Zhou, and Jason Wei. Challenging BIG-bench tasks and whether chain-of-thought can solve them. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, editors, *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13003–13051, Toronto, Canada, 2023. Association for Computational Linguistics.
- [61] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model, 2023.
- [62] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *ArXiv preprint*, abs/2307.09288, 2023.

- [63] Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khachabi, and Hannaneh Hajishirzi. Self-instruct: Aligning language models with self-generated instructions. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13484–13508, Toronto, Canada, 2023. Association for Computational Linguistics.
- [64] Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Atharva Naik, Arjun Ashok, Arut Selvan Dhanasekaran, Anjana Arunkumar, David Stap, Eshaan Pathak, Giannis Karamanolakis, Haizhi Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krima Doshi, Kuntal Kumar Pal, Maitreya Patel, Mehrad Moradshahi, Mihir Parmar, Mirali Purohit, Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma, Ravsehaj Singh Puri, Rushang Karia, Savan Doshi, Shailaja Keyur Sapat, Siddhartha Mishra, Sujan Reddy A, Sumanta Patro, Tanay Dixit, and Xudong Shen. Super-NaturalInstructions: Generalization via declarative instructions on 1600+ NLP tasks. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5085–5109, Abu Dhabi, United Arab Emirates, 2022. Association for Computational Linguistics.
- [65] Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. Finetuned language models are zero-shot learners. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [66] Bin Wu, Jinyuan Fang, Xiangxiang Zeng, Shangsong Liang, and Qiang Zhang. Adaptive compositional continual meta-learning. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 37358–37378. PMLR, 23–29 Jul 2023.
- [67] Bin Wu, Zaiqiao Meng, Qiang Zhang, and Shangsong Liang. Meta-learning helps personalized product search. In *Proceedings of the ACM Web Conference 2022, WWW ’22*, page 2277–2287, New York, NY, USA, 2022. Association for Computing Machinery.
- [68] Mengzhou Xia, Sadhika Malladi, Suchin Gururangan, Sanjeev Arora, and Danqi Chen. Less: Selecting influential data for instruction tuning, 2024.
- [69] Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, Qingwei Lin, and Daxin Jiang. WizardLM: Empowering large pre-trained language models to follow complex instructions. In *The Twelfth International Conference on Learning Representations*, 2024.
- [70] Yang Xu, Yongqiang Yao, Yufan Huang, Mengnan Qi, Maoquan Wang, Bin Gu, and Neel Sundaresan. Rethinking the instruction quality: Lift is what you need, 2023.
- [71] Fuzhao Xue, Yao Fu, Wangchunshu Zhou, Zangwei Zheng, and Yang You. To repeat or not to repeat: Insights from scaling LLM under token-crisis. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [72] Adam X. Yang, Maxime Robeyns, Thomas Coste, Jun Wang, Haitham Bou-Ammar, and Laurence Aitchison. Bayesian reward models for llm alignment, 2024.
- [73] Adam X Yang, Maxime Robeyns, Xi Wang, and Laurence Aitchison. Bayesian low-rank adaptation for large language models. In *ICLR*, 2024.
- [74] Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. HellaSwag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy, 2019. Association for Computational Linguistics.
- [75] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. *ArXiv preprint*, abs/2205.01068, 2022.

- [76] Hao Zhao, Maksym Andriushchenko, Francesco Croce, and Nicolas Flammarion. Long is more for alignment: A simple but tough-to-beat baseline for instruction fine-tuning, 2024.
- [77] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging LLM-as-a-judge with MT-bench and chatbot arena. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023.
- [78] Chunting Zhou, Pengfei Liu, Puxin Xu, Srinu Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, LILI YU, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. LIMA: Less is more for alignment. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [79] Han Zhou, Xingchen Wan, Lev Proleev, Diana Mincu, Jilin Chen, Katherine A Heller, and Subhrajit Roy. Batch calibration: Rethinking calibration for in-context learning and prompt engineering. In *The Twelfth International Conference on Learning Representations*, 2024.
- [80] Han Zhou, Xingchen Wan, Ivan Vulić, and Anna Korhonen. AutoPEFT: Automatic Configuration Search for Parameter-Efficient Fine-Tuning. *Transactions of the Association for Computational Linguistics*, 12:525–542, 2024.

Appendix Overview

The appendix is structured as follows:

Appendix §A provides a brief description (with statistical summaries) for instruction tuning datasets.

