

Revisiting Catastrophic Forgetting in Large Language Model Tuning

Hongyu Li
Wuhan University
hongyuli@whu.edu.cn

Liang Ding*
The University of Sydney
liangding.liam@gmail.com

Meng Fang
University of Liverpool
Meng.Fang@liverpool.ac.uk

Dacheng Tao
Nanyang Technological University
dacheng.tao@ntu.edu.sg

Abstract

Catastrophic Forgetting (CF) means models forgetting previously acquired knowledge when learning new data. It compromises the effectiveness of large language models (LLMs) during fine-tuning, yet the underlying causes have not been thoroughly investigated. This paper takes the first step to reveal the direct link between the flatness of the model loss landscape and the extent of CF in the field of LLMs. Based on this, we introduce the sharpness-aware minimization to mitigate CF by flattening the loss landscape. Experiments on three widely-used fine-tuning datasets, spanning different model scales, demonstrate the effectiveness of our method in alleviating CF. Analyses show that we nicely complement the existing anti-forgetting strategies, further enhancing the resistance of LLMs to CF.

1 Introduction

Instruction fine-tuning is key to improving the capabilities and controllability of large language models (LLMs) (Brown et al., 2020; Fedus et al., 2022), which have already demonstrated strong performance in various tasks (Zhong et al., 2023; Peng et al., 2023; Lu et al., 2023; Ren et al., 2024). One major obstacle to tuning LLMs is catastrophic forgetting (CF, Kirkpatrick et al., 2017), which means LLMs forget prior knowledge when learning new data. Recent works have provided substantial evidence confirming the negative impact of CF on LLMs, e.g., Bi et al. (2024) empirically show that the fine-tuned model is even worse than their foundation counterpart on several tasks, and Zeng et al. (2023) reveal the dropped general performance of LLMs after developing their agent capabilities.

Understanding the effectiveness of LLMs during fine-tuning is important for downstream tasks, however, the underlying causes of CF remain largely unexplored. There are two types of work addressing

CF, from data and model perspectives, respectively. Scialom et al. (2022) propose continual learning with rehearsal in instruction tuning, though its effectiveness varies with task selection. Lin et al. (2023) indicate that both continual learning and weight averaging (Wise-FT) effectively preserve generality.

While the above techniques are somewhat successful, they require expensive extra data-constructing and training costs, and are even sometimes impractical in LLMs, because a) the data cards of many pretrained models are unclear (Shi et al., 2023a), making rehearsal unfeasible, and b) anti-forgetting training brings an unstable and expensive training process (Datta et al., 2023).

In this work, we turn to finding a cheap, stable, and orthogonal solution to alleviate the CF in tuning LLMs. In particular, we ① reveal the high correlation between the extent of CF and the flatness of the loss landscape (LLS, Keskar et al., 2017; Dziugaite and Roy, 2017), ② mitigate the CF in LLMs by flattening the LLS from optimization perspective, ③ show the complementarity between our method and existing anti-forgetting works. Specifically, we designed three probing analyses to achieve ①, and found that a flatter LLS could reduce the severity of CF. Based on our observation, we introduced “Sharpness-Aware Minimization” (SAM, Foret et al., 2021) to flatten the model LLS to approach ②. For ③, we found that our introduced optimizer nicely complements a series of anti-forgetting methods, including rehearsal (Scialom et al., 2022) and Wise-FT (Lin et al., 2023). Our **contributions** are as follows:

- To the best of our knowledge, we are the first to empirically reveal a direct correlation between the flatness of the model LLS and CF in the field of LLMs.
- We present the first work that mitigates CF of LLMs from the perspective of optimization.

*Corresponding author.




Figure 1: **Visualization of different models’ loss landscapes (LLS)** with contour lines. We can see that with the data/task gaps of (a), (b), (c) gradually increase, the disturbance of their contours becomes more obvious.

- The proposed method can be synergistically combined with existing methods to enhance the resilience of LLMs against forgetting.

2 Unveiling the Hidden Nexus: CF and the Model LLS

We investigate the relationship between CF and the flatness of LLS through three analyses: *i) loss landscape visualization*, *ii) flatness degree of LLS*, and *iii) general task performance*, where the visualized comparison of *i)* and results of *ii) & iii)* are discussed in § 2.1 and § 2.2, respectively, to intuitively and quantitatively reveal their correlation.

We simulate continue-learning scenarios with three models: Llama2-7b (Touvron et al., 2023) tuned on 1) “Alpaca (Taori et al., 2023)”, 2) “Alpaca→Open-Platypus (Lee et al., 2023)”, and 3) “Alpaca→Auto-Wiki (Jiang et al., 2020; Alva-Manchego et al., 2020)”, where “→” means continue-train on the following dataset and the data gaps are gradually increased. For the experimental details, please refer to the Appendix A.1.1.

2.1 Visualized LLS comparison of different CF settings

Following Hao et al. (2019); Zan et al. (2022), we visualize the 2D loss surface to explore the differences in the LLS of models under different settings. We represent the weight of a well-trained model as θ_0 and define the loss surface with the function: $f(\alpha, \beta) = \mathcal{L}(\theta_0 + \alpha\delta_1 + \beta\delta_2)$, where \mathcal{L} is the loss function, α and β are scalar values representing the current coordinate, with δ_1 and δ_2 as two direction vectors. We choose Gaussian noise as the basis vectors for exploring the parameter space, and the axes correspond to α and β values. This method allows us to simulate a random direction in

	Flatness Degree			General Perf.
	SC	AG	MAG	MMLU
(a)Alpaca	52.87	105.37	65.53	40.53
(b)→Open-Platypus	52.98	106.41	68.91	33.46 \downarrow -7.1
(c)→Auto-Wiki	53.77	111.03	70.71	23.31 \downarrow -17.2

Table 1: **Quantitative results of the LLS flatness (“Flatness Degree”) of LLMs and their general task performance (“General Perf.”).**

the parameter space for the loss landscape visualization, which helps ensure a broad and generalized exploration of the LLS’ curvature.

