

# ON THE LOSS OF CONTEXT-AWARENESS IN GENERAL INSTRUCTION FINE-TUNING

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Pretrained Large Language Models (LLMs) require post-training methods such as supervised fine-tuning (SFT) on instruction-response pairs to enable instruction following. However, this process can potentially harm existing capabilities learned during pretraining. In this paper, we investigate the loss of context awareness after SFT, defined as the capability to extract and understand information from the user-provided context and respond accordingly. We are the first to identify and show that the loss of context-awareness appears on instruction-finetuned LLMs when the chat template is applied to the input prompts. We identify the performance decline is partially caused by the bias embedded into the chat template to focus less on the user-provided context. Based on these observations, we propose two methods to mitigate the loss of context awareness in instruct models: post-hoc attention steering on user tokens and conditional instruction fine-tuning with a context-dependency indicator. Empirical experiments on 4 context-dependent downstream tasks and 3 pretrained LLMs of different sizes show that our methods effectively mitigates the loss of context awareness without compromising the general ability to follow instructions. Our findings also strongly advocate the necessity to carefully benchmark context awareness after instruction fine-tuning.

## 1 INTRODUCTION

Large Language Models (LLMs) pretrained on large-scale datasets acquire a diverse set of language modeling capabilities in pretraining. To enhance these models' ability to follow general instructions, further fine-tuning is typically required such as supervised instruction fine-tuning (SFT) (Wei et al., 2021; Ouyang et al., 2022) and reinforcement learning from human feedback (RLHF) (Christiano et al., 2017) to better understand and respond to human requests. However, the additional fine-tuning can potentially harm existing capabilities learned in pretraining, as pointed out by several existing works (Lin et al., 2024; Bai et al., 2022; Fu et al., 2024).

In this paper, we particularly investigate the loss of context-awareness after instruction fine-tuning, which is the capability to understand and retrieve exact information from the user-provided context and respond accordingly. Context awareness is crucial for many real-world use cases, including retrieval augmented generalization (RAG) (Khandelwal et al., 2020; Izacard et al., 2023; Xu et al., 2023b), in-context learning (Agarwal et al., 2024), and contextual question-answering (QA) (Rajpurkar et al., 2016; Choi et al., 2018; Dua et al., 2019). We first illustrate the loss of context awareness in Figure 2 with the Needle-in-a-Haystack test on four popular instruction-tuned models. We demonstrate that the performance degradation is consistent on both long-context and relatively short-context LLMs, which cannot be solely explained by the distribution difference in context lengths between the instruction dataset and the evaluation benchmarks, as suggested by prior works (Dubey et al., 2024).

We identify that the bias embedded within the chat template to focus less on the user tokens is a major cause of context-awareness degradation. Normal instruction fine-tuning dataset contains a mixture of both model-dependent and context-dependent queries. The former consists of queries the model can respond to relying mostly on its **internal knowledge** learned during pretraining. On the other hand, responding to the second type of query requires exact information retrieval and processing from the **user-provided context** in the input prompt, such as in-context learning, long-form instruction-following, and contextual QA tasks. However, a query accompanied by a context

054 can still be a model-dependent query as the model may have learned the context during pretraining,  
055 and is able to respond without relying on the user given context. Therefore, it is challenging for  
056 the model to differentiate between these two type of queries from the prompt only, and incorrect  
057 identification could lead to hallucination or being over-reliant on user-provided context

058 We validate the bias embedded in chat templates with the Needle-In-a-Haystack test (NIH) (Kam-  
059 radt, 2023), which requires a model to retrieve a given text “needle” from a long paragraph of  
060 irrelevant text. Our experiments show that the needle retrieval performance drops on instruction-  
061 finetuned models *only* when the chat template is applied to the input, which is however crucial for  
062 the model to distinguish different roles in conversations. We further show that the performance drop  
063 can be attributed to the drop in attention value allocated to the whole user input section. There-  
064 fore, the context retrieval capability is not lost in the model, but “inhibited” by the chat-formatted  
065 fine-tuning.

066 Based on these observations, a straightforward approach is to directly steer the attention value during  
067 inference time to emphasize the user inputs on instruction-tuned models. This can be achieved by  
068 manually intervening attention scores of the user context tokens during generation. Experiments  
069 show that while performance on simple retrieval tasks can be significantly boosted by attention  
070 steering, manipulating the attention value in the inference stage can harm other capabilities of the  
071 model, deteriorating performance on more complex tasks.

072 To further improve upon the undesirable trade-off of post-hoc techniques, we are motivated to steer  
073 the attention allocation in the fine-tuning stage. To achieve this goal, we identify context-dependent  
074 conversations before fine-tuning and add a special token to the prompt as a hint to the model. The  
075 special token can then be added at inference time when more attention is demanded to be allocated  
076 to the user-provided context. Empirical experiments show the effectiveness of our method on several  
077 open-source, pretrained LLMs and instruction fine-tuning datasets.

078 Our contributions are summarized as follows:

- 080 • We identify that supervised instruction fine-tuning causes pretrained language models to  
081 deteriorate in context awareness (even for short context lengths).
- 082 • We pinpoint the worsened context awareness is associated with attention allocation bias  
083 embedded within the chat template.
- 084 • We propose an inference-time technique to partially recover the context-awareness of gen-  
085 eral instruction-tuned language models by manually intervening attention scores.
- 086 • We propose a training-time technique utilizing conditional indicators to mitigate the loss of  
087 context awareness of pretrained language models during instruction-tuning.

## 090 2 RELATED WORK

091  
092 **Instruction fine-tuning and chat templates.** Large language models only learn language model-  
093 ing on general corpus during pretraining. To enable instruction-following, they usually require su-  
094 pervised fine-tuning on instruction-following datasets (SFT) (Wei et al., 2021; Ouyang et al., 2022),  
095 followed by reinforcement learning with human feedback (RLHF) (Christiano et al., 2017). In this  
096 paper, we mainly focus on the SFT stage. Instruction fine-tuning datasets consist of user instruc-  
097 tion and target model response pairs, which can be collected from modified NLP tasks (Wei et al.,  
098 2021; Longpre et al., 2023), human annotations (Ouyang et al., 2022; Chiang et al., 2023) or synthe-  
099 sized data from existing LLMs (Ding et al., 2023; Xu et al., 2023a). Instruction fine-tuning usually  
100 converts training examples into a dialog format with a chat template, which consists of a user role  
101 indicator (e.g. [INST] in llama-2 models), an assistant role indicator (e.g. [/INST] in llama-2  
102 models) and an optional system role indicator (e.g. <<SYS>> in llama-2 models). Other role indi-  
103 cators are also used in more complicated scenarios such as tool using (Schick et al., 2023). However,  
104 these role indicators and role partition in the conversation are not sufficiently presented and learned  
105 during pretraining, making them prone to bias during fine-tuning.

106 **Evaluation and improvement of context-awareness** The capability to retrieve and understand  
107 information from the context and respond accordingly is important for many tasks including in-  
context learning (Agarwal et al., 2024; Brown et al., 2020), retrieval augmented generation (Lewis

et al., 2020) and contextual QA (Dua et al., 2019; Rajpurkar et al., 2016; Choi et al., 2018). These tasks are commonly included in standard evaluation benchmarks of recent LLMs. To improve the performances on context-dependent tasks, a common practice is to collect or synthesize context-dependent data, especially contextual QA data, and mix them into the fine-tuning dataset (An et al., 2024; Dubey et al., 2024). Apart from adding more data, Hsieh et al. (2024) also explores calibrating the attention weights to compensate for the drop of attention weight in the middle of a long context. However, to date, limited existing works have examined how chat templates affect context-awareness when fine-tuning language models on conversational instruction-following data. The most relevant conclusion from existing works is mentioned by Dubey et al. (2024) that long-context capabilities learned during pretraining drops significantly after SFT. However, they only attribute this performance decline to a lack of long-context data in SFT stage.

**Conditional Supervised Fine-Tuning** Conditional Supervised Fine-Tuning (CSFT) involves fine-tuning a pretrained model on specific tasks while conditioning on additional information or context. The conditions represent specific attributes or styles of the demonstration response and are concatenated to the original input as prefixes. The model is expected to associate the condition prefix with the style or attribute such that whenever the condition prefix is added, the model generates responses accordingly. Dong et al. (2023) proposes SteerLM, where they applied CSFT to align models with human values (e.g. helpfulness, humor, and creativity) by conditioning the model on attribute prefixes. Korbak et al. (2023) adds two control tokens as indicators for good and bad demonstration responses. As opposed to attribute tags or control tokens, Chain of Hindsight (Liu et al., 2023) conditions demonstration responses on natural language comments on the quality or styling of the response.