Appendix §B provides details of evaluation benchmarks and settings.

Appendix §C provides the experimental setting, implementation details and hyperparameters for all comparison methods used in our experiments.

Appendix §D provides the supplementary experimental results to investigate the effect of our approach on training and testing losses.

Appendix §E provides the supplementary experimental results to investigate the relationship between the win rate on the AlpacaEval 1.0 and the number of epochs.

Appendix §F provides the mathematical formula for the Kullback-Leibler (KL) divergence used in our paper.

Appendix §G provides the supplementary experimental results to investigate the relationship between the output length and the number of epochs.

A Instruction Tuning Dataset

In this work, we use 13 popular datasets from previous instruction tuning research. For the WizardLM, Sharegpt, Science Literature, and Code Alpaca datasets, we directly use the subset provided by the previous work [30]. Refer to the dataset statistics in Table 5. In addition, we provide an analysis of the output length distribution for LIMA, Alpapasus Dolly 3k, Alpapasus Dolly 9k, Alpapasus Alpaca 5k, Less MMLU Chat, Less Tydiqa, and Less BBH ICL datasets, as shown in Figure 7.

Table 5: Statistical summary for various instruction tuning datasets. The table includes sample sizes, the average total length of instructions and outputs, the average output length, and the average instruction length with their standard deviations, and ratio calculations.

Dataset	Size	Total	Output	Output Std	Instruction	Instruction Std	Output/Instruction	Instruction/Output
LIMA	1,030	484.47	442.75	491.34	41.72	79.28	10.6124	0.0942
Less MMLU Chat	13,533	225.19	8.24	16.42	216.95	301.64	0.0380	26.3316
Less Tydiqa	13,533	172.44	25.13	42.62	147.31	235.37	0.1706	5.862
Less BBH ICL	13,533	262.03	61.44	92.55	200.60	196.79	0.3063	3.265
Alpapasus Dolly 3k	2,996	111.91	68.08	106.38	43.83	107.53	1.5530	0.6439
Alpapasus Dolly 9k	9,229	73.40	56.62	48.91	16.79	11.33	3.3727	0.2965
Alpapasus Alpaca 5k	5,305	48.29	30.81	34.44	17.48	12.45	1.7631	0.5672
Tulu V2	326,181	541.16	343.56	575.32	197.60	345.99	1.7387	0.5751
Tulu V2 (10%)	32,618	517.45	338.96	562.74	178.49	345.72	1.8991	0.5266
Tulu V2 (50%)	163,090	515.63	340.67	571.06	174.97	343.45	1.9470	0.5136
Tulu V2 (20%)	65,236	504.56	336.89	562.46	167.68	331.24	2.0092	0.4977
WizardLM	30,000	350.05	258.35	182.98	91.71	86.09	2.8170	0.3550
Sharegpt	50,000	1035.39	831.15	757.10	204.24	344.51	4.0695	0.2457
Science Literature	7,544	1196.08	46.46	57.34	1149.62	905.99	0.0404	24.7417
Stanford Alpaca	52,002	63.77	45.18	44.97	18.59	12.42	2.4302	0.4115
Code Alpaca	20,022	49.74	27.40	27.35	22.34	10.67	1.2262	0.8156

B Evaluation Datasets and Details

We use the open-source repositories, LM-Evaluation Harness³ and Huggingface Dataset⁴ as the evaluation tools. We describe our evaluation setup below:

³<https://github.com/EleutherAI/lm-evaluation-harness>

⁴<https://huggingface.co/docs/datasets>

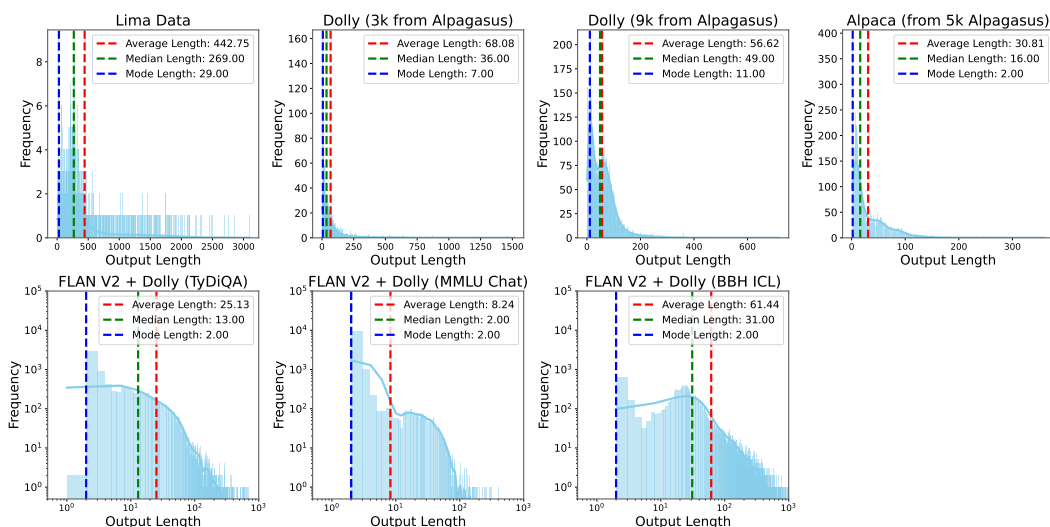


Figure 7: Distribution of output lengths of instruction tuning datasets. This figure presents histograms for the distribution of output lengths across seven datasets, including LIMA, Alpapasus Dolly 3k, Alpapasus Dolly 9k, Alpapasus Alpaca 5k, Less MMLU Chat, Less Tydiqa, and Less BBH ICL. Each subplot displays the frequency of output lengths with key statistical indicators: the average (red dashed line), median (green dashed line), and mode (blue dashed line) of each dataset. The last three subplots employ a logarithmic scale on both axes to better illustrate data spread.

MMLU. We evaluate the model using the dataset at the huggingface dataset ⁵. We follow the protocol outlined in HuggingFace Open LLM Leaderboard ⁶. The evaluation uses multiple-choice questions formatted as the question followed by four choices (A, B, C, D) and prompting for an answer. We calculate the mean accuracy (acc) across test examples.

BBH. The model evaluation utilises the dataset at the huggingface dataset ⁷, specifically tested on the ‘test’ split without the use of few-shot examples. We follow the setup in previous works [30, 60]. The evaluation metric is the exact match score, averaged (mean) to assess performance. Generation is constrained to a maximum of 1024 tokens, with termination upon encountering specific delimiters such as “</s>”, “Q”, or double newlines. The generation is greedy decoding (temperature set to 0.0) and does not use sampling. Answer extraction employs regex patterns to identify responses immediately following “the answer is” and captures only the first occurrence.

GSM8K. We evaluate using the dataset at the huggingface dataset ⁸, focusing on arithmetic problem-solving in the ‘test’ split. We follow the HuggingFace Open LLM Leaderboard to 8 few-shot examples. Exact match is the chosen metric, with case insensitivity and select regex-based filtering of common punctuation and formatting characters to ensure precise validation of numerical answers. The primary focus is on extracting and comparing the final numerical answer to the model’s output using a strict regex-based match setup.

HumanEval. We evaluate using the dataset and the evaluation code from the previous work [30]. We report the performance of the pass@1. We perform the decoding using two different temperatures, 0.1 and 0.7. We report the better pass@1 from these two decoding results.

⁵https://huggingface.co/datasets/hails/mmlu_no_train

⁶https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard

⁷<https://huggingface.co/datasets/lucaemon/bbh>

⁸<https://huggingface.co/datasets/gsm8k>

ARC. The evaluation setup for the dataset at the huggingface dataset ⁹ utilises a multiple-choice format. We follow the HuggingFace Open LLM Leaderboard to 25 few-shot examples. The performance metric used is mean normalised accuracy (acc_norm).

CoQA. We conduct the model evaluation on the dataset at the huggingface dataset ¹⁰. We follow the HuggingFace Open LLM Leaderboard to 0 few-shot examples. The output generation terminates upon encountering a new line followed by "Q:". The mean F1 score is used as the evaluation metric.

PIQA. Evaluation on the dataset at the huggingface dataset ¹¹ involves a multiple-choice. The evaluation incorporates 10 few-shot examples, according to the LIMIT [32]. Performance is measured using the mean normalized accuracy (acc_norm).

OpenBookQA. The dataset at the huggingface dataset ¹² is evaluated in a multiple-choice format. The mean normalized accuracy (acc_norm) is used as the evaluation metric.