From the LLS visualization in Fig. 1, we observe that with the continue-learned task gap increases, i.e., (a)≤(b)≤(c), their landscape and contour lines become sharper and more disturbed, *intuitively* demonstrating the **high correlation between the flatness degree** (of model loss landscape) **and the forgetting degree** (of continually tuning LLMs).

2.2 Quantitative flatness degree and general performance of different CF settings

Besides intuitive visualization, we quantitatively measure the *flatness degree of loss landscape* (with three metrics: ‘SC’ (Surface Curvature, %), ‘AG’ (Average Gradient, %), ‘MAG’ (Mean Absolute Gradient, %)), and the *model performance on general tasks* (in terms of MMLU benchmark (Hendrycks et al., 2021a)). Details of our proposed metrics can see Appendix A.3.

As seen in Tab. 1, for different CF settings (difficulty rank: (a)≤(b)≤(c)), their performance on downstream tasks, i.e., MMLU tasks, is significantly dropped (up to -17.2 on average), confirming our hypothesis that **continually training on harder tasks (with larger data gaps) will exacerbate CF** problem. Also, we find that the flatness degree

rank is as follows: (a) \geq (b) \geq (c), where their average values are 74.59, 76.1, and 78.50, respectively, implying that *there is a highly positive correlation between the severity of CF and the sharpness of the loss landscape*.

3 Improving CF with Sharpness-Aware Optimization

Motivated by the above findings, we naturally consider introducing the Sharpness-Aware Minimization (SAM, Foret et al., 2021) method to flatten the model loss landscape, thereby mitigating the CF problem. Given model weights as $w \in \mathbb{R}^d$, loss function as f , and a training dataset $S = \{(x_i, y_i)\}_{i=1}^n$ *i.i.d.* sampled from the distribution D , the working mechanism of SAM can be formulated as follows:

$$\min_w \max_{\|\epsilon\|_2 \leq \rho} f(w + \epsilon), \quad (1)$$

where ϵ is the perturbation in a neighbourhood ball area with a radius denoted as ρ . The optimization objective is to ensure that the loss f does not increase substantially with ϵ constrained by ρ . Further utilizing Taylor series expansion, ϵ can be approximate as $\rho \nabla_w f(w) / \|\nabla_w f(w)\|_2$. In case, Eq. 1 can be simplified as:

$$\min_w f\left(w + \rho \frac{\nabla_w f(w)}{\|\nabla_w f(w)\|_2}\right). \quad (2)$$

Eq. 2 could be solved through a two-step gradient descent process. Firstly, the perturbation ϵ is computed by $\rho \nabla_w f(w) / \|\nabla_w f(w)\|_2$. Following this, the second phase of gradient descent performs the actual update of the weights.

4 Experimental Setting

Dataset To validate the capability of SAM on diverse datasets, we conduct experiments on four widely-used instruction fine-tuning datasets: Alpaca (Taori et al., 2023), ShareGPT52K (Chiang et al., 2023), MetaMathQA (Yu et al., 2023), and Open-Platypus (Lee et al., 2023). For detailed introduction, please refer to the Appendix A.1.1.

Model Selection To further demonstrate the compatibility of SAM with models of different sizes, we conducted experiments on models including TinyLlama-1.1B (Zhang et al., 2024), Llama2-7B, and Llama2-13B (Touvron et al., 2023).

Implementation Details For a fair comparison across the same pre-trained language model, our initial phase follows (Taori et al., 2023) to train an instruction fine-tuning model on the Alpaca dataset. This baseline model underpins our subsequent fine-tuning experiments on various datasets. Besides, all hyper-parameters and settings are maintained uniformly in the following fine-tuning. We opt for the AdamW (Loshchilov and Hutter, 2019) optimizer in scenarios without SAM. All experiments are executed on 16 NVIDIA A800 GPUs.

Evaluation Tasks To effectively assess the extent of the general knowledge retained in LLMs and to quantify the degree of CF, we have implemented a collection of general assessment tasks. These tasks are categorized as follows: (1) **Domain knowledge benchmark (DK)**: MMLU (Hendrycks et al., 2021a), including STEM, Humanities, Social Sciences, and Other domains. (2) **Reasoning**: SuperGlue (Wang et al., 2019) include AX_b , AX_g , RTE, COPA, and datasets like Hellaswag (Zellers et al., 2019), Boolq (Clark et al., 2019), and Siqa (Sap et al., 2019). (3) **Understanding**: RACE (middle & high) (Lai et al., 2017), Openbookqa-fact (Mihaylov et al., 2018), and Csl-dev (Li et al., 2022). (4) **Exams**: ARC-c (Clark et al., 2018). Besides, we also conduct experiments on domain-specific datasets like TruthfulQA (Lin et al., 2022) to assess the degree of retention of specific abilities. Refer to Appendix A.1.2 for more details.

5 Experimental Results

We demonstrate the superiority of our proposed method through four experimental analyses:

1. Mitigating CF on different datasets
2. Mitigating CF across different model sizes
3. Comparison with other advanced methods
4. Complementary to existing works

The average results for 4 evaluation categories are presented in Tab. 2 (values in %), where “(w/o)” and “(w/)” refers to train without and with SAM, respectively. The four experiments are as follows, and the total experimental results can be referred to the Appendix A.4.