### 3 LOSS OF CONTEXT-AWARENESS EMBEDDED IN CHAT TEMPLATES

#### 3.1 LOSS OF CONTEXT-AWARENESS IN NEEDLE-IN-A-HAYSTACK (NIH) TEST

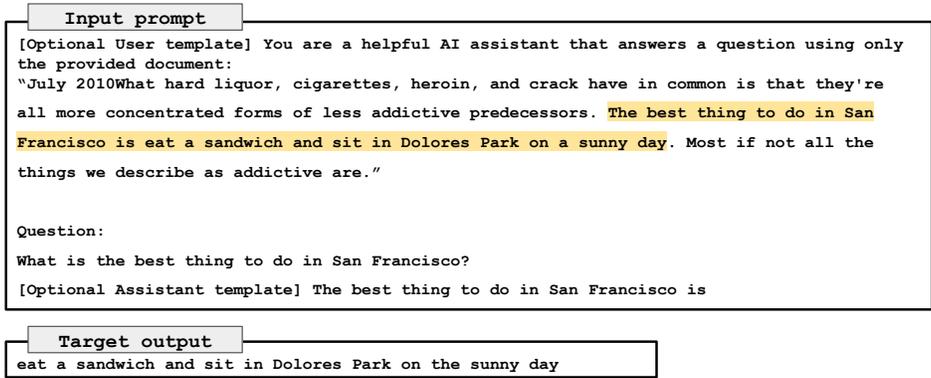


Figure 1: An example of the Needle-in-a-haystack (NIH) test used in our work. [Optional User template] and [Optional Assistant template] are user and assistant role indicators used in instruction finetuned models. The inserted needle is highlighted in yellow.

In this section, we show the loss of context awareness after instruction fine-tuning with the needle-in-a-haystack (NIH) test.

**Experimental settings.** NIH evaluates the performance of language models to extract a given sentence (the needle) from an irrelevant context. The needle can be inserted at different locations in contexts with different lengths. We report the recall error:

$$err = \frac{1}{|K|} \sum_{w \in K} \mathbb{1}(w \in output)$$

where  $K$  is the set of keywords in the targeted output and  $output$  is the output of the LLM. We average the recall error across 400 NIH tests with different insertion locations and context lengths within the model’s context window. More details about the NIH tests can be found in Appendix A.2.

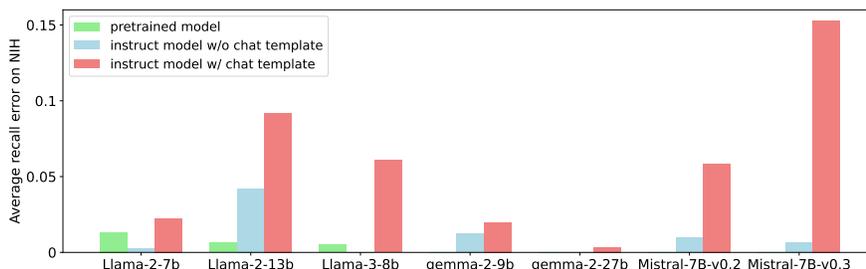


Figure 2: Average recall error (1 - recall) on NIH for different model series. We report the performance of official instruct models (with and without chat templates) and their corresponding pre-trained models from 5 model families with sizes ranging from 7B to 27B. Recall errors for some models are too small to be visible from the figure. Detailed numerical values can be found in Appendix B.1

We evaluate the results on four open-source pretrained language models (not instruction-tuned) and their corresponding instruction-finetuned versions from Llama and Mistral series. Here we don't consider stronger close-source models as their pretrained versions are not available. The context window lengths of these models range from 4,096 to 32,768. We show a typical NIH prompt used in our experiments in Figure 1. When the chat template is applied to the prompt, the whole input prompt is partitioned into the user instruction input and model response, indicated by special role markers in the chat template (e.g. `<|user|>` and `<|assistant|>`). These chat templates are learned only in the instruction fine-tuning data, to teach the model to perform dialog conversation in response to the user instructions.

**Results.** We report the results in Figure 2. As we can see from the figure, the recall error increases significantly on instruction finetuned models when the chat template is added. When the chat template is removed, the recall error on the instruction-tuned model is comparable to or even better than pretrained models. These results indicate that the context retrieval capabilities are not wiped away in the model during instruction fine-tuning, but instead impacted by the bias embedded in the chat templates. However, simply removing the chat templates is not practical as they are necessary for the model to distinguish different roles in conversations. Therefore, we are motivated to recover the context-awareness of instruction-finetuned models by mitigating the negative impacts of chat templates on context-dependent tasks.

Moreover, the aforementioned phenomenon is consistent among models with different context window lengths, different model families and chat templates, and model sizes. In the remaining sections, we will focus on the context awareness of relatively short-context LLMs.

### 3.2 ATTENTION ALLOCATION BIAS FROM CHAT TEMPLATES

Based on the observation that the performance on NIH only drops when the chat template is present, we hypothesize that the chat templates, which are newly introduced in instruction fine-tuning, may encode bias to downweight the importance of user inputs when generating responses.

The distribution of the dataset adopted for pretraining and instruction-tuning is fundamentally different. In the pretraining stage, the model is trained on general-domain texts where all tokens are equally important in the context. However, during instruction fine-tuning, the model learns to distinguish between user input and its responses through chat templates. This user and model partition is not predominantly present in the pretraining corpus, and therefore difficult for the model to benefit from pretraining to learn a generalizable mapping. As a result, the model may tend to memorize the global bias to attend less on the user input if the instruction finetuning dataset is dominated by queries that do not require exact information retrieval.

To validate our hypothesis above, we visualize the attention allocated to each part of the model input in the NIH test, with and without the chat template. An input to an instruction-finetuned LLM

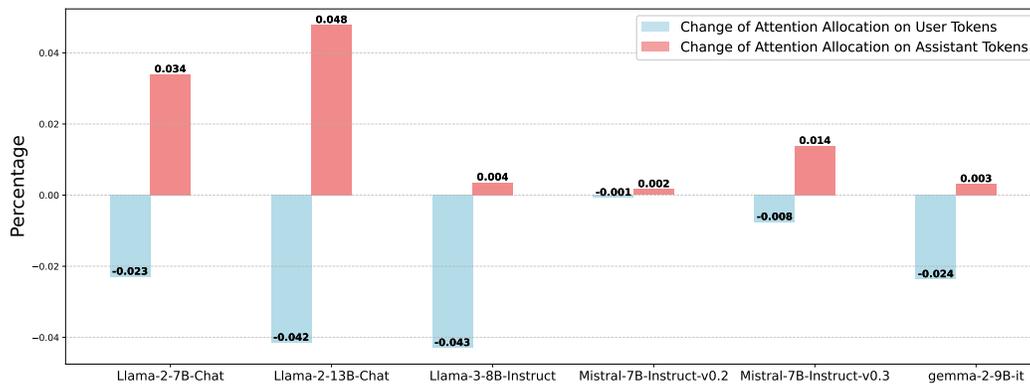


Figure 3: We visualize the changes in attention allocation on user tokens and assistant tokens after applying the chat templates. The attention allocation is calculated when the model is generating the first answer token in its response. For the case where the chat template is applied, we normalize the attention value on user tokens, assistant tokens, and the BOS token such that attention scores allocated to the three sums up to 1. The attention weight is averaged on 400 tests with context lengths ranging from 200 to 4000 and needle depth from 0% to 100%. Details for absolute attention scores allocated to each parts can be found in Appendix B.2.

can be divided into four parts: chat template, including role indicators used to separate user inputs and model response; special tokens, specifically the BOS token denoting the beginning of the text; user input and model response prefix. We calculate and visualize the changes of attention weight allocation on 6 models from 5 model families in Figure 3. Specifically, we calculate the attention weight from the last token before the answer token, which denotes the word “is” in Figure 1, to other tokens in the input prefix. We calculate the attention weight allocation on the most representative head for context retrieval, which is the head with the highest attention weight on the user input. From the figure, we can see a clear trend that the attention weight allocated to user tokens decreases when the chat template is added and attention allocated to assistant tokens increases.

The decreased attention weight allocated to user input context implies the model is relying more on its knowledge when generating the next token. For queries that require information retrieval from the user-provided context, the probability of hallucination also increases.