LAMBADA. The evaluation of the model on the dataset at the huggingface dataset ¹³ is performed using a loglikelihood output type. The mean accuracy is used as the evaluation metric.

HellaSwag. In the ‘hellaswag’ dataset at the huggingface dataset ¹⁴, model evaluation is conducted using a multiple-choice format. We follow the HuggingFace Open LLM Leaderboard to 10 few-shot examples. The mean normalized accuracy (acc_norm) is used as the evaluation metric.

The Winograd Schema Challenge. The evaluation is conducted using a multiple-choice format on the ‘test’ split at the huggingface dataset ¹⁵. The mean accuracy is used as the evaluation metric.

Winogrande. The ‘winogrande’ dataset is assessed using a multiple-choice format at the huggingface dataset ¹⁶. We follow the HuggingFace Open LLM Leaderboard to 5 few-shot examples. The mean accuracy is used as the evaluation metric.

LAMBADA. For this dataset, evaluation is conducted using the loglikelihood output type on the ‘test’ split at the huggingface dataset ¹⁷. This variant focuses on predicting the last word of text passages in English. The mean accuracy is used as the evaluation metric.

Translation Benchmarks WMT. The evaluation of the translation capabilities is performed on the WMT 2014¹⁸ and WMT 2016¹⁹ datasets at the huggingface dataset. Here we use the ‘ter’ score as the evaluation metric.

TruthfulQA. We use the dataset at the huggingface dataset ²⁰. We follow the setup at the HuggingFace Open LLM Leaderboard using the 6 few-shot examples. The mean accuracy is used as the evaluation metric.

ToxiGen. We use the dataset at the huggingface dataset ²¹. The task is assessed using a multiple-choice framework to evaluate the model’s capability to identify hateful content in text statements. The mean accuracy is used as the evaluation metric.

⁹https://huggingface.co/datasets/allenai/ai2_arc

¹⁰<https://huggingface.co/datasets/EleutherAI/coqa>

¹¹<https://huggingface.co/datasets/piqa>

¹²<https://huggingface.co/datasets/openbookqa>

¹³<https://huggingface.co/datasets/lambada>

¹⁴<https://huggingface.co/datasets/hellaswag>

¹⁵https://huggingface.co/datasets/winograd_wsc

¹⁶<https://huggingface.co/datasets/winogrande>

¹⁷https://huggingface.co/datasets/EleutherAI/lambada_openai

¹⁸<https://huggingface.co/datasets/wmt14>

¹⁹<https://huggingface.co/datasets/wmt16>

²⁰https://huggingface.co/datasets/truthful_qa

²¹<https://huggingface.co/datasets/skg/toxigen-data>

Hendrycks Ethics. We use the dataset at the huggingface dataset ²², with a multiple-choice format. The model aims to detect whether described actions in various contexts are ethically wrong. The prompt format integrates a specific scenario followed by a structured question: "Is this wrong?" and then prompts for an answer with options 'no' or 'yes'. The mean accuracy is used as the evaluation metric.

C Implementation Details

Experimental Design for Figure 2 Left. Here we present a detailed experimental design for Figure 2 Left. We perform experiments on a variety of datasets, including LIMA, Alpapasus Dolly 3k, Alpapasus Dolly 9k, Alpapasus Alpaca 5k, Less MMLU Chat, Less Tydiqa, Less BBH ICL, Tulu V2, Code Alpaca, Stanford Alpaca, Science Literature, WizardLM, and Sharegpt. Furthermore, to evaluate the effectiveness of IM on datasets with different instruction-to-output length ratios, we select three subsets from Tulu V2. Each subset contains 3,000 training examples, with instruction-to-output length ratios of approximately 5, 10, and 15, respectively.

Experimental Design for Figure 2 Right. Here we provide a detailed experimental design for Figure 2 Right. We strategically sampled varying sizes of training examples from the Tulu V2 dataset to investigate the effectiveness of IM with different sizes of training examples. Starting with approximately 320,000 examples in the Tulu V2 dataset, we create subsets of data ranging from as few as 1,000 to as many as 35,000 examples. These subsets were selected randomly, ensuring a representative mix across different scales. We adhered to a fixed instruction-to-output length ratio of approximately 10 to maintain consistency in training conditions across all samples. We train the LLAMA-2-7B-BASE on all these subsets and evaluate them respectively.

Table 6: Hyperparameters and configurations for supervised fine-tuning.