Performance of SAM in Mitigating CF on different datasets As presented in Tab. 2(a),

Dataset	DK	Understanding	Reasoning	Exams	AVG	Δ
Alpaca	40.53	58.74	63.33	45.08	51.92	—
→ShareGPT(w/o)	26.08	52.84	58.68	45.76	45.84	-6.08
→ShareGPT(w/)	40.08	57.91	63.78	44.41	51.55	+5.71
→Open-Platypus(w/o)	31.13	50.70	61.09	37.97	45.22	-6.70
→Open-Platypus(w/)	41.07	58.28	64.50	45.08	52.23	+7.01
→Meta-Math(w/o)	33.13	52.61	58.46	36.27	45.12	-6.80
→Meta-Math(w/)	34.79	55.77	62.38	42.71	48.91	+3.79

(a) General Performance of Evaluation Tasks on SAM (w/o) & (w/) across Different Datasets.

Model	Dataset	DK	Understanding	Reasoning	Exams	AVG	Δ
TinyLlama	Alpaca	23.16	30.62	46.16	26.78	31.68	—
	→Open-Platypus(w/o)	23.04	31.06	46.23	27.12	31.86	+0.18
	→Open-Platypus(w/)	23.14	30.39	46.91	26.10	31.64	-0.22
Llama2-7B	Alpaca	40.53	58.74	63.33	45.08	51.92	—
	→Open-Platypus(w/o)	31.13	50.70	61.09	37.97	45.22	-6.70
	→Open-Platypus(w/)	41.07	58.28	64.50	45.08	52.23	+7.01
Llama2-13B	Alpaca	48.15	69.90	64.37	63.73	61.54	—
	→Open-Platypus(w/o)	28.23	61.92	64.12	54.58	52.21	-9.33
	→Open-Platypus(w/)	49.38	69.80	65.72	63.05	61.99	+9.78

(b) General Performance of Evaluation Tasks on SAM (w/o) & (w/) on Different Model Sizes.

Method	DK	Understanding	Reasoning	Exams	AVG	Δ
Alpaca	40.53	58.74	63.33	45.08	51.92	—
→Open-Platypus(w/o)	31.13	50.70	61.09	37.97	45.22	-6.70
→Open-Platypus(w/)	41.07	58.28	64.50	45.08	52.23	+7.01
Wise-FT (w/o)	37.75	56.64	62.65	47.12	51.04	-0.88
Wise-FT (w/)	40.59	58.14	64.56	44.75	52.01	+0.97
Rehearsal (w/o)	33.38	54.69	61.25	43.05	48.09	-3.83
Rehearsal (w/)	40.35	57.09	63.27	43.73	51.11	+3.02

(c) General Performance of Evaluation Tasks on SAM Comparing/Combining with Wise-FT and Rehearsal.

Table 2: General Performance of Evaluation Tasks on Four Analyses of SAM.

for different CF settings (Alpaca as the baseline, difficulty rank: ShareGPT52K \leq Open-Platypus \leq MetaMathQA), their performance on evaluation tasks experiences a significant decline (Δ^1 CF: ShareGPT: 11.71%; Open-Platypus: 12.90%; Meta-Math: 13.10%), confirming that different degrees of CF occur in the fine-tuning phase. **With the introduction of SAM, the CF is markedly mitigated** and the performance is comparable to baseline. Notably, the Tab. 2 shows the results in balanced training cost, i.e., (w/o) epoch=2 and (w/) epoch=1.

Given the need for a fair comparison and to balance computational cost of SAM, the number of epochs for training differs between SAM (w/o) and (w/). It is crucial to present the SAM (w/o) epoch=1 results in Tab. 3 to uphold our experimental rigor. Results show SAM (w/o) epoch=2 generally outperforms SAM (w/o) epoch=1. Remarkably, **SAM (w/) epoch=1 consistently achieved the best results**, highlighting our method’s superiority.

¹ Δ represents the percentage of the increase relative to the original value.

Another important concern is preserving capability within the domain. Our approach effectively alleviates CF while minimally impacting the acquisition of new knowledge, which is supported on the Open-Platypus dataset. As shown in Tab. 4, it is clear that **SAM mitigates CF with little impact of specific domain performance in LLMs** (e.g., Δ AVG: SAM (w/o): 10.7%; SAM (w/): 11.2%).

Performance of SAM on LLMs of Different Sizes For the results in Tab. 2(b), we conduct Alpaca \rightarrow Open-Platypus across models of different sizes, and they exhibited different degrees of performance dropping (e.g., Δ DK: TinyLlama-1.1B: 0.52%; Llama2-7B: 23.19%; Llama2-13B: 41.37%), indicating that **the severity of the forgetting increasing with model size**. Compared to the performance of SAM in mitigating CF across models, **SAM could significantly mitigate the extent of CF with model size increasing**.

In Comparison with Other Advanced Anti-Forgetting Methods For the comparative experimental results represents in Tab.2(c), **SAM sig-**

Epoch	Dataset	DK	Understanding	Reasoning	Exams	AVG
2	→ShareGPT(w/o)	26.08	52.84	58.68	45.76	45.84
1	→ ShareGPT(w/o)	24.98	52.40	59.12	45.08	45.39
2	→Open-Platypus(w/o)	31.13	50.70	61.09	37.97	45.22
1	→ Open-Platypus(w/o)	29.89	49.37	60.60	40.34	45.05
2	→Meta-Math(w/o)	33.13	52.61	58.46	36.27	45.12
1	→ Meta-Math(w/o)	32.84	52.69	58.43	38.64	45.65

Table 3: General Performance of **Different Epochs** on Llama2-7b.

Dataset	Hellaswag	TruthfulQA	AVG
Alpaca	65.00	22.64	43.82
→Open-Platypus(w/o)	73.16	23.87	48.52 ^{↑+4.7}
→Open-Platypus(w/)	73.49	23.99	48.74 ^{↑+4.9}

Table 4: Performance in Fine-tuning Domain.

nificantly outperforms Wise-FT and Rehearsal, although the other two also achieve decent performance (Δ Mitigating CF: SAM: 7.01%; Wise-FT: 5.82%; Rehearsal: 3.79%).

