

# Better supervised fine-tuning of closed-source large models

Fei Ding<sup>1</sup>, Correspondence: [dingfei@email.ncu.edu.cn](mailto:dingfei@email.ncu.edu.cn)

## Abstract

The recent proliferation of so-called open-source large language models (such as LLaMA, Falcon, Mistral) has introduced a broader range of alternatives for AI practitioners and researchers. However, the majority of these models cannot be considered truly open-source, as they often provide only partial artifacts, such as final model weights or inference code. Furthermore, technical documentation accompanying these models tends to focus on high-level architectural decisions and superficial metrics, leaving critical aspects of the training process, including dataset composition, distribution, model checkpoints, and intermediate results, largely undisclosed. This lack of transparency presents a significant barrier to progress in the field, restricting the potential for open, collaborative research. In the absence of access to original datasets, attempts to further train or fine-tune these models by third parties are susceptible to issues such as catastrophic forgetting. In response to this challenge, we propose a method that facilitates more effective supervised fine-tuning of these closed-source models, without requiring access to the original data, while mitigating the risk of catastrophic forgetting.

## 1 Introduction

Catastrophic forgetting represents a critical challenge for large language models (LLMs) and neural networks (NNs). This phenomenon is characterized by the models' propensity to abruptly lose previously acquired knowledge when assimilating new information. Such a limitation significantly impedes the development of robust and reliable artificial intelligence systems, particularly in dynamic contexts where ongoing learning from novel data is imperative.

Catastrophic forgetting—the tendency of deep neural networks to "forget" previously acquired knowledge when introduced to new information—has been a subject of investigation since 1989

McCloskey and Cohen, 1989. This phenomenon is most evident when models are sequentially trained on distinct tasks; however, it also occurs whenever a model learns information in a sequential manner, particularly when there are shifts in data distribution over time. In practical machine learning applications, it is common for new training data to be introduced continuously. To incorporate this new information into model training, developers face a choice: they can either retrain the entire model from scratch, starting with randomly initialized weights and utilizing all available training data, a process that is computationally intensive, or they can take an existing model trained on prior data and perform fine-tuning on the newly acquired data. However, since new data typically originates from a distribution that is slightly different from that of the old data, significant changes in distribution can exacerbate the effects of catastrophic forgetting during the fine-tuning process.

The landscape of Large Language Models (LLMs) has undergone a remarkable transformation over the past year, characterized by an unprecedented surge in both their popularity and capabilities. Leading this evolution are proprietary LLMs such as GPT-4 [OpenAI, 2023](#) and Claude [Claude, 2023](#), which have garnered significant attention within the AI community owing to their exceptional power and versatility. Concurrently, the recent emergence of openly accessible yet highly capable LLMs, including LLaMA ([Touvron et al., 2023a,b](#)), Falcon ([Penedo et al., 2023](#)), and Mistral ([Jiang et al., 2023](#)), has empowered researchers and practitioners to easily acquire, customize, and deploy LLMs across a broader range of environments and applications.

Catastrophic forgetting and overtraining (or overfitting) represent distinct challenges encountered in the training of neural networks and large language models. Catastrophic forgetting occurs when a model discards previously acquired knowledge