Hyperparameter	Assignment
GPUs	2 or 4 A100 80G GPUs, 2 48G A6000 GPUs
Batch size per GPU	1
Total batch size	128
Number of epochs	2, 3, or 10
Maximum sequence length	2048
Learning rate	2×10^{-5}
Learning rate optimizer	AdamW
Adam epsilon	1e-6
Adam beta weights	0.9, 0.98
Learning rate scheduler	Linear with warmup
Warmup proportion	0.03
Weight decay	0
Mixed precision	bf16
Gradient accumulation steps	Calculated dynamically

Implementation Details. In our study, we fine-tune the LLaMA-2-7B, LLaMA-2-13B and OPT-6.7 model using four A100 80G GPUs, with a per-GPU batch size of 1 and a total batch size of 128, employing a learning rate of $2e-5$. Training typically proceeds for 2 epochs with a maximum sequence length of 2048 tokens. We utilise gradient accumulation, calculated to effectively distribute training steps across the available hardware, resulting in larger batch sizes despite hardware limitations. We employ mixed precision (bf16), linear learning rate scheduling with a warm-up ratio of 0.03, and a weight decay of 0. To optimise our training, we use DeepSpeed with a stage 3 configuration without offloading. Our setup also includes the usage of Flash Attention [19] and slow tokenization

²²https://huggingface.co/datasets/ElleutherAI/hendrycks_ethics

to enhance training efficiency and compatibility. Our code is implemented using Open-Instruct²³, Pytorch²⁴ and Huggingface²⁵. Table 6 lists the hyperparameters.

D Train and Test Loss

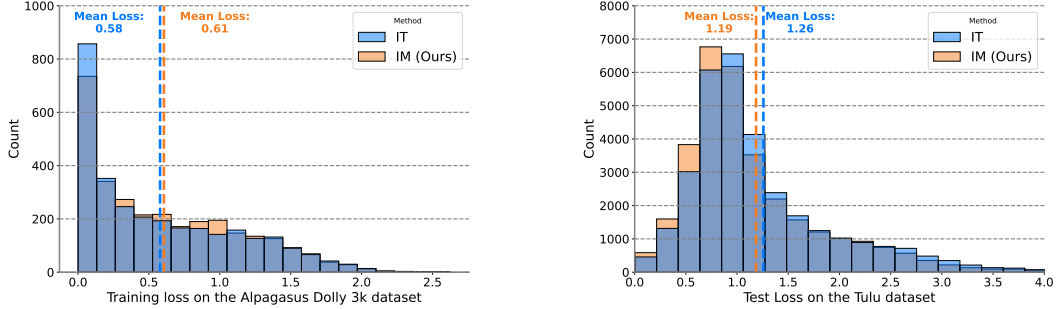


Figure 8: **(Left)** Training loss distribution for each example between our approach INSTRUCTION MODELLING (IM) and INSTRUCTION TUNING (IT) on the Alpapasus Dolly 3k dataset. **(Right)** Test loss distribution for each example between IM and IT on the Tulu V2 dataset, using a 10% sampled data. Mean losses are marked by dashed lines. For both IM and IT, here we only compute the loss over the output part. IM has a higher train loss with lower test loss, suggesting that IM effectively mitigates the overfitting issues compared to IT.

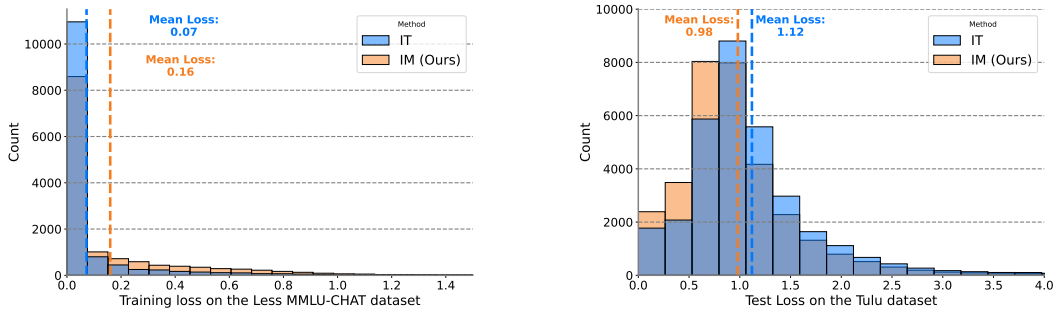


Figure 9: **(Left)** Training loss distribution for each example between our approach INSTRUCTION MODELLING (IM) and INSTRUCTION TUNING (IT) on the Less MMLU Chat dataset. **(Right)** Test loss distribution for each example between IM and IT on the Tulu V2 dataset, using a 10% sampled data. Mean losses are marked by dashed lines. For both IM and IT, here we only compute the loss over the output part. IM has a higher train loss with lower test loss, suggesting that IM effectively mitigates the overfitting issues compared to IT.