Complementary to Existing Advanced Anti-Forgetting Methods For the results in Tab.2(c), SAM could effectively improve the performance of CF along with the two methods (Δ Mitigating CF: Wise-FT(w/): 0.97%; Rehearsal(w/): 3.02%), which illustrates that *SAM can be orthogonally combined with other methods*, providing incremental benefits to mitigate CF.

6 Related Work

In continual learning, approaches to mitigate CF have been extensively explored, including a) replay-based methods for knowledge retention (Shin et al., 2017; Lopez-Paz and Ranzato, 2017; Chaudhry et al., 2018); b) regularization strategies to safeguard existing information (Deng et al., 2021; Shi et al., 2023b); and c) architectural modifications to facilitate new learning (Shi et al., 2023b; Wen et al., 2020). While traditional techniques have achieved some success with small-scale models, their effectiveness diminishes when applied to LLMs due to the high predictability and unpredictability of LLMs (Ganguli et al., 2022). LLMs encounter a more severe CF problem in continue learning.

Research on CF in LLMs primarily follows two approaches. The first is data-centric, utilizing replay techniques to incorporate past data or generate synthetic experiences (Scialom et al., 2022; Sun et al., 2019), which helps in retaining old knowledge. The success of this method heavily depends on the quality of the selected data, as poor

choices can undermine its effectiveness. The second approach is model-centric, employing integration strategies to combine fine-tuned and original models (Lin et al., 2023; Xiao et al., 2023), enhancing generalization and preventing forgetting. However, the effectiveness is constrained by the model’s architecture, capacity, and the quality of initial training.

7 Conclusion

We reexamine the phenomenon of Catastrophic Forgetting (CF) and reveal the direct link between the flatness of the model loss landscape (LLS) and the extent of CF in LLMs. Based on our findings, we introduce sharpness-aware optimization to flatten the LLS, thereby mitigating CF. Experiments show its effectiveness, especially in larger models, suggesting our method may become a standard strategy to mitigate CF during LLMs tuning.

In future work, we would like to explore more strategies to orthogonally address the forgetting problem in LLM tuning, such as advanced optimizers (Sun et al., 2024; Zhong et al., 2022; Yang et al., 2023), proper anti-forgetting learning curricula (Ding et al., 2021; Zhou et al., 2021; Faber et al., 2024), and model merging techniques (Li et al., 2023; He et al., 2023).

Acknowledgments

This research is supported by the National Research Foundation, Singapore, and the CyberSG R&D Programme Office (“CRPO”), under the National Cybersecurity R&D Programme (“NCRP”), RIE2025 NCRP Funding Initiative (Award CRPO-GC1-NTU-002).

Limitations

Despite the progress we made, there still exist limitations in our work. On the one hand, our experiments are concentrated on a specific aspect of the Catastrophic Forgetting (CF) — the direct link between the flatness of the model’s loss landscape

and CF. While this finding makes us introduce the sharpness-aware minimization technique to flatten the loss landscape and thereby mitigate CF, it also means that other potentially influential factors in CF have not been comprehensively explored. On the other hand, our research primarily addresses CF during the fine-tuning phase of LLMs, but it does not directly tackle CF that may occur during other stages. For instance, CF can also be a significant challenge when models are updated with new data post-deployment. Therefore, the scope of our findings and the applicability of our solution may be limited to specific phases of the LLMs lifecycle.

Ethics Statement

We take ethical considerations seriously and strictly adhere to the ACL Ethics Policy. This paper revisits a major obstacle of fine-tuning the large language models – Catastrophic Forgetting, and provides an optimization-level solution to alleviate it. All employed models and datasets in this paper are publicly available and have been widely adopted by researchers. All experimental results upon these open models and datasets are reported accurately and objectively. Thus, we believe that this research will not pose any ethical issues.

Reproducibility. We provide the detailed experimental setup in the Appendix, such as hyperparameters and statistic descriptions. Also, we release code to facilitate reproducing the experimental results: https://github.com/Li-Hyn/LLM_CatastrophicForgetting.

References

- Fernando Alva-Manchego, Louis Martin, Antoine Bordes, Carolina Scarton, Benoît Sagot, and Lucia Specia. 2020. *Asset: A dataset for tuning and evaluation of sentence simplification models with multiple rewriting transformations*. *arXiv preprint*.
- Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiushi Du, Zhe Fu, et al. 2024. *Deepseek llm: Scaling open-source language models with longtermism*. *arXiv preprint*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. *Language models are few-shot learners*. In *NeurIPS*.
- Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. 2018. *Efficient lifelong learning with a-gem*. *arXiv preprint*.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. 2023. *Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality*. *blog*.
- Clark, Christopher, Lee, Kenton, Chang, Ming-Wei, Kwiatkowski, Tom, Collins, Michael, and Kristina Toutanova. 2019. *Boolq: Exploring the surprising difficulty of natural yes/no questions*. In *NAACL*.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. *Think you have solved question answering? try arc, the ai2 reasoning challenge*. *arXiv preprint*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. *Training verifiers to solve math word problems*. *arXiv preprint*.
- Arghya Datta, Subhrangshu Nandi, Jingcheng Xu, Greg Ver Steeg, He Xie, Anoop Kumar, and Aram Galstyan. 2023. *Measuring and mitigating local instability in deep neural networks*. *arXiv preprint*.
- Danruo Deng, Guangyong Chen, Jianye Hao, Qiong Wang, and Pheng-Ann Heng. 2021. *Flattening sharpness for dynamic gradient projection memory benefits continual learning*. In *NeurIPS*.
- Liang Ding, Longyue Wang, Xuebo Liu, Derek F Wong, Dacheng Tao, and Zhaopeng Tu. 2021. *Progressive multi-granularity training for non-autoregressive translation*. In *Findings of ACL*.
- Gintare Karolina Dziugaite and Daniel M. Roy. 2017. *Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data*. In *UAI*.
- Kamil Faber, Dominik Zurek, Marcin Pietron, Nathalie Japkowicz, Antonio Vergari, and Roberto Corizzo. 2024. *From mnist to imagenet and back: benchmarking continual curriculum learning*. *Machine Learning*.
- William Fedus, Barret Zoph, and Noam Shazeer. 2022. *Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity*. *JMLR*.
- Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. 2021. *Sharpness-aware minimization for efficiently improving generalization*. In *ICLR*.
- Deep Ganguli, Danny Hernandez, Liane Lovitt, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova Dassarma, Dawn Drain, Nelson Elhage, et al. 2022. *Predictability and surprise in large generative models*. In *ACM FAccT*.

- Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and William B Dolan. 2007. [The third pascal recognizing textual entailment challenge](#). In *ACL-PASCAL*.
- Yaru Hao, Li Dong, Furu Wei, and Ke Xu. 2019. [Visualizing and understanding the effectiveness of bert](#). In *EMNLP-IJCNLP*.
- Shwai He, Run-Ze Fan, Liang Ding, Li Shen, Tianyi Zhou, and Dacheng Tao. 2023. [Merging experts into one: Improving computational efficiency of mixture of experts](#). In *EMNLP*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021a. [Measuring massive multitask language understanding](#). In *ICLR*.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021b. [Measuring mathematical problem solving with the MATH dataset](#). In *NeurIPS*.
- Chao Jiang, Mounica Maddela, Wuwei Lan, Yang Zhong, and Wei Xu. 2020. [Neural crf model for sentence alignment in text simplification](#). *arXiv preprint*.
- Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. 2017. [On large-batch training for deep learning: Generalization gap and sharp minima](#). In *ICLR*.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. [Overcoming catastrophic forgetting in neural networks](#). *PNAS*.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard H. Hovy. 2017. [RACE: large-scale reading comprehension dataset from examinations](#). In *EMNLP*.
- Ariel N. Lee, Cole J. Hunter, and Nataniel Ruiz. 2023. [Platypus: Quick, cheap, and powerful refinement of llms](#). *arXiv preprint*.
- Weishi Li, Yong Peng, Miao Zhang, Liang Ding, Han Hu, and Li Shen. 2023. [Deep model fusion: A survey](#). *arXiv preprint*.
- Yudong Li, Yuqing Zhang, Zhe Zhao, Linlin Shen, Weijie Liu, Weiwan Mao, and Hui Zhang. 2022. [Csl: A large-scale chinese scientific literature dataset](#). *arXiv preprint*.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. [Truthfulqa: Measuring how models mimic human falsehoods](#). In *ACL*.
- Yong Lin, Lu Tan, Hangyu Lin, Zeming Zheng, Renjie Pi, Jipeng Zhang, Shizhe Diao, Haoxiang Wang, Han Zhao, Yuan Yao, et al. 2023. [Speciality vs generality: An empirical study on catastrophic forgetting in fine-tuning foundation models](#). *arXiv preprint*.
- David Lopez-Paz and Marc’Aurelio Ranzato. 2017. [Gradient episodic memory for continual learning](#). In *NIPS*.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *ICLR*.
- Qingyu Lu, Baopu Qiu, Liang Ding, Kanjian Zhang, Tom Kocmi, and Dacheng Tao. 2023. [Error analysis prompting enables human-like translation evaluation in large language models: A case study on chatgpt](#). *arXiv preprint*.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. [Can a suit of armor conduct electricity? a new dataset for open book question answering](#). In *EMNLP*.
- Keqin Peng, Liang Ding, Qihuang Zhong, Li Shen, Xuebo Liu, Min Zhang, Yuanxin Ouyang, and Dacheng Tao. 2023. [Towards making the most of chatgpt for machine translation](#). In *Findings of EMNLP*.
- Zhiyao Ren, Yibing Zhan, Baosheng Yu, Liang Ding, and Dacheng Tao. 2024. [Healthcare copilot: Eliciting the power of general llms for medical consultation](#). *arXiv preprint*.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan LeBras, and Yejin Choi. 2019. [Socialliqa: Commonsense reasoning about social interactions](#). *arXiv preprint*.
- Thomas Scialom, Tuhin Chakrabarty, and Smaranda Muresan. 2022. [Fine-tuned language models are continual learners](#). In *EMNLP*.
- Weijia Shi, Anirudh Ajith, Mengzhou Xia, Yangsibo Huang, Daogao Liu, Terra Blevins, Danqi Chen, and Luke Zettlemoyer. 2023a. [Detecting pretraining data from large language models](#). In *NeurIPS*.
- Wenheng Shi, Yiren Chen, Zhe Zhao, Wei Lu, Kimmo Yan, and Xiaoyong Du. 2023b. [Create and find flatness: Building flat training spaces in advance for continual learning](#). In *ECAI*.
- Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. 2017. [Continual learning with deep generative replay](#). In *NIPS*.
- Fan-Keng Sun, Cheng-Hao Ho, and Hung-Yi Lee. 2019. [Lamol: Language modeling for lifelong language learning](#). *arXiv preprint*.
- Hao Sun, Li Shen, Qihuang Zhong, Liang Ding, Shixiang Chen, Jingwei Sun, Jing Li, Guangzhong Sun, and Dacheng Tao. 2024. [Adasam: Boosting sharpness-aware minimization with adaptive learning rate and momentum for training deep neural networks](#). *Neural Networks*.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. [Stanford alpaca: An instruction-following llama model](#). *Stanford blog*.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint*.

Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. [Superglue: A stickier benchmark for general-purpose language understanding systems](#). In *NeurIPS*.

Yeming Wen, Dustin Tran, and Jimmy Ba. 2020. [Batchensemble: an alternative approach to efficient ensemble and lifelong learning](#). *arXiv preprint*.