085	upon assimilating new information, particularly in	<b>2 Background and Related Work</b>	134
086	sequential learning contexts. This phenomenon		
087	is attributed to the modifications in model weights	<b>2.1 Data Rehearsal</b>	135
088	that disrupt the performance of earlier tasks. In con-	Robins (ROBINS, 1995) introduced the concept of	136
089	trast, overtraining arises when a model becomes	rehearsal in 1995, shortly following the advent of	137
090	excessively attuned to the training data, leading	the notion of catastrophic forgetting. In essence,	138
091	it to capture noise and specific details rather than	this approach entails incorporating data from pre-	139
092	generalizable patterns, ultimately resulting in poor	vious tasks during the training of new ones. While	140
093	performance on new, unseen data. While catast-	this method has demonstrated considerable efficacy,	141
094	rophic forgetting undermines knowledge retention	it necessitates maintaining access to historical data,	142
095	in dynamic learning environments, overfitting sig-	or at the very least, an independent and identically	143
096	nificantly restricts the model’s ability to generalize	distributed (i.i.d.) subsample of such data, which	144
097	effectively from the training set to novel data.	may not always be feasible. Furthermore, inte-	145
098	When continuing training, in order to address	grating past data increases the overall volume of	146
099	both catastrophic forgetting and overfitting, it re-	training data, resulting in longer training durations	147
100	quires us to have knowledge of both the original	for each epoch during model fine-tuning.	148
101	data and its distribution.	Since most large models do not have publicly	149
102	Despite the increasing prominence and acces-	available datasets for rehearsal, the common ap-	150
103	sibility of open-source large language models	proach is to use some public sft datasets mixed	151
104	(LLMs), a significant trend has emerged towards	with their own sft datasets to simulate a review	152
105	restricting visibility and access to the intricacies of	process. However, this approach can lead to cer-	153
106	their training, fine-tuning, and evaluation method-	tain issues. Our approach involves extracting the	154
107	ologies. This includes critical components such as	concealed data distribution of the supervised fine-	155
108	the underlying training code and datasets, which	tuning (SFT) instructions directly from the model	156
109	are essential for a comprehensive understanding of	parameters.	157
110	model behavior and performance.	<b>2.2 Continue Fine-tuning</b>	158
111	This approach limits our ability to perform SFT	Our methodology addresses the challenge of con-	159
112	(Supervised Fine-Tuning) on these models.	tinual fine-tuning, wherein the model undergoes	160
113	Because using the same data easily leads to over-	successive fine-tuning with newly acquired data	161
114	fitting, while differences in data distribution can	post-initial fine-tuning. Continual learning is es-	162
115	cause catastrophic forgetting, better SFT (Super-	sential for models that must adapt to dynamic en-	163
116	vised Fine-Tuning) requires an alternative approach	vironments, assimilating information from a continu-	164
117	for models that do not disclose their original SFT	ous data stream while retaining previously learned	165
118	data. We can reverse-engineer the model paramet-	knowledge. A critical obstacle in this domain is	166
119	ers to extract the distribution of the original SFT	the issue of catastrophic forgetting, which refers	167
120	data, then generate new SFT data based on this	to the pronounced degradation in performance on	168
121	distribution, and mix it with our own SFT data in a	earlier tasks when the model is exposed to novel	169
122	certain proportion. This allows for more effective	data. As the model adjusts its parameters to in-	170
123	fine-tuning.	corporate new information, it inadvertently over-	171
124	This paper presents the following contributions:	writes previously acquired knowledge, thereby di-	172
125		minishing its effectiveness on prior tasks. To ad-	173
126	• We deciphered the hidden data distribution of	dress this, the research community has proposed	174
127	open-source models through model paramet-	a range of strategies, typically classified into four	175
128	ers and used it for experience replay during	main categories: Replay-Based (Shin et al., 2017;	176
129	SFT fine-tuning to better mitigate catastrophic	Ren et al., 2024), Regularization-Based (Mi et al.,	177
	forgetting.	2020), Gradient-Based (Lee et al., 2021), and	178
130		Architecture-Based (Geng et al., 2021) approaches.	179
131	• We obtained the optimal instruction responses	In our experiments, we adopted a basic experience	180
132	through mutual scoring among three models,	replay mechanism, reduced the initial learning rate	181
133	significantly improving the response quality	to avoid overfitting.	182
	and enhancing the effectiveness of SFT.		

183  
184  
185  
186  
187  
188  
189  
190  
191  
  
192  
193  
194  
195  
196  
197  
198  
199  
200  
201  
202  
203  
204  
205  
206  
207  
208  
209  
210  
211  
212  
213  
214  
215  
216  
217  
218  
219  
220  
221  
222  
223  
224  
225  
226  
227  
228  
229  
230  
231  
232

### 3 Methods

Our experiments are divided into three parts: the first part involves extracting the original SFT data distribution from the model; the second part mixes the extracted SFT data with new data for training; and the third part uses commonly available general SFT data mixed with new data for training, comparing the results with those from the second part.