In this section, we provide additional experiments regarding training and testing loss distributions. Figure 8 focuses on the Alpapasus Dolly 3k and Tulu V2 datasets, displaying how IM tends to exhibit higher training losses yet achieves lower test losses compared to IT. Similarly, Figure 9 compares these methods on the Less MMLU Chat and Tulu V2 datasets under analogous conditions.

E The impact of Epochs on the Win Rate

Figure 10 illustrates the comparative analysis of AlpacaEval 1.0 scores across different epochs for two datasets, LIMA and Alpapasus Dolly 9k datasets. We evaluate the performance of IM and IT

²³<https://github.com/allenai/open-instruct>

²⁴<https://pytorch.org/>

²⁵<https://huggingface.co/>

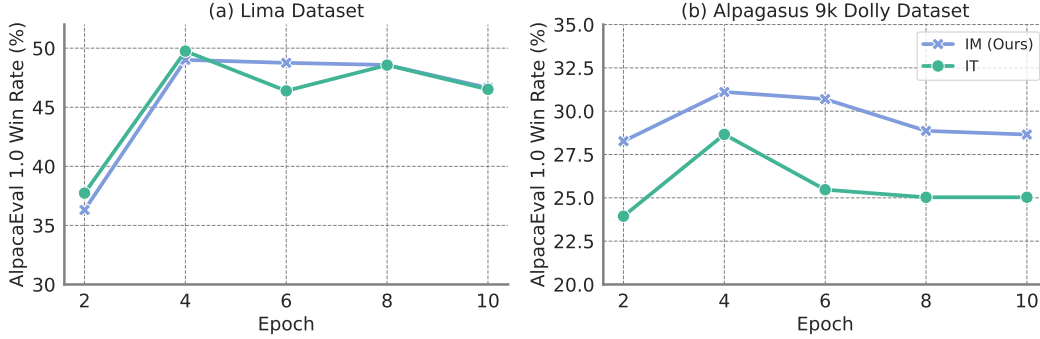


Figure 10: AlpacaEval 1.0 performance trends for IM and IT approaches on the LIMA and Alpapasus Dolly 9k datasets across different epochs.

over different numbers of epochs. IM consistently surpasses IT in performance on the Alpapasus Dolly 9k dataset, while the performance of both approaches is comparable on the LIMA dataset.

F Applying KL Divergence Loss for Instruction Tuning

In this section, we first briefly introduce the Kullback-Leibler (KL) divergence and then introduce the experimental details. Future work can investigate the effectiveness of parameter-efficient fine-tuning approaches as the regularisation [80, 55, 28].

Kullback-Leibler Divergence. Kullback-Leibler (KL) divergence is commonly employed as a regularisation method in the fine-tuning of LMs, helping to mitigate overfitting by constraining the fine-tuned model to remain similar to the pre-trained model [44]. Specifically, the KL divergence is added to the fine-tuning objective as a per-token regularisation term between the fine-tuned LM $\pi_\theta(x)$, and the pre-trained LM, $\pi^{\text{pre}}(x)$. For supervised fine-tuning with the next token prediction loss, the training objective incorporating KL divergence is computed as follows:

$$\mathcal{L}_{\text{KL}}(\theta) = \mathbb{E}_{x \sim \mathcal{D}} \left[\sum_t -\log \pi_\theta(x_t | x_{0:t-1}) + \lambda \sum_t \text{KL}(\pi_\theta(x_t | x_{0:t-1}) || \pi^{\text{pre}}(x_t | x_{0:t-1})) \right], \quad (5)$$

where λ is a regularisation parameter that balances the loss due to the next token prediction and the KL divergence, and $\pi(x_t | x_{0:t-1})$ represents the next token distribution of the fine-tuned or pre-trained LM conditioned on the preceding context.

Table 7: Performance on 18 NLP tasks and AlpacaEval 2.0, with various values of λ , trained on the (LLAMA-2-7B-BASE).