#### 4 MITIGATING THE LOSS OF CONTEXT-AWARENESS

In the previous sections, we have established that the loss of context awareness to the user input is apparent when the chat template is applied and is directly reflected in the attention allocated. The next step is to mitigate this loss. The first straightforward approach is to directly intervene and increase the attention weight allocated to the user input to compensate for the attention allocation impacted by the chat template. We call this attention steering in this paper, which can be generally performed at inference time, regardless of how the model is fine-tuned. However, since the attention layers encode additional information beyond merely context dependency, manipulating the attention weights (especially with a larger magnitude) may also introduce unintended artifacts.

Therefore, in the second method, we want to encourage the model to focus more on the user input during training for context-dependent queries. Recall from Section 1 that both model-dependent and context-dependent queries can have context provided in the input, the model is unable to reliably distinguish between these two types of queries from the input prompts only. Therefore, we propose a metric to identify context-dependent queries with a seed LLM. We then explicitly inform the model of *when* to attend more to the user-provided context by appending an indicator token to the selected subset of context-dependent instructions during fine-tuning.

#### 4.1 POST-HOC ATTENTION STEERING ON USER INPUT

In this section, we explain in detail how to increase the attention allocation level on the user input with post hoc modification to the attention weight during model inference.

Specifically, on each transformer layer, we modify the self-attention weight:

$$\hat{\text{Att}}(\mathbf{x}, \mathbf{y}) = \begin{cases} \frac{1}{Z} \cdot \text{Att}(\mathbf{x}, \mathbf{y}) & \text{if } \mathbf{y} \in U \\ \frac{1}{Z} \cdot \alpha \text{Att}(\mathbf{x}, \mathbf{y}) & \text{otherwise,} \end{cases} \quad (1)$$

where  $\mathbf{x}$  and  $\mathbf{y}$  are two tokens in the input sequence,  $0 < \alpha < 1$ ,  $U$  is the subset of user input tokens intended to emphasize, and  $Z$  is the normalization constant.

We adopt the same implementation of attention steering as Zhang et al. (2024b). They steered the attention of pretrained language models to emphasize a user-specified portion of the user instruction, enabling models to follow user instructions without explicit instruction fine-tuning. In our setting, we up-weight the attention of instruction fine-tuned models on the whole user input prompt which consequently down-weights other tokens (chat template role tokens, BOS/EOS tokens, and partially generated model responses).

Although post hoc attention steering to improve context-awareness can be generally applied during inference time after a language model has been instruction fine-tuned, it negatively impacts other capabilities of the model that also depend on the attention score. This is a fundamental constraint of post-hoc model editing techniques, as components within a language model simultaneously provide multiple different functions. The trade-off of compromising the performance of other functions limits the magnitude of steering, which in turn limits the effectiveness of improving context awareness.

In the next section, we propose a training-stage method to train a model to better respond to context-dependent queries with an indicator token optionally provided by the user.

#### 4.2 INSTRUCTION FINE-TUNING WITH CONTEXT-DEPENDENCY INDICATORS

Instead of post-hoc attention steering which introduces negative side effects, it is more desirable to directly control the attention during the fine-tuning process. Our goal is to nudge the model during training to teach the model to pay more attention to the user input when an indicator is appended to the user instruction by the user.

Specifically, we process the instruction fine-tuning examples such that responses that demand *more* attention on user instruction context are prepended with a special indicator token in its corresponding user instruction. Thus, after training whenever this indicator token is added and the model generation is conditioned on this token, the model will attend more to the user-provided context.

##### 4.2.1 IDENTIFYING CONTEXT-DEPENDENT EXAMPLES

We first introduce a metric to identify context-dependent training examples from a general instruction fine-tuning dataset. A training sample in the instruction fine-tuning dataset is a conversation between the user and model assistant, which may consist of multiple turns of instruction and response pairs. For a sample with  $n$  turns of instructions and responses, a maximum of  $n$  indicator tokens can be prepended in front of each of the responses to indicate this particular turn of response demands significant user instruction as context.

Let us denote user instructions as  $X$ , assistant responses as  $Y$ , and a conversation of  $n$  total turns as  $C = [X_1, Y_1, \dots, X_n, Y_n]$ . We start by preparing a seed instruction-finetuned model  $M$ , which can be the same or a weaker pretrained model fine-tuned with the original instruction fine-tuning dataset. We then define the context-dependency score for the  $m^{\text{th}}$  turn response  $Y_m$  as follows

$$s_M(Y_m) = \frac{1}{|Y_m|} \sum_{\mathbf{y} \in Y_m} \max_{h \in H} \left( \sum_{\mathbf{x} \in X_1 \cup \dots \cup X_m} \text{Att}_h(\mathbf{y}, \mathbf{x}) \right), \quad (2)$$

where  $H$  is the set of attention heads in model  $M$ . As different heads learn different capabilities, we keep the max user attention weight across all heads to select the most representative head for attention allocation on the user input. The final score measures the sum of attention scores allocated to all user instructions in prior turns  $X_1 \cup \dots \cup X_m$ , averaged over response tokens  $\mathbf{y} \in Y_m$ . In

Datasets	Avg. conversation length	# conversations	# instructions	# instructions in $\hat{S}$
ShareGPT	1,567.68	93,645	331,722	38,542
UltraChat-200k	1,437.33	207,865	657,794	108,646
WizardLM-70K	484.00	57,523	57,523	12,938

Table 1: Statistics of instruction fine-tuning datasets in our experiments. # denotes the number of conversations and individual instructions in the dataset. We report the statistics after performing preprocessing as detailed in Section 5.1.2. Average length is measured in the number of tokens with TinyLlama tokenization.  $\hat{S}$  is the selected subset of context-dependent conversation turns.

practice, we compute the score on a single middle layer for efficiency. We defer the discussion on layer selection to the Appendix B.3.

#### 4.2.2 INSTRUCTION FINE-TUNING WITH CONTEXT-DEPENDENCY INDICATORS

With the context-dependency scores calculated for each model response in the instruction fine-tuning dataset, a threshold  $0 < \beta < 1$  can then be selected such that each conversation turn  $(\mathbf{X}_m, \mathbf{Y}_m)$  with  $s_M(\mathbf{Y}_m) > \beta$  is considered highly context-dependent and added into the subset  $\hat{S}$ .

For each conversation turn  $(\mathbf{X}_m, \mathbf{Y}_m) \in \hat{S}$ , we append a special token [IND] to the user instruction  $\mathbf{X}_m$ . In our implementation, the special token [IND] is added as an additional special token to the vocabulary to avoid conflicting with existing ones. The modified subset  $\hat{S}$  is then reincorporated into the instruction fine-tuning dataset as the final training data. Our method is more general in terms of the task and input format compared to related works, which synthesize context-relevant examples targeted toward specific tasks (e.g. contextual QA). This makes our method more applicable to different types of user queries that require more attention to the user-provided context.

## 5 EXPERIMENTS

We compare our method to vanilla instruction-tuning to demonstrate the efficacy of mitigating loss of context awareness. We instruction-finetuned 3 open-source, pretrained models on 3 different instruction datasets and evaluated the models on both context-dependency and general instruction-following tasks.

### 5.1 EXPERIMENTAL SETTINGS

#### 5.1.1 MODELS

We consider 3 open-source, pretrained large language models in our experiments: TinyLlama-1.1B (Zhang et al., 2024a), Llama-2-7B (Touvron et al., 2023), and Llama-3-8B (Dubey et al., 2024). TinyLlama-1.1B is a 1.1B Llama model pretrained on 3 trillion tokens with a context window length of 2048. Llama-2-7B and Llama-3-8B have a context window of 4096 and 8192 tokens, respectively. Due to limited computational resources, we only include Llama-2-7B and Llama-3-8B in our experiments and truncate the training examples up to 4096 tokens. We fine-tune Llama-2 and Llama-3 models using QLoRA (Dettmers et al., 2024) on transformer layers. Detailed hyperparameters can be found in Appendix A.1.1.

#### 5.1.2 INSTRUCTION FINE-TUNING DATASETS

We include 3 open-source instruction fine-tuning datasets: ShareGPT datasets adopted by Vicuna (Chiang et al., 2023), UltraChat-200k (Ding et al., 2023), and WizardLM-70K (Xu et al., 2023a). For ShareGPT, we follow the same preprocessing process as Chiang et al. (2023). We also removed refusal responses in ShareGPT and WizardLM-70K as the fine-tuned models will become oversensitive otherwise.

For all 3 datasets, we removed model responses without user input instructions in incomplete conversation chunks. Statistics of the processed datasets are presented in Table 1.