Shitao Xiao, Zheng Liu, Peitian Zhang, and Xingrun Xing. 2023. [Lm-cocktail: Resilient tuning of language models via model merging](#). *arXiv preprint*.

Enneng Yang, Li Shen, Zhenyi Wang, Shiwei Liu, Guibing Guo, and Xingwei Wang. 2023. [Data augmented flatness-aware gradient projection for continual learning](#). In *ICCV*.

Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. 2023. [Meta-math: Bootstrap your own mathematical questions for large language models](#). *arXiv preprint*.

Changtong Zhan, Liang Ding, Li Shen, Yu Cao, Weifeng Liu, and Dacheng Tao. 2022. [On the complementarity between pre-training and random-initialization for resource-rich machine translation](#). In *COLING*.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [Hellaswag: Can a machine really finish your sentence?](#) In *ACL*.

Aohan Zeng, Mingdao Liu, Rui Lu, Bowen Wang, Xiao Liu, Yuxiao Dong, and Jie Tang. 2023. [Agenttuning: Enabling generalized agent abilities for llms](#). *arXiv preprint*.

Peiyuan Zhang, Guangtao Zeng, Tianduo Wang, and Wei Lu. 2024. [Tinyllama: An open-source small language model](#). *arXiv preprint*.

Qihuang Zhong, Liang Ding, Juhua Liu, Bo Du, and Dacheng Tao. 2023. [Can chatgpt understand too? a comparative study on chatgpt and fine-tuned bert](#). *arXiv preprint*.

Qihuang Zhong, Liang Ding, Li Shen, Peng Mi, Juhua Liu, Bo Du, and Dacheng Tao. 2022. [Improving sharpness-aware minimization with fisher mask for better generalization on language models](#). In *Findings of EMNLP*.

Lei Zhou, Liang Ding, Kevin Duh, Shinji Watanabe, Ryohhei Sasano, and Koichi Takeda. 2021. [Self-guided curriculum learning for neural machine translation](#). In *IWSLT*.

A Appendix

A.1 Details of Training and Evaluation Datasets

As outlined in §4 and §2, our study involves extensive experiments across five widely-utilized datasets. Additionally, we assess our model’s performance on selected tasks from the SuperGLUE benchmark along with other distinct datasets. In this section, we provide detailed descriptions of the datasets and tasks employed. The description of each task includes:

A.1.1 Training Datasets

1. **Alpaca**: Alpaca (Taori et al., 2023) is a dataset of 52,000 instructions and demonstrations generated by OpenAI’s text-davinci-003 engine. This instruction data can be used to conduct instruction-tuning for language models and make the language model follow instruction better.
2. **ShareGPT52K**: ShareGPT52K (Chiang et al., 2023) is a collection of approximately 52,000 conversations scraped via the ShareGPT API. These conversations include both user prompts and responses from OpenAI’s ChatGPT.
3. **MetaMathQA**: MetaMathQA (Yu et al., 2023) is a dataset of 395,000 instructions which augmented from the training sets of GSM8K (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021b).
4. **Open-Platypus**: Open-Platypus (Lee et al., 2023) is focused on improving LLM logical reasoning skills and has the item amount to 24,900.
5. **Auto-Wiki**: Auto-Wiki (Jiang et al., 2020; Alva-Manchego et al., 2020) provides a set of aligned sentences from English Wikipedia and Simple English Wikipedia as a resource to train sentence simplification systems.

A.1.2 Evaluation Tasks

1. **MMLU**: Massive Multitask Language Understanding benchmark (MMLU) (Hendrycks et al., 2021a) to evaluate the knowledge stored in the LLMs, which can be divided into STEM, Human, Social, and Other.

2. **AX_b**: AX-b (Wang et al., 2019) is a broad-coverage diagnostic task, which requires to determine the logical relation between the given sentence pair, with three relations: entailment, contradiction and neutral.
3. **AX_g**: AX-g (Wang et al., 2019) is a Winogender diagnostic task, which requires to determine which noun the pronoun refers to according to the given sentence and pronoun.
4. **RTE**: Recognizing Textual Entailment (Giampiccolo et al., 2007), given a premise and a hypothesis, is a task to predict whether the premise entails the hypothesis.
5. **COPA**: COPA (Wang et al., 2019) is a causal inference task, which requires to select the correct causal relation based on the given premise.
6. **Hellaswag**: HellaSwag (Zellers et al., 2019) is a challenge dataset for evaluating common-sense natural language inference, which is specially hard for state-of-the-art models, though its questions are trivial for humans (>95% accuracy).
7. **Boolq**: Boolean Question (Clark et al., 2019) is a question-answering task where each sample consists of a short passage and a yes/no question about the passage.
8. **Siqa**: Siqa (Sap et al., 2019) is a social interaction question-answering task, which requires selecting the most reasonable behaviour based on the given scenario and three possible subsequent behaviours.
9. **RACE**: RACE (Lai et al., 2017) is a large-scale reading comprehension dataset with more than 28,000 passages and nearly 100,000 questions.
10. **Openbookqa-fact**: OpenBookQA (Mihaylov et al., 2018) contains questions that require multi-step reasoning, application of common-sense knowledge, and in-depth comprehension of text.
11. **Csl-dev**: CSL (Li et al., 2022) is a large-scale Chinese Scientific Literature dataset, which contains the titles, abstracts, keywords and academic fields of 396k papers
12. **ARC-c**: The AI2’s Reasoning Challenge (ARC) (Clark et al., 2018) is a multiple-choice question-answering dataset, containing questions from science exams from grade 3 to grade 9.
13. **TruthfulQA**: TruthfulQA (Lin et al., 2022) is a benchmark to measure whether a language model is truthful in generating answers to questions.