#### 3.1 Extracting the instruction distribution.

Cracking the instruction distribution consists of three steps: (1) instruction generation, (2) response generation, and (3) filtering high-quality responses. The pipeline can be fully automated without any human intervention.

##### Step 1: Instruction Generation.

The objective of this step is to extract unreleased training data from the model’s parameters. Given an open-weight aligned large language model (e.g., Llama-3-70B-Instruct), we design a pre-query template in the format of the predefined instruction template.

We input the prompt "`<start_header_id>user<end_header_id>`" into the large model (Llama-3-70B-Instruct), which generates a single instruction in response. By repeating this process 100,000 times, we obtain a total of 100,000 instructions, which collectively represent the current instruction distribution of the large model.

**Step 2: Response Generation.** The objective of this step is to generate responses to the instructions obtained in Step 1.

We send these instructions to Llama-3-70B-Instruct and two additional powerful large language models( such as gpt4 and Qwen2-72B-Instruct). For each instruction, each model generates three responses, resulting in a total of nine responses for each instruction.

**Step 3: Filtering High-quality Responses.** For each instruction, we evaluate nine generated responses using the three models previously mentioned, assigning quality scores to each response. The scores from the three models are then averaged to identify the response with the highest overall score.

Combining the optimal response with the corresponding instruction forms the instruction dataset. The exact prompt we use for scoring is provided in Table 2.

#### 3.2 Data mixing and training.

Mix the extracted SFT data with our new SFT data, then proceed with training. The new SFT data accounts for 17% of all the data. The learning rate is set to 1e-6.

#### 3.3 Comparative experiment.

Use other open-source SFT datasets instead of the extracted SFT data for comparative experiments to identify which dataset used for experience replay results in less catastrophic forgetting.

**Baselines for Supervised Fine-Tuning and Preference Optimization.** These datasets include: **Evol Instruct** (Xu et al., 2023), **UltraChat** (Ding et al., 2023), **ShareGPT** (Chiang et al., 2023), **WildChat** (Zhao et al., 2024), **GenQA** (Chen et al., 2024), **OpenHermes 1** (Teknium, 2023b), **OpenHermes 2.5** (Teknium, 2023a), and **Tulu V2 Mix** (Iverson et al., 2023). ShareGPT and WildChat are representative human-written datasets containing 112K and 652K high-quality multi-round conversations between humans and GPT, respectively. Evol Instruct, UltraChat, and GenQA are representative open-source synthetic datasets. Following (Meng et al., 2024), we use the 208K sanitized version of Ultrachat provided by HuggingFace<sup>1</sup>. OpenHermes 1, OpenHermes 2.5, and Tulu V2 Mix are crowd-sourced datasets consisting of a mix of diverse open-source instruction datasets, with 243K, 1M, and 326K conversations, respectively.

We evaluated a variety of tasks featured on the Hugging Face Open LLM Leaderboard (Beeching et al., 2023), as presented in Table 1. The tasks include MMLU-PRO (Massive Multitask Language Understanding - Professional) (Wang et al., 2024), GPQA (Graduate-Level Google-Proof Q&A Benchmark) (Rein et al., 2023), IFEval (Zhou et al., 2023) and MATH level 5 (Hendrycks et al., 2021). Our experimental results demonstrate that employing our approach (extracting instruction distributions from the model) yields improved fine-tuning performance.

#### 3.4 Ablation Study

We tested the responses generated directly by the target model without using the three models for filtering, and the results are presented in Table 1. We also tested generating three responses solely by the target model without using the other two models

<sup>1</sup>[https://huggingface.co/datasets/HuggingFaceH4/ultrachat\\_200k](https://huggingface.co/datasets/HuggingFaceH4/ultrachat_200k)

233  
234  
235  
236  
237  
  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
  