## 5.1.3 BENCHMARKS AND METRICS

SFT Dataset	Model Name	$\alpha$	NIH	SQuAD	QuAC	DROP
ShareGPT (Vicuna)	TinyLlama	1.0	0.9846	0.5918	0.1130	0.2739
		0.9	<b>1.0</b>	<b>0.5972</b>	<b>0.1140</b>	<b>0.2781</b>
		0.8	0.4661	0.3657	0.0810	0.2539
	Llama-2	1.0	0.3378	0.7601	0.1590	<b>0.3390</b>
		0.9	<b>0.4018</b>	<b>0.7920</b>	<b>0.1740</b>	0.3309
		0.8	0.2686	0.7699	0.1670	0.3353
	Llama-3	1.0	0.8957	0.8216	0.1560	<b>0.4215</b>
		0.9	<b>0.91</b>	<b>0.8307</b>	<b>0.1650</b>	0.4110
		0.8	0.5975	0.5363	0.0650	0.2057
UltraChat-200K	TinyLlama	1.0	<b>1.0</b>	<b>0.7475</b>	0.1570	<b>0.3096</b>
		0.9	<b>1.0</b>	0.7408	<b>0.1590</b>	0.3016
		0.8	0.8368	0.6243	0.1400	0.2801
	Llama-2	1.0	<b>0.9850</b>	0.8272	0.1540	<b>0.3791</b>
		0.9	0.9429	<b>0.8529</b>	<b>0.1700</b>	0.3774
		0.8	0.3303	0.7739	0.1180	0.2261
	Llama-3	1.0	<b>1.0</b>	0.8393	0.1610	<b>0.5099</b>
		0.9	<b>1.0</b>	<b>0.8604</b>	<b>0.1720</b>	0.4952
		0.8	0.9182	0.7095	0.1450	0.2801
WizardLM-70K	TinyLlama	1.0	0.9250	<b>0.5994</b>	0.0990	<b>0.2753</b>
		0.9	<b>1.0</b>	0.5878	<b>0.1040</b>	0.2697
		0.8	0.6132	0.4624	0.0510	0.2484
	Llama-2	1.0	0.7375	0.8229	0.1550	<b>0.3407</b>
		0.9	<b>0.8264</b>	<b>0.8430</b>	<b>0.1820</b>	0.3360
		0.8	0.5621	0.8144	0.1850	0.3326
	Llama-3	1.0	<b>0.9846</b>	0.8765	0.1560	<b>0.4687</b>
		0.9	0.9700	<b>0.8851</b>	<b>0.1650</b>	0.4203
		0.8	0.9625	0.7173	0.1050	0.3042

Table 2: NIH and reading comprehension performances with different attention intervention factors.

**Benchmarks on context-awareness** In addition to NIH, we additionally include 3 closed-book QA tasks to benchmark context awareness: SQuAD (Rajpurkar et al., 2016), QuAC (Choi et al., 2018) and DROP (Dua et al., 2019).

SQuAD is a reading comprehension benchmark where the answer to each question can be found in the context. We only evaluate the answerable subset of questions in SQuAD 1.0. QuAC is similar to SQuAD but the questions are more open-ended and the answers contain a longer span in the context. We only evaluate the model on the first conversation round of QuAC, since WizardLM-70K is a single-round conversational dataset. DROP requires a more comprehensive understanding and analysis of the given context and the answers require retrieval of multiple information from the context and discrete operations such as addition, sorting, or counting. As instruction fine-tuned models are not specifically trained on QA tasks to provide concise answers, the models’ responses are generally more verbose. Therefore, we report the containing score (whether the model response contains the ground-truth answer) instead of the F1 score to exclude the impacts of response styles of different models. Prompt templates for QA tasks are listed in Appendix A.3.

For Needle-in-a-haystack, we report the recall error as defined in Section 3.1. We set the maximum NIH context length to 1k for models fine-tuned on WizardLM-70K, due to the shorter instructions in the dataset. For models fine-tuned on ShareGPT and UltraChat-200k, we set the maximum NIH context length to the maximum context window considered in finetuning, which is 2k for TinyLlama and 4k for Llama-2 / Llama-3. The prompt template used in NIH is the same as Section 3.1 except that we remove the response prefix and keeps only the user input prompt.

**Benchmark on general instruction-following** We evaluate the general instruction-following performance after instruction fine-tuning on MT-Bench (Zheng et al., 2023). The quality ratings of the responses are judged by a GPT-4 judge regarding the helpfulness, relevance, accuracy, depth, creativity, and level of detail of the response. We report the average rating across 80 responses.

## 5.2 ATTENTION STEERING ON USER INPUT

We report the performance of direct attention steering in Table 2. According to the results, a medium intervention factor  $\alpha$  can boost the performance of NIH and most QA tasks except DROP. DROP is a more challenging closed-book QA dataset that requires additional operations or calculations based on information retrieved from the context. The worsening performance on DROP suggests that intervening attention scores, albeit improving context awareness, might hurt other capabilities of the model. For general instruction-following tasks that does not require heavy context-dependency, the user can choose to turn off the attention steering, and keep the same performance on general tasks.

## 5.3 INSTRUCTION FINE-TUNING WITH CONTEXT-DEPENDENCY INDICATORS

SFT dataset	Pretrained Model	Method	NIH	SQuAD	QuAC	DROP	Avg.	MT-Bench
ShareGPT (Vicuna)	TinyLlama-1.1B	Vanilla	0.9846	0.5918	0.1130	0.2739	0.4908	3.7250
		+ Indicator	<b>0.9921</b>	<b>0.6144</b>	<b>0.1290</b>	<b>0.2784</b>	<b>0.5035</b>	<b>3.8250</b>
	Llama-2-7b	Vanilla	0.3378	0.7601	0.1590	<b>0.3390</b>	0.3990	<b>6.4875</b>
		+ Indicator	<b>0.7007</b>	<b>0.7830</b>	<b>0.1600</b>	0.3390	<b>0.4957</b>	5.7375
	Llama-3-8b	Vanilla	0.8957	0.8216	0.1560	0.4215	0.5737	<b>7.4375</b>
		+ Indicator	<b>0.9404</b>	<b>0.8394</b>	<b>0.1660</b>	<b>0.4317</b>	<b>0.5943</b>	7.1875
UltraChat-200K	TinyLlama-1.1B	Vanilla	<b>1.0</b>	0.7289	0.1520	0.3096	0.5476	3.9000
		+ Indicator	<b>1.0</b>	<b>0.7475</b>	<b>0.1570</b>	<b>0.3096</b>	<b>0.5535</b>	<b>4.0250</b>
	Llama-2-7b	Vanilla	<b>0.9850</b>	0.8272	0.1540	<b>0.3791</b>	0.5863	5.7125
		+ Indicator	0.9725	<b>0.8508</b>	<b>0.1570</b>	0.3758	<b>0.5890</b>	<b>5.8250</b>
	Llama-3-8b	Vanilla	<b>1.0</b>	0.8393	0.1610	<b>0.5099</b>	0.6276	7.2375
		+ Indicator	<b>1.0</b>	<b>0.8510</b>	<b>0.1660</b>	0.5022	<b>0.6298</b>	<b>7.2750</b>
WizardLM-70K	TinyLlama-1.1B	Vanilla	0.9250	0.5994	0.0990	0.2753	0.4747	4.2750
		+ Indicator	<b>0.9925</b>	<b>0.6279</b>	<b>0.1140</b>	<b>0.2836</b>	<b>0.5045</b>	<b>4.5625</b>
	Llama-2-7b	Vanilla	0.7375	0.8229	0.1550	0.3407	0.5140	5.7750
		+ Indicator	<b>0.9254</b>	<b>0.8260</b>	<b>0.1640</b>	<b>0.3444</b>	<b>0.5650</b>	<b>5.9875</b>
	Llama-3-8b	Vanilla	0.9846	0.8765	0.1560	0.4687	0.6215	7.1125
		+ Indicator	<b>0.9871</b>	<b>0.8792</b>	<b>0.1710</b>	<b>0.4785</b>	<b>0.6290</b>	<b>7.2750</b>
(Internal datasets)	Llama-2-7b-chat	-	0.8264*	0.8301	0.1330	0.4422	0.5579	6.9375
	Llama-3-8b-Instruct	-	1.0*	0.8612	0.1850	0.4654	0.6279	8.0750

\* Here NIH is evaluated without the response prefix used in Section 3.1 and the maximum context length is set to 4096 for fair comparison. Therefore, the numbers are different from the numbers in Figure 2.