A.2 Experimental Setup

In this paper, we trained from the base model (TinyLlama, Llama2-7b, and Llama2-13B). For a fair comparison, we follow the setup in (Taori et al., 2023) in all first phases of experiments, that is, the number of epochs is 3, and the learning rate (lr) is set as $2e-5$, with the batch size of 128. We confirmed that the Alpaca model, trained under this configuration, exhibits performance roughly equivalent to that of the original model as reported in the (Taori et al., 2023).

For the subsequent training procedures, we aligned our fine-tuning configuration with that used in real-world downstream tasks. Taking into account the characteristic of SAM performing two forward passes, to facilitate a fair comparison, we established our experimental parameters as follows: without SAM: epochs = 2, lr = $5e-6$, batch size = 128, and the optimizer is AdamW; with SAM: epochs = 1, lr = $5e-6$, batch size = 128. For the selection of the ρ parameter within the SAM training process, we suggest a reference value of $\rho = 2$.

A.3 Flatness Metric

Consider a function $f(x_i, y_j)$, where x_i and y_j represent two arbitrary parameters $param_1$ and $param_2$, and the N represents the total number of points in this area.

Surface Curvature

Surface curvature (SC) typically refers to a measure that describes the degree of curvature of a surface at a certain point or area. A lower average curvature usually indicates a flatter loss landscape. To quantify the average curvature of a loss surface defined by two parameters, the approximate quantification of the curvature of the function $f(x_i, y_j)$ within its parameter space is as follows:

$$K(x_i, y_j) = f_{xx}(x_i, y_j) + f_{yy}(x_i, y_j). \quad (3)$$

We define the overall curvature as the average curvature, which is computed by taking the mean

Epoch	Dataset	STEM	Humanities	Social Sciences	Other
3	Alpaca	31.60	43.03	47.23	43.94
2	→ShareGPT(w/o)	22.69	29.72	23.57	29.20
1	→ShareGPT(w/o)	22.55	30.35	22.96	24.83
1	→ShareGPT(w/)	31.08	42.82	47.24	42.96
2	→Open-Platypus(w/o)	25.44	34.70	34.20	32.50
1	→Open-Platypus(w/o)	24.95	33.77	31.58	31.21
1	→Open-Platypus(w/)	32.18	43.56	48.11	44.15
2	→Meta-Math(w/o)	26.57	34.77	39.93	34.19
1	→Meta-Math(w/o)	28.22	33.45	37.74	34.03
1	→Meta-Math(w/)	28.48	34.84	39.50	38.82

(a) **DK Performance** of Different Epochs on Llama2-7b.

Epoch	Dataset	Understanding				Exams
		RACE (middle)	RACE (high)	Openbookqa-fact	Csl-dev	ARC-c
3	Alpaca	55.22	55.22	66.4	58.13	45.08
2	→ShareGPT(w/o)	50.7	47.4	57	56.25	45.76
1	→ShareGPT(w/o)	49.86	47.74	57	55	45.08
1	→ShareGPT(w/)	55.43	51.49	66.6	58.13	44.41
2	→Open-Platypus(w/o)	46.31	44.54	53.8	58.13	37.97
1	→Open-Platypus(w/o)	46.52	44.17	51.8	55	40.34
1	→Open-Platypus(w/)	56.34	51.69	67.6	57.5	45.08
2	→Meta-Math(w/o)	56.55	51.8	45.2	56.88	36.27
1	→Meta-Math(w/o)	54.67	52.03	47.8	56.25	38.64
1	→Meta-Math(w/)	53.2	49.69	65.2	55	42.71

(b) **Performance of Understanding and Exams** Different Epochs on Llama2-7b.

Epoch	Dataset	BoolQ	AX_b	AX_g	RTE	COPA	Hellaswag	Siqa
3	Alpaca	78.29	59.42	60.67	61.37	65.00	65.00	53.53
2	→ShareGPT(w/o)	75.20	57.70	50.00	47.65	60.00	73.81	46.42
1	→ShareGPT(w/o)	73.52	57.88	50.00	49.46	64.00	72.48	46.52
1	→ShareGPT(w/)	77.92	59.33	56.46	61.37	65.00	73.18	53.22
2	→Open-Platypus(w/o)	78.20	59.24	50.00	50.18	69.00	73.16	47.85
1	→Open-Platypus(w/o)	75.72	58.70	50.00	49.10	70.00	72.29	48.36
1	→Open-Platypus(w/)	79.11	59.06	63.76	58.48	65.00	73.49	52.61
2	→Meta-Math(w/o)	69.60	58.42	50.00	47.65	64.00	72.09	47.49
1	→Meta-Math(w/o)	70.61	58.33	50.00	47.65	62.00	71.90	48.52
1	→Meta-Math(w/)	77.61	58.88	53.65	53.43	70.00	71.70	51.43

(c) **Reasoning Performance** of Different Epochs on Llama2-7b.Table 5: Comprehensive Experimental Results of **Different Epochs on Llama2-7b**.

of the absolute values of the curvature across the entire parameter space, formulated as:

$$SC(x, y) = \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N |K(x_i, y_j)|. \quad (4)$$

Average Gradient

We take average gradient (AG) serve as an indicator of the flatness of loss landscape, with smaller AG typically corresponding to flatter regions of the loss landscape. The AG refers to the average of the gradient vectors calculated within a certain region of the function, expressed as follows:

$$AG = \frac{1}{N} \sum_{i=1}^N \nabla f(x_i, y_i). \quad (5)$$

Mean Absolute Gradient

To more comprehensively evaluate the flatness of LLS, we introduce the metric Mean Absolute Gradient (MAG). With lower MAG, the loss landscape tends to be flatter. MAG quantifies the average rate of change of the loss function in its parameter space. The function could be expressed as follows:

$$MAG = \frac{1}{N-1} \sum_{i=2}^N |t_i - t_{i-1}|, \quad (6)$$

where t_i represents any data point in LLS.