	NLP Tasks	AlpacaEval 2.0
LLAMA-2-7B-BASE	49.32	0.01
$\lambda = 0.01$	48.81	2.58
$\lambda = 0.1$	48.77	2.44
$\lambda = 1.0$	49.26	0.06

Ablation study on the effect of λ . In Table 3, we set the value of the λ as 1.0. Here we provide additional experiments with different values of λ . Table 7 presents the model performance on the NLP tasks and AlpacaEval 2.0. This aligns our observations in §4.3.

G The impact of Epochs on Output Lengths

Figure 11 illustrates the average output length of various models across different epochs. We report the output length on four different datasets, including Alpapasus Dolly 3k, Alpapasus Dolly

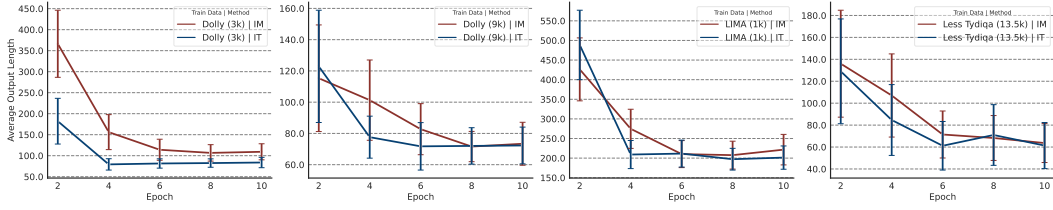


Figure 11: Comparative analysis of output lengths for IM and IT across different epochs on Alpagasus Dolly 3k, Alpagasus Dolly 9k, LIMA, and Less Tydiqa datasets.

9k, LIMA, and Less Tydiqa. Each line represents the average output length of a model, with epochs ranging from 2 to 10, and is accompanied by error bars that denote the normalised standard deviation (10%) of the output lengths. Our experimental results show that our approach IM does not consistently increase the output length and that win rates are not necessarily associated with the length of the output.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS paper checklist",**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: [In the introduction, we explicitly detail our contributions and provide hyperlinks that direct readers to the corresponding sections of our paper.](#)

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: [Please refer to §5.](#)

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: [Our paper does not include theoretical results.](#)

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: [Please refer to Appendix §A for training data details, Appendix §B for evaluation details, Appendix §C for implementation details. We also provide all the codes at <https://github.com/ZhengxiangShi/InstructionModelling>, with comprehensive instructions to reproduce our work, including details about the coding environment, data, training procedures, evaluation, and analysis.](#)

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Please refer to [Appendix §A for training data details](#), [Appendix §B for evaluation details](#), [Appendix §C for implementation details](#). We also provide all the codes at <https://github.com/ZhengxiangShi/InstructionModelling>, with comprehensive instructions to reproduce our work, including details about the coding environment, data, training procedures, evaluation, and analysis.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.

- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Please refer to Appendix §A for training data details, Appendix §B for evaluation details, Appendix §C for implementation details. We also provide all the codes at <https://github.com/ZhengxiangShi/InstructionModelling>, with comprehensive instructions to reproduce our work, including details about the coding environment, data, training procedures, evaluation, and analysis.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We report error bars for some results (see Figure 6). For other experiments, error bars are not reported because it would be too computationally expensive.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: [Please refer to Appendix §C for implementation details, including GPUs we used in this paper.](#)

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: [Our research conforms, in every respect, with the NeurIPS Code of Ethics.](#)

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g. if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: [Please refer to §5.](#)

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.

- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: [Our paper does not address or focus on topics such as data security, privacy, or ethical considerations regarding the release of sensitive technology.](#)

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: [The creators and original owners of the assets used in the paper are properly credited. This is achieved by acknowledging their contributions directly within the text, referencing their original works, and by including citations to their publications or sources. For example, we acknowledge the contribution of other open-source repositories in our repository at <https://github.com/ZhengxiangShi/InstructionModelling>.](#)

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: [We will open-source our repository at the GitHub repository. Please check our code at https://github.com/ZhengxiangShi/InstructionModelling.](https://github.com/ZhengxiangShi/InstructionModelling)

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: [Our paper does not involve crowdsourcing nor research with human subjects.](#)

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: [Our paper does not involve crowdsourcing nor research with human subjects.](#)

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.