Table 3: Comparing vanilla instruction finetuning with finetuning with context-relevant indicators (+ indicator). On ‘+ Indicator’ models, [IND] is added for NIH and contextual QA tasks and removed for MT-Bench. As a reference, we also list the performances evaluated on official Llama-2 and Llama-3 instruct models, which are finetuned with internal datasets.

**Settings and selected subset  $\hat{S}$ .** We adopt a TinyLlama model finetuned on the original ShareGPT (Vicuna) dataset as the seed model  $M$ . We compute the context-dependency score on an arbitrary middle layer (15 in all of our main experiments) for faster computation. As we show in Appendix B.3, subsets selected by scores calculated on different middle layers are highly consistent. We compute the context-dependency scores on all three instruction fine-tuning datasets and append [IND] to prompts associated with responses with context-dependency scores above the threshold  $s_M(\mathbf{Y}_m) > \beta$ . We set  $\beta = 0.6$  in all our main experiments. Ablation studies on the threshold  $\beta$  can be found in Appendix B.4. We report the number of selected instructions of each fine-tuning dataset in Table 1. More statistics can be found in Appendix B.5.

**Evaluation results on context-awareness.** We report the performance of NIH and three QA tasks in Table 3. For the “vanilla” fine-tuning setting, we train and evaluate the model without the in-

486 indicator token. For the “+ Indicator” setting, we add the indicator token to the selected subset of  
 487 prompts in fine-tuning and all queries for evaluation. According to the table, “+ Indicator” outper-  
 488 forms “vanilla” fine-tuning in most cases. Particularly, “+ Indicator” consistently outperforms direct  
 489 attention steering on DROP, reinforcing that the training-time solution handles trade-offs between  
 490 language model capabilities better than inference-time editing. The results demonstrate that mod-  
 491 els can learn to focus more on the user-provided context when the indicator token is present in the  
 492 prompt.

493  
 494 **Evaluation on general instruction-following benchmarks.** We also report the evaluation results  
 495 on MT-Bench in Table 3. As samples in MT-Bench are mostly open-ended questions without heavy  
 496 context dependency, the context-dependency indicator is not added for the “+ Indicator” models  
 497 during evaluation. Models finetuned with our method show comparable or even better performance  
 498 on MT-Bench in most cases. Therefore, our methods can effectively mitigate the loss of context  
 499 awareness without losing the general ability of instruction following. The context indicator enables  
 500 more fine-grained control over language models’ behaviors during inference time.

## 501 6 CONCLUSION

502  
 503 This work highlights the detrimental effects of supervised instruction fine-tuning on the context-  
 504 awareness of pretrained language models, even in scenarios involving short context lengths. We  
 505 have identified that this decline is closely linked to inherent attention allocation biases within chat  
 506 templates. To combat these challenges, we propose an inference-time technique that allows for  
 507 the manual adjustment of attention scores, facilitating a partial recovery of context awareness in  
 508 instruction-tuned models. Furthermore, we introduce a training-time approach that employs condi-  
 509 tional indicators to help preserve context awareness during the instruction-tuning process. Together,  
 510 these contributions aim to enhance the performance of language models in maintaining contextual  
 511 understanding while benefiting from supervised instruction.

512  
 513 **Limitation** Our technique of associating context-dependent user instructions with the indicator  
 514 token may also encode other unintended styles of the selected subset of instructions. The issue  
 515 is particularly aggravated if the subset of context-relevant samples is small and significantly dif-  
 516 ferent from the remaining dataset. Utilizing the indicator token also requires the users to know  
 517 which instructions require paying more attention to the user context during inference time. Tasks  
 518 demonstrated in our experiment (e.g. keyword retrieval and closed-book QA) evidently require more  
 519 attention to the user context. However, complex tasks may involve multiple language skills which  
 520 complicates whether the addition of the indicator token will benefit the performance. A future direc-  
 521 tion is to automatically determine whether appending the indicator to a user instruction is necessary.  
 522 Besides, we only validate the effectiveness of our proposed methods on relatively small-size models  
 523 due to limited computational resources.

## 524 REFERENCES

- 525  
 526 Rishabh Agarwal, Avi Singh, Lei M. Zhang, Bernd Bohnet, Luis Rosias, Stephanie Chan, Biao  
 527 Zhang, Ankesh Anand, Zaheer Abbas, Azade Nova, John D. Co-Reyes, Eric Chu, Feryal Be-  
 528 hbahani, Aleksandra Faust, and Hugo Larochelle. Many-shot in-context learning, 2024. URL  
 529 <https://arxiv.org/abs/2404.11018>.
- 530 Shengnan An, Zexiong Ma, Zeqi Lin, Nanning Zheng, and Jian-Guang Lou. Make your llm fully  
 531 utilize the context, 2024. URL <https://arxiv.org/abs/2404.16811>.
- 532  
 533 Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn  
 534 Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless  
 535 assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*,  
 536 2022.
- 537 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhari-  
 538 wal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agar-  
 539 wal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh,  
 Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz

- 540 Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec  
541 Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In  
542 H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural  
543 Information Processing Systems*, volume 33, pp. 1877–1901. Curran Associates, Inc.,  
544 2020. URL [https://proceedings.neurips.cc/paper\\_files/paper/2020/  
545 file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf).
- 546 Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng,  
547 Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An  
548 open-source chatbot impressing gpt-4 with 90%\* chatgpt quality, March 2023. URL [https://  
549 lmsys.org/blog/2023-03-30-vicuna/](https://lmsys.org/blog/2023-03-30-vicuna/).
- 550 Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wen tau Yih, Yejin Choi, Percy Liang, and Luke  
551 Zettlemoyer. Quac : Question answering in context, 2018. URL [https://arxiv.org/abs/  
552 1808.07036](https://arxiv.org/abs/1808.07036).
- 553 Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep  
554 reinforcement learning from human preferences. *Advances in neural information processing systems*,  
555 30, 2017.
- 556 Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning  
557 of quantized llms. *Advances in Neural Information Processing Systems*, 36, 2024.
- 558 Ning Ding, Yulin Chen, Bokai Xu, Yujia Qin, Zhi Zheng, Shengding Hu, Zhiyuan Liu, Maosong  
559 Sun, and Bowen Zhou. Enhancing chat language models by scaling high-quality instructional  
560 conversations. *arXiv preprint arXiv:2305.14233*, 2023.
- 561 Yi Dong, Zhilin Wang, Makesh Narsimhan Sreedhar, Xianchao Wu, and Oleksii Kuchaiev.  
562 Steerlm: Attribute conditioned sft as an (user-steerable) alternative to rlhf. *arXiv preprint  
563 arXiv:2310.05344*, 2023.
- 564 Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner.  
565 DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs. In  
566 Jill Burstein, Christy Doran, and Thamar Solorio (eds.), *Proceedings of the 2019 Conference of  
567 the North American Chapter of the Association for Computational Linguistics: Human Language  
568 Technologies, Volume 1 (Long and Short Papers)*, pp. 2368–2378, Minneapolis, Minnesota, June  
569 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1246. URL [https://  
570 aclanthology.org/N19-1246](https://aclanthology.org/N19-1246).
- 571 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha  
572 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony  
573 Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark,  
574 Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere,  
575 Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris  
576 Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong,  
577 Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny  
578 Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino,  
579 Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael  
580 Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Ander-  
581 son, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah  
582 Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan  
583 Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Ma-  
584 hadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy  
585 Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak,  
586 Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Al-  
587 wala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini,  
588 Krithika Iyer, Kshitiz Malik, Kuenlev Chiu, Kunal Bhalla, Lauren Rantala-Yeary, Laurens van der  
589 Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo,  
590 Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Man-  
591 nat Singh, Manohar Paluri, Marcin Kardas, Mathew Oldham, Mathieu Rita, Maya Pavlova,  
592 Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal,  
593