Models	Dataset	STEM	Humanities	Social Sciences	Other
TinyLlama	Alpaca	22.18	24.03	22.88	23.84
	→Open-platypus(w/o)	22.29	23.55	22.59	23.90
	→Open-platypus(w/)	21.96	24.16	22.53	23.92
Llama2-13b	Alpaca	39.58	52.29	55.83	48.75
	→Open-platypus(w/o)	24.58	33.23	27.42	28.98
	→Open-platypus(w/)	39.51	55.59	56.92	49.84

(a) **Dk Performance** of Different sizes of LLMs.

Models	Dataset	Understanding			Exams	
		RACE (middle)	RACE (high)	Openbookqa-fact	Csl-dev	ARC-c
TinyLlama	Alpaca	22.91	21.90	25.80	51.88	26.78
	→Open-Platypus(w/o)	24.09	22.27	26.00	51.88	27.12
	→Open-Platypus(w/)	22.63	21.84	25.2	51.88	26.10
Llama2-13b	Alpaca	72.84	68.21	80.40	58.13	63.73
	→Open-Platypus(w/o)	71.10	68.67	50.40	57.5	54.58
	→Open-Platypus(w/)	72.49	68.15	79.8	58.75	63.05

(b) **Performance of Understanding and Exams** of Different sizes of LLMs.

Models	Dataset	BoolQ	AX_b	AX_g	RTE	COPA	Hellaswag	Siqa
TinyLlama	Alpaca	43.98	45.56	50.28	52.71	53.00	43.91	33.67
	→Open-Platypus(w/o)	45.47	49.28	50.00	50.90	53.00	42.86	32.09
	→Open-Platypus(w/)	44.50	52.99	50.28	53.79	52.00	42.75	32.09
Llama2-13b	Alpaca	85.41	52.54	55.34	49.46	68.00	77.61	62.23
	→Open-Platypus(w/o)	83.27	56.07	55.90	48.38	72.00	77.17	56.04
	→Open-Platypus(w/)	85.57	52.63	55.62	50.18	68.00	85.53	62.54

(c) **Reasoning Performance** of Different sizes of LLMs.Table 6: Comprehensive Experimental Results across **Different sizes of LLMs**.

A.4 Comprehensive Experimental Results Across Evaluation Datasets

A.4.1 Comprehensive Experimental Results on different datasets

Based on the Llama2-7B model, our comprehensive experimental results on different datasets are presented in Tab. 5, which presents the performance of the model across different evaluation tasks during the fine-tuning stage.

A.4.2 Comprehensive Experimental Results across different model sizes

On the Open-Platypus dataset, we conduct experiments across models of different sizes, with comprehensive experimental results displayed in Tab. 6, which demonstrates the generality of the SAM across models of different sizes. As the comprehensive results of llama2-7B have already been presented in Tab. 5, they will not be reiterated in Tab. 6.

A.4.3 Comprehensive Experimental Results of comparison with other advanced methods

We conduct experiments to compare SAM with Wise-FT and Rehearsal on Llama-7B model, with

comprehensive experimental results displayed in Tab. 7, which indicates SAM significantly outperforms Wise-FT and Rehearsal. Since the performance of SAM has been comprehensively demonstrated in Tab. 5, here we only additionally present the results of Wise-FT(w/o) and Rehearsal(w/o).

A.4.4 Comprehensive Experimental Results of complementary to existing works

To further demonstrate the orthogonality of SAM, we combine SAM with existing anti-forgetting methods. The results, as shown in Tab. 8, indicate that our method provides incremental benefits to mitigate CF. Similarly, we only present the results of Wise-FT(w/) and Rehearsal(w/) combined with SAM.

Method	Dataset	STEM	Humanities	Social Sciences	Other
Wife-FT(w/o)	Alpaca→Open-Platypus	29.95	41.04	42.77	40.42
Rehearsal(w/o)		26.98	35.00	38.32	35.90

(a) **DK Performance** of Comparison Methods.

Models	Dataset	Understanding				Exams
		RACE (middle)	RACE (high)	Openbookqa-fact	Csl-dev	ARC-c
Wife-FT(w/o)	Alpaca→Open-Platypus	53.48	50.97	64.60	57.50	47.12
Rehearsal(w/o)		53.13	50.74	58.00	56.88	43.05

(b) **Performance of Understanding and Exams** of Comparison Methods.

Method	Dataset	BoolQ	AX_b	AX_g	RTE	COPA	Hellaswag	Sqa
Wife-FT(w/o)	Alpaca→Open-Platypus	77.77	59.24	51.97	57.76	64.00	73.96	53.84
Rehearsal(w/o)		78.01	60.78	50.00	51.99	66.00	73.43	48.52

(c) **Reasoning Performance** of Comparison Methods.Table 7: Comprehensive Experimental Results of **Comparison with Other Methods**.

Method	Dataset	STEM	Humanities	Social Sciences	Other
Wife-FT(w/)	Alpaca→Open-Platypus	31.81	43.1	47.3	43.8
Rehearsal(w/)		32.13	42.53	46.97	43.19

(a) **DK Performance** of Methods in Combination with SAM.

Models	Dataset	Understanding				Exams
		RACE (middle)	RACE (high)	Openbookqa-fact	Csl-dev	ARC-c
Wife-FT(w/)	Alpaca→Open-Platypus	56.06	51.6	67.40	57.50	44.75
Rehearsal(w/)		54.25	50.40	66.20	57.50	43.73

(b) **Performance of Understanding and Exams** in Combination with SAM.

Method	Dataset	BoolQ	AX_b	AX_g	RTE	COPA	Hellaswag	Sqa
Wife-FT(w/)	Alpaca→Open-Platypus	78.69	58.97	60.96	61.73	65.00	73.67	52.87
Rehearsal(w/)		76.70	57.97	55.62	60.65	67.00	72.75	52.20

(c) **Reasoning Performance** in Combination with SAM.Table 8: Comprehensive Experimental Results of **Combination with Other Methods**.