274  
275  
276  
277  
278  
279

Alignment Setup	MMLU-PRO (5)	GPQA (0)	IFEval(0)	Math Lvl 5 (4)	Average
Llama-3-70B-Instruct	46.74	4.92	80.99	23.34	<b>39.00</b>
Extracted-Instructions-Unfiltered	46.11	4.72	81.31	23.11	38.81
One-Model-Filtered	46.55	4.91	81.72	23.06	39.06
<b>Three-Models-Mix-Filtered</b>	46.73	4.88	81.93	23.29	<b>39.21</b>
ShareGPT	46.14	4.31	81.31	20.24	38.00
Evol Instruct	45.76	4.64	82.52	22.30	38.81
GenQA	43.33	4.48	80.43	15.41	35.91
OpenHermes 1	45.31	4.21	81.91	15.52	36.74
OpenHermes 2.5	45.63	4.79	82.33	15.62	37.09
Tulu V2 Mix	46.47	4.19	82.69	16.62	37.49
WildChat	45.83	4.12	81.32	22.11	38.35
UltraChat	45.15	4.08	81.57	20.31	37.78

Table 1: This table compares the performance of models fine-tuned with supervision using the extracted instruction dataset for experience replay against baseline models and the official instruction model across various downstream benchmarks. All models are fine-tuned with supervision on the Llama-3-70B-Instruct model.

and then selecting the best one, and the results are presented in Table 1.

## 4 Conclusion

In this paper, we developed a method to extract instruction distributions from a model trained on an unpublished instruction dataset. We then leveraged two additional powerful models to collaboratively generate high-quality responses, forming an instruction dataset used as experience replay data during model fine-tuning. Compared to other baseline methods, our approach mitigates catastrophic forgetting and enhances fine-tuning performance.

## 5 Limitations

We conducted experiments only on Llama-3-70B-Instruct, achieving favorable results. Due to computational constraints, we did not perform extensive testing on other models.

## References

Edward Beeching, Clémentine Fourier, Nathan Habib, Sheon Han, Nathan Lambert, Nazneen Rajani, Omar Sanseviero, Lewis Tunstall, and Thomas Wolf. 2023. Open llm leaderboard.

Jiuhai Chen, Rifaa Qadri, Yuxin Wen, Neel Jain, John Kirchenbauer, Tianyi Zhou, and Tom Goldstein. 2024. Genqa: Generating millions of instructions from a handful of prompts. *arXiv preprint arXiv:2406.10323*.

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. 2023. *Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality*. 310  
311  
312

Claude. 2023. *Claude 2.1 model card*. Technical report, Claude Inc. 313  
314

Ning Ding, Yulin Chen, Bokai Xu, Yujia Qin, Zhi Zheng, Shengding Hu, Zhiyuan Liu, Maosong Sun, and Bowen Zhou. 2023. Enhancing chat language models by scaling high-quality instructional conversations. *arXiv preprint arXiv:2305.14233*. 315  
316  
317  
318  
319

Binzong Geng, Fajie Yuan, Qiancheng Xu, Ying Shen, Ruifeng Xu, and Min Yang. 2021. Continual learning for task-oriented dialogue system with iterative network pruning, expanding and masking. *arXiv preprint arXiv:2107.08173*. 320  
321  
322  
323  
324

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*. 325  
326  
327  
328  
329

Hamish Ivison, Yizhong Wang, Valentina Pyatkin, Nathan Lambert, Matthew Peters, Pradeep Dasigi, Joel Jang, David Wadden, Noah A Smith, Iz Beltagy, et al. 2023. Camels in a changing climate: Enhancing lm adaptation with tulu 2. *arXiv preprint arXiv:2311.10702*. 330  
331  
332  
333  
334  
335

Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*. 336  
337  
338  
339  
340

Seanie Lee, Hae Beom Lee, Juho Lee, and Sung Ju Hwang. 2021. Sequential reptile: Inter-task gradient alignment for multilingual learning. *arXiv preprint arXiv:2110.02600*. 341  
342  
343  
344