594 Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Olivier Duchenne, Onur  
595 Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhar-  
596 gava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong,  
597 Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic,  
598 Roberta Raileanu, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sum-  
599 baly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa,  
600 Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang,  
601 Sharath Rapparthi, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende,  
602 Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney  
603 Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom,  
604 Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta,  
605 Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vladan Petro-  
606 vic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang,  
607 Xiaoqing Ellen Tan, Xinfeng Xie, Xuchao Jia, Xuwei Wang, Yaelle Goldschlag, Yashesh Gaur,  
608 Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre  
609 Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aaron Grattafiori, Abha  
610 Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay  
611 Menon, Ajay Sharma, Alex Boesenberg, Alex Vaughan, Alexei Baevski, Allie Feinstein, Amanda  
612 Kallet, Amit Sangani, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew  
613 Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Franco, Aparajita  
614 Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh  
615 Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De  
616 Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Bran-  
617 don Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina  
618 Mejia, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai,  
619 Chris Tindal, Christoph Feichtenhofer, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li,  
620 Danny Wyatt, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana  
621 Liskovich, Didem Foss, DingKang Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil,  
622 Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Erik Brinkman, Esteban Ar-  
623 caute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Firat Ozgenel, Francesco  
624 Caggioni, Francisco Guzmán, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella  
625 Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Govind Thattai, Grant Herman, Grigory  
626 Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hamid Shojanazeri, Han Zou, Hannah Wang,  
627 Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Gold-  
628 man, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Irina-Elena Veliche, Itai Gat, Jake Weissman,  
629 James Geboski, James Kohli, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer  
630 Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe  
631 Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie  
632 Wang, Kai Wu, Kam Hou U, Karan Saxena, Karthik Prasad, Kartikay Khandelwal, Katayoun  
633 Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kun Huang, Kunal  
634 Chawla, Kushal Lakhota, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva,  
635 Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian  
636 Khabza, Manav Avalani, Manish Bhatt, Maria Tsimpoukelli, Martynas Mankus, Matan Hasson,  
637 Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Ke-  
638 neally, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel  
639 Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mo-  
640 hammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navy-  
641 ata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikolay Pavlovich Laptev, Ning Dong,  
642 Ning Zhang, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli,  
643 Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux,  
644 Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao,  
645 Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Raymond Li,  
646 Rebeekah Hogan, Robin Battey, Rocky Wang, Rohan Maheswari, Russ Howes, Ruty Rinott,  
647 Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Sa-  
648 tadrū Pan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lind-  
649 say, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shiva Shankar, Shuqiang Zhang, Shuqiang  
650 Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen  
651 Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Sungmin Cho,  
652 Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser,

- 648 Tamara Best, Thilo Kohler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Tim-  
649 othy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan,  
650 Vinay Satish Kumar, Vishal Mangla, Vitor Albiero, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu  
651 Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Con-  
652 stable, Xiaocheng Tang, Xiaofang Wang, Xiaojian Wu, Xiaolan Wang, Xide Xia, Xilun Wu,  
653 Xinbo Gao, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi,  
654 Youngjin Nam, Yu, Wang, Yuchen Hao, Yundi Qian, Yuzi He, Zach Rait, Zachary DeVito, Zef  
655 Rosnbrick, Zhaoduo Wen, Zhenyu Yang, and Zhiwei Zhao. The llama 3 herd of models, 2024.  
656 URL <https://arxiv.org/abs/2407.21783>.
- 657 Tingchen Fu, Deng Cai, Lema Liu, Shuming Shi, and Rui Yan. Disperse-then-merge: Pushing  
658 the limits of instruction tuning via alignment tax reduction. In Lun-Wei Ku, Andre Martins, and  
659 Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics ACL 2024*, pp.  
660 2967–2985, Bangkok, Thailand and virtual meeting, August 2024. Association for Computational  
661 Linguistics. doi: 10.18653/v1/2024.findings-acl.175. URL [https://aclanthology.org/  
662 2024.findings-acl.175](https://aclanthology.org/2024.findings-acl.175).
- 663 Cheng-Yu Hsieh, Yung-Sung Chuang, Chun-Liang Li, Zifeng Wang, Long Le, Abhishek Kumar,  
664 James Glass, Alexander Ratner, Chen-Yu Lee, Ranjay Krishna, and Tomas Pfister. Found in  
665 the middle: Calibrating positional attention bias improves long context utilization. In Lun-  
666 Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association for Computa-  
667 tional Linguistics ACL 2024*, pp. 14982–14995, Bangkok, Thailand and virtual meeting, August  
668 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.890. URL  
669 <https://aclanthology.org/2024.findings-acl.890>.
- 670 Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane  
671 Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. Atlas: Few-shot learning  
672 with retrieval augmented language models. *Journal of Machine Learning Research*, 24(251):  
673 1–43, 2023.
- 674 Greg Kamradt. Llmtest\_needleinahaystack, November 2023. URL [https://github.com/  
675 gkamradt/LLMTest\\_NeedleInAHaystack](https://github.com/gkamradt/LLMTest_NeedleInAHaystack).
- 676 Urvashi Khandelwal, Angela Fan, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. Nearest neigh-  
677 bor machine translation. *arXiv preprint arXiv:2010.00710*, 2020.
- 678 Tomasz Korbak, Kejian Shi, Angelica Chen, Rasika Vinayak Bhalerao, Christopher Buckley, Jason  
679 Phang, Samuel R Bowman, and Ethan Perez. Pretraining language models with human prefer-  
680 ences. In *International Conference on Machine Learning*, pp. 17506–17533. PMLR, 2023.
- 681 Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal,  
682 Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented genera-  
683 tion for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:  
684 9459–9474, 2020.
- 685 Yong Lin, Hangyu Lin, Wei Xiong, Shizhe Diao, Jianmeng Liu, Jipeng Zhang, Rui Pan, Haoxiang  
686 Wang, Wenbin Hu, Hanning Zhang, Hanze Dong, Renjie Pi, Han Zhao, Nan Jiang, Heng Ji, Yuan  
687 Yao, and Tong Zhang. Mitigating the alignment tax of rlhf, 2024. URL [https://arxiv.  
688 org/abs/2309.06256](https://arxiv.org/abs/2309.06256).
- 689 Hao Liu, Carmelo Sferrazza, and Pieter Abbeel. Chain of hindsight aligns language models with  
690 feedback. *arXiv preprint arXiv:2302.02676*, 2023.
- 691 Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V  
692 Le, Barret Zoph, Jason Wei, et al. The flan collection: Designing data and methods for effective  
693 instruction tuning. In *International Conference on Machine Learning*, pp. 22631–22648. PMLR,  
694 2023.
- 695 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong  
696 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to fol-  
697 low instructions with human feedback. *Advances in neural information processing systems*, 35:  
698 27730–27744, 2022.

- 702 Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions  
703 for machine comprehension of text, 2016. URL <https://arxiv.org/abs/1606.05250>.  
704
- 705 Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Eric Hambro,  
706 Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can  
707 teach themselves to use tools. In *Thirty-seventh Conference on Neural Information Processing*  
708 *Systems*, 2023. URL <https://openreview.net/forum?id=Yacmpz84TH>.
- 709 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Niko-  
710 lay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open founda-  
711 tion and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- 712 Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du,  
713 Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. *arXiv preprint*  
714 *arXiv:2109.01652*, 2021.
- 715  
716 Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and  
717 Daxin Jiang. Wizardlm: Empowering large language models to follow complex instructions.  
718 *arXiv preprint arXiv:2304.12244*, 2023a.
- 719 Peng Xu, Wei Ping, Xianchao Wu, Lawrence McAfee, Chen Zhu, Zihan Liu, Sandeep Subramanian,  
720 Evelina Bakhturina, Mohammad Shoeybi, and Bryan Catanzaro. Retrieval meets long context  
721 large language models. *arXiv preprint arXiv:2310.03025*, 2023b.
- 722  
723 Peiyuan Zhang, Guangtao Zeng, Tianduo Wang, and Wei Lu. Tinyllama: An open-source small  
724 language model, 2024a.
- 725 Qingru Zhang, Chandan Singh, Liyuan Liu, Xiaodong Liu, Bin Yu, Jianfeng Gao, and Tuo Zhao.  
726 Tell your model where to attend: Post-hoc attention steering for LLMs. In *The Twelfth Interna-*  
727 *tional Conference on Learning Representations*, 2024b. URL [https://openreview.net/](https://openreview.net/forum?id=xZDW00oejD)  
728 [forum?id=xZDW00oejD](https://openreview.net/forum?id=xZDW00oejD).
- 729  
730 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang,  
731 Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica.  
732 Judging llm-as-a-judge with mt-bench and chatbot arena, 2023. URL [https://arxiv.org/](https://arxiv.org/abs/2306.05685)  
733 [abs/2306.05685](https://arxiv.org/abs/2306.05685).  
734  
735  
736  
737  
738  
739  
740  
741  
742  
743  
744  
745  
746  
747  
748  
749  
750  
751  
752  
753  
754  
755

## A APPENDIX

### A.1 EXPERIMENTAL DETAILS

#### A.1.1 INSTRUCTION FINE-TUNING

We adopted the fine-tuning recipes from the Huggingface alignment-handbook<sup>1</sup> for Llama-2 and Llama-3 QLoRA tuning. For the TinyLlama model, we use the fine-tuning recipe provided by the author<sup>2</sup>. We finetune the models for 1 epoch on ShareGPT and UltraChat-200K and 2 epochs on WizardLM-70K as it has a smaller training set. We use the TinyLlama chat template for all instruct models finetuned in Table 2 and 3.

### A.2 NIH EVALUATION DETAILS

For all NIH evaluations, we average the recall error on 400 tests. Specifically, we evaluate on 20 context lengths uniformly distributed between 200 and the maximum context length, and 20 needle insertion depths uniformly located within 0% and 100%.

### A.3 CONTEXTUAL QA EVALUATION DETAILS

We list the prompts used in contextual QA tasks in Table 4 and Table 5. For contextual QA tasks, we generate answers up to 100 tokens and truncate them at the end of the first complete sentence. For NIH tests, we generate the answers up to 50 tokens.

As UltraChat-200K constructs their data with a fixed set of prompt templates similar to our default ones used in evaluation (The templates used for ShareGPT and WizardLM models in Table 5 and 4), we evaluate UltraChat-200K finetuned models with a simpler template to exclude the impact from overfitting on finetuning prompt templates.

Instruct Finetuning Dataset	Template for SQuAD and DROP
ShareGPT & WizardLM-70K	{context}\nAnswer the question according to the above passage: {question}
UltraChat-200K	{context} {question}

Table 4: Prompt templates used for SQuAD and DROP in Table 2 and Table 3 when the model is finetuned on different instruction finetuning datasets.

Instruct Finetuning Dataset	Template for QuAC
ShareGPT & WizardLM-70K	{context}\nAnswer the question with pieces from the the above passage: {question}
UltraChat-200K	{context} {question}

Table 5: Prompt templates used for QuAC in Table 2 and Table 3 when the model is finetuned on different instruction finetuning datasets.

## B ADDITIONAL EXPERIMENT RESULTS

### B.1 FULL NIH RESULTS ON OPEN-SOURCE OFFICIAL MODELS

In Figure 2, we only report the NIH performances when the response prefix is added for fair comparison. In Table 6 we show the exact numbers for Figure 2 as well as additional evaluation results without the response prefix. When the response prefix is removed, the performance drop on NIH is even more significant compared to without chat templates.

<sup>1</sup><https://github.com/huggingface/alignment-handbook>

<sup>2</sup><https://github.com/jzhang38/TinyLlama>

Model Name	Context window	w/o chat template		
		w/ response prefix	w/ response prefix	w/o response prefix
Llama-2-7b	4096	98.67%	-	-
Llama-2-7b-chat		99.75%	97.75%	82.71%
Llama-2-13b	4096	99.35%	-	-
Llama-2-13b-chat		95.78%	90.79%	92.71%
Llama-3-8b	8192	99.50%	-	-
Llama-3-8b-instruct		100%	95.25%	95.35%
mistral-v0.2	32768	100%	-	-
mistral-v0.2-instruct		99.00%	94.14%	93.92%
mistral-v0.3	32768	100%	-	-
mistral-v0.3-instruct		99.32%	84.71%	72.00%
gemma-2-9b	8196	100%	-	-
gemma-2-9b-it		98.75%	98.03%	98.25%
gemma-2-27b	8196	100%	-	-
gemma-2-27b-it		100%	99.64%	99.25%

Table 6: NIH performance with and without chat templates on different models.

### B.2 FULL RESULTS FOR FIGURE 2

In Figure 3, we only show the changes of attention allocation with and without chat templates. In Figure 4 we show the absolute numbers of attention allocation to each part of input prompts. When the chat template is added, we normalize the attention weight on user tokens, response tokens and BOS token only, with a sum of attention allocation to be 1.

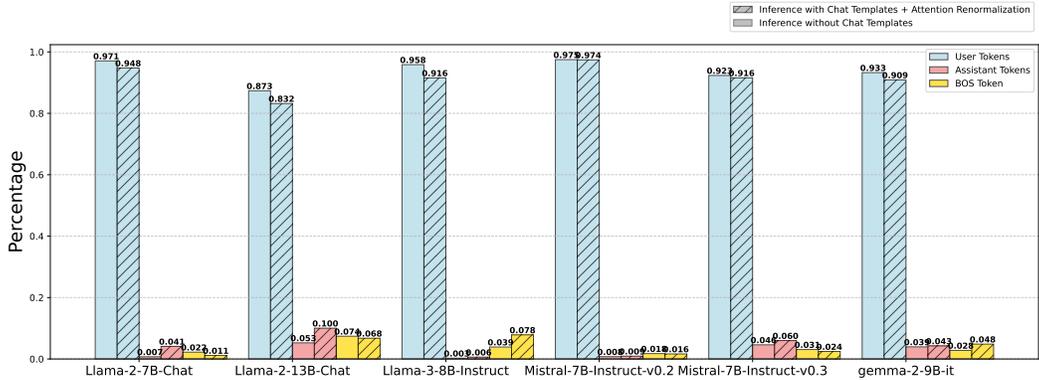


Figure 4: We visualize the full attention allocation on user tokens, assistant tokens and BOS token with and without applying the chat templates. The attention allocation is calculated when the model is generating the first answer token in its response. For the case where the chat template is applied, we normalize the attention value on user tokens, assistant tokens, and the BOS token such that attention scores allocated to the three sums up to 1. The attention weight is averaged on 400 tests with context lengths ranging from 200 to 4000 and needle depth from 0% to 100%.

### B.3 AGREEMENT BETWEEN DIFFERENT LAYERS

In Figure 5, we calculate and visualize the disagreement heatmap in  $\hat{S}$  selection when the context-dependency score is calculated on different layers. We use the same TinyLlama model fine-tuned on the vanilla ShareGPT dataset as the seed model  $M$ . Specifically, we first calculate the context-dependency scores for each conversation turn in 500 randomly sampled examples from the ShareGPT dataset on different layers. We then select the top 10% conversation turn with the highest context-dependency scores on the layer as the subset  $\hat{S}$ . We compute the disagreement between two

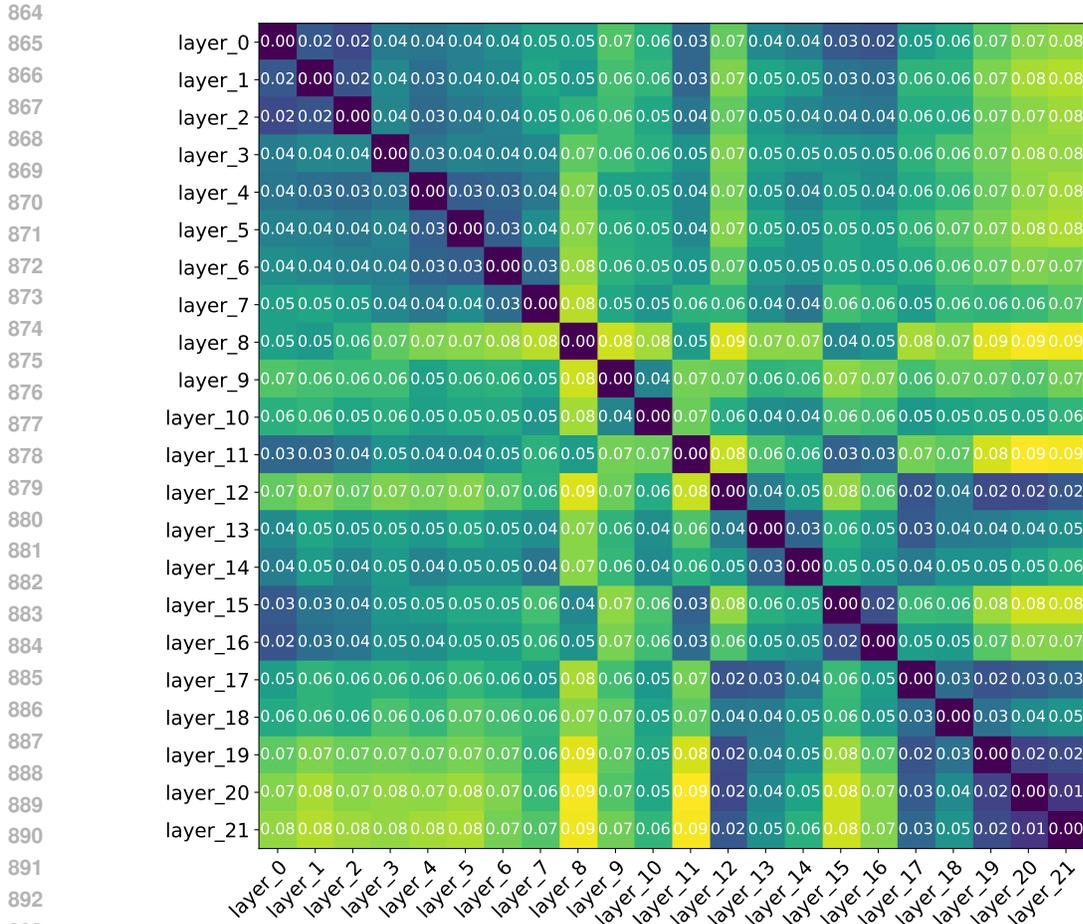


Figure 5: We visualize the disagreement heatmap of  $\hat{S}$  selection when the context-dependency score  $S_M(\mathbf{Y}_m)$  is calculated on different layers. We select 10% of conversation turns with the highest context-dependency scores on each layer as  $\hat{S}$ . The disagreement is measured by the number of non-overlapped conversation turns in  $\hat{S}$  selected by any two layers.