345	Michael McCloskey and Neal J. Cohen. 1989. <a href="#">Catastrophic interference in connectionist networks: The sequential learning problem</a> . 24:109–165.	Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. <i>arXiv preprint arXiv:2406.01574</i> .	398
346			399
347			400
348	Yu Meng, Mengzhou Xia, and Danqi Chen. 2024. Simpo: Simple preference optimization with a reference-free reward. <i>arXiv preprint arXiv:2405.14734</i> .	Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. 2023. Wizardlm: Empowering large language models to follow complex instructions. <i>arXiv preprint arXiv:2304.12244</i> .	401
349			402
350			403
351			404
352	Fei Mi, Liangwei Chen, Mengjie Zhao, Minlie Huang, and Boi Faltings. 2020. Continual learning for natural language generation in task-oriented dialog systems. <i>arXiv preprint arXiv:2010.00910</i> .	Wenting Zhao, Xiang Ren, Jack Hessel, Claire Cardie, Yejin Choi, and Yuntian Deng. 2024. <a href="#">Wildchat: 1m chatGPT interaction logs in the wild</a> . In <i>The Twelfth International Conference on Learning Representations</i> .	406
353			407
354			408
355			409
356	OpenAI. 2023. <a href="#">Gpt-4 technical report</a> .		410
357	Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. The refinedweb dataset for falcon llm: outperforming curated corpora with web data, and web data only. <i>arXiv preprint arXiv:2306.01116</i> .	Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023. Instruction-following evaluation for large language models. <i>arXiv preprint arXiv:2311.07911</i> .	411
358			412
359			413
360			414
361			415
362			
363			
364	David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. 2023. Gpqa: A graduate-level google-proof q&a benchmark. <i>arXiv preprint arXiv:2311.12022</i> .	<b>A Appendix</b>	416
365			
366			
367			
368			
369	Weijieying Ren, Xinlong Li, Lei Wang, Tianxiang Zhao, and Wei Qin. 2024. Analyzing and reducing catastrophic forgetting in parameter efficient tuning. <i>arXiv preprint arXiv:2402.18865</i> .		
370			
371			
372			
373	ANTHONY ROBINS. 1995. <a href="#">Catastrophic forgetting, rehearsal and pseudorehearsal</a> . <i>Connection Science</i> , 7(2):123–146.		
374			
375			
376	Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. 2017. Continual learning with deep generative replay. <i>Advances in neural information processing systems</i> , 30.		
377			
378			
379			
380	Teknum. 2023a. <a href="#">Openhermes 2.5: An open dataset of synthetic data for generalist llm assistants</a> .		
381			
382	Teknum. 2023b. <a href="#">Openhermes dataset</a> .		
383	Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. <i>arXiv preprint arXiv:2302.13971</i> .		
384			
385			
386			
387			
388			
389	Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. <i>arXiv preprint arXiv:2307.09288</i> .		
390			
391			
392			
393			
394			
395	Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, et al. 2024.		
396			
397			

<p>Below is a user instruction and an AI response. Evaluate the quality of the AI's response based on how well it fulfills the user's request. Assign a score based on the following 5-point scale:</p> <p>1: The response is incomplete, off-topic, or contains irrelevant, vague, or missing information. It may repeat the user's question, include personal opinions, or be written from a non-AI perspective (e.g., blog-like). It may also have promotional or irrelevant content.</p> <p>2: The response addresses some of the user's request but lacks detail or direct relevance. It provides only a general approach instead of a specific solution.</p> <p>3: The response is helpful but lacks an AI perspective. It covers the user's request but appears taken from a personal blog, webpage, or similar source. It may include personal opinions, experiences, or mentions of external content.</p> <p>4: The response is clear, complete, and written from an AI's perspective. It directly addresses the user's request, but there may be minor room for improvement, such as clarity or conciseness.</p> <p>5: The response is excellent, written from an AI's perspective, with a clear focus on the user's request. It is thorough, well-organized, and shows expert knowledge without irrelevant content. The response is logical, easy to follow, and engaging.</p> <p>Provide a brief justification for your score and then write "Score: &lt;rating&gt;" in the last line.</p> <p>&lt;generated instruction&gt;</p> <p>&lt;output&gt;</p>
---

Table 2: A prompt used to evaluate the quality of a response.