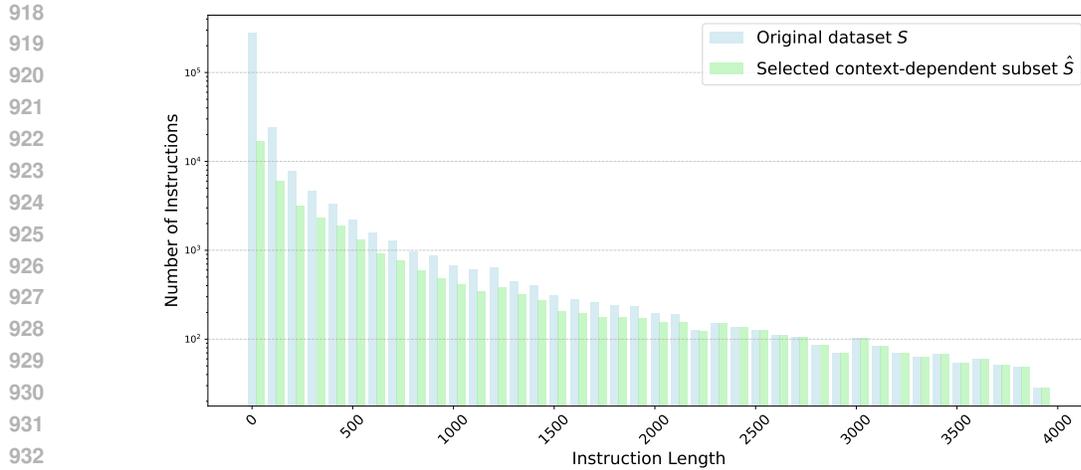
layers by calculating the ratio of non-overlapped conversation turns in their selected  $\hat{S}$ . We can see from the figure that the disagreement between 9 middle layers are low, indicating that we can safely choose an arbitrary layer for the context-dependency score calculation.

#### B.4 ABLATION STUDY FOR DIFFERENT THRESHOLD $\beta$

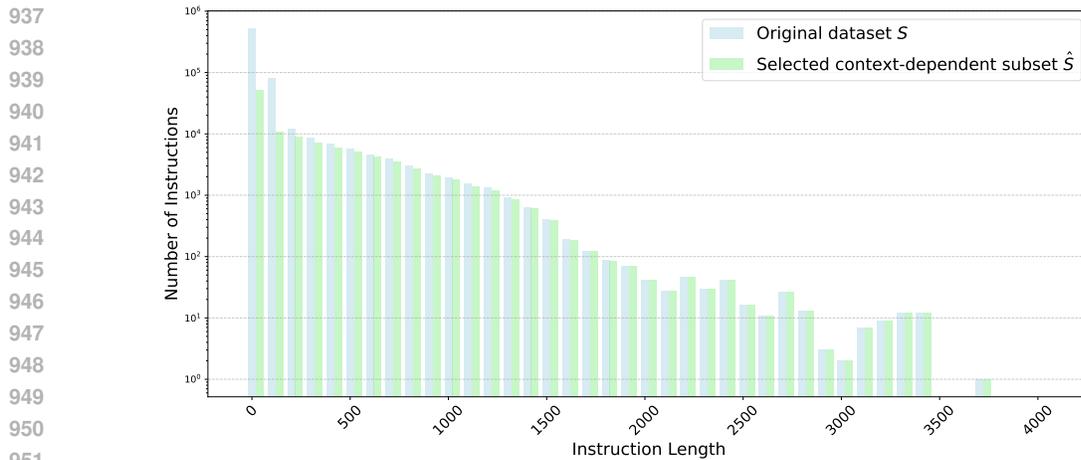
Threshold $\beta$	SQuAD	QuAC	DROP	MT-Bench
1.0 (Vanilla)	0.5918	0.1130	0.2739	3.725
0.5	0.6207	0.1270	0.2872	4.075
0.6	0.6144	0.1290	0.2784	3.825
0.7	0.6160	0.1290	0.2786	3.675

Table 7: Ablation study with different threshold  $\beta$ , which is used in Section 4.2.

We use  $\beta = 0.6$  in all our main experiments. To evaluate the sensitivity to the threshold  $\beta$ , we select  $\hat{S}$  with different thresholds and prepare the final modified instruction finetuning dataset. We finetune a TinyLlama-1.1B model on these three datasets and evaluate them on three contextual QA



934 Figure 6: Change of instruction lengths between the original and the selected subset from ShareGPT  
935 dataset.



952 Figure 7: Change of instruction lengths between the original and the selected subset from UltraChat-  
953 200K dataset.

954  
955  
956 tasks and MT-Bench. As we can see from Table 7, all three models outperforms vanilla finetuning  
957 on the contextual QA tasks. However, performance on MT-Bench shows a decreasing trend when  
958 the threshold increases from 0.5 to 0.7, potentially due to a more drastic difference between  $\hat{S}$  and  
959 the unselected subset.

## 961 B.5 DISTRIBUTION OF INSTRUCTION LENGTHS

962  
963 Here we visualize the change of distribution of instruction length between original instruction fine-  
964 tuning dataset and the selected context-dependent subset  $\hat{S}$ . Although a higher context-dependency  
965 is to some extent correlated with longer instruction length, there are still a large amount of short  
966 instructions showing high context dependency and selected into  $\hat{S}$ .

967  
968  
969  
970  
971

972  
 973  
 974  
 975  
 976  
 977  
 978  
 979  
 980  
 981  
 982  
 983  
 984  
 985  
 986  
 987  
 988  
 989  
 990  
 991  
 992  
 993  
 994  
 995  
 996  
 997  
 998  
 999  
 1000  
 1001  
 1002  
 1003  
 1004  
 1005  
 1006  
 1007  
 1008  
 1009  
 1010  
 1011  
 1012  
 1013  
 1014  
 1015  
 1016  
 1017  
 1018  
 1019  
 1020  
 1021  
 1022  
 1023  
 1024  
 1025

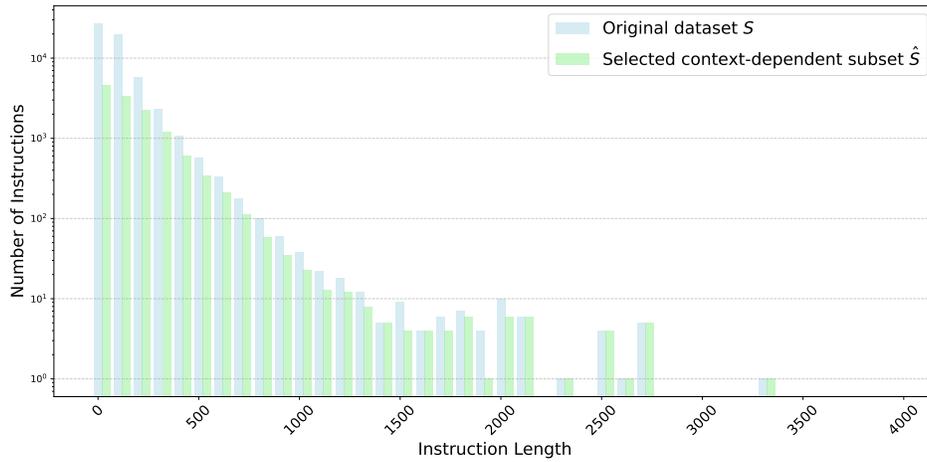


Figure 8: Change of instruction lengths between the original and the selected subset from WizardLM-70K dataset.