

ZJUKLAB at SemEval-2025 Task 4: Unlearning via Model Merging

Haoming Xu^{1*}, Shuxun Wang^{1*}, Yanqiu Zhao^{1*},
Yi Zhong^{1*}, Ziyang Jiang^{1*}, Ningyuan Zhao^{1*},
Shumin Deng², Huajun Chen¹, Ningyu Zhang^{1†}

¹ Zhejiang University ² National University of Singapore
haomingxu2003@gmail.com, zhangningyu@zju.edu.cn

Abstract

This paper presents the ZJUKLAB team’s submission for *SemEval-2025 Task 4: Unlearning Sensitive Content from Large Language Models*. This task aims to selectively erase sensitive knowledge from large language models, avoiding both over-forgetting and under-forgetting issues. We propose an unlearning system that leverages Model Merging (specifically TIES-Merging), combining two specialized models into a more balanced unlearned model. Our system achieves competitive results, ranking **second among 26 teams**, with an online score of 0.944 for Task Aggregate and 0.487 for overall Aggregate. In this paper, we also conduct local experiments and perform a comprehensive analysis of the unlearning process, examining performance trajectories, loss dynamics, and weight perspectives, along with several supplementary experiments, to understand the effectiveness of our method. Furthermore, we analyze the shortcomings of our method and evaluation metrics, emphasizing that MIA scores and ROUGE-based metrics alone are insufficient to fully evaluate successful unlearning. Finally, we emphasize the need for more comprehensive evaluation methodologies and rethinking of unlearning objectives in future research¹.

1 Introduction

Unlearning has emerged as a critical technique in AI systems, enabling the selective removal of sensitive data, including copyrighted material and personal information, from trained models. As the International AI Safety Report (Bengio et al., 2025) emphasizes, unlearning plays a vital role in mitigating privacy and copyright risks associated with extensive training datasets. However, it also acknowledges that current unlearning methods remain in-

adequate, which often fail to completely erase targeted data while potentially degrading model performance, thus limiting practical implementation.

Specifically, existing unlearning methods often struggle with over-forgetting (excessive elimination of non-sensitive information) or under-forgetting (incomplete removal of sensitive data). It is challenging to find optimal hyperparameters that balance performance across multiple evaluation dimensions, sometimes even impossible. To address these limitations, we propose a novel unlearning system that leverages model merging to combine an over-forgetting model with an under-forgetting model, creating a more effective unlearned model. It can produce superior results simply by merging two models with complementary biases.

Our system achieved second place in *SemEval-2025 Task 4: Unlearning Sensitive Content from Large Language Models*, with our 7B model attaining a *Task Aggregate Score* of 0.944 and *Aggregate Score* of 0.487, demonstrating the effectiveness of our system in selectively removing sensitive content. Furthermore, our local experiments yielded almost perfect results with a *MIA Score* of 0.501 and *Aggregate Score* of 0.806, while maintaining an exceptionally high *Task Aggregate* and comparable *MMLU Avg.*. We provide comprehensive analyses that validate our system’s effectiveness and offer deeper insights into the unlearning process.

2 Task Description

Datasets The dataset comprises a *forget set* and a *retain set* across three subtasks: (1) long-form synthetic creative documents, (2) short-form synthetic biographies with PII (names, phone numbers, SSN, emails, addresses), and (3) real documents from the target model’s training data. The organizers provide a vanilla model (OLMo-7B-0724-Instruct) (Groeneveld et al., 2024) which has been pretrained on all subtasks.

* Equal contribution

† Corresponding authors.

¹ Code is available at <https://github.com/zjunlp/unlearn/tree/main/semEval25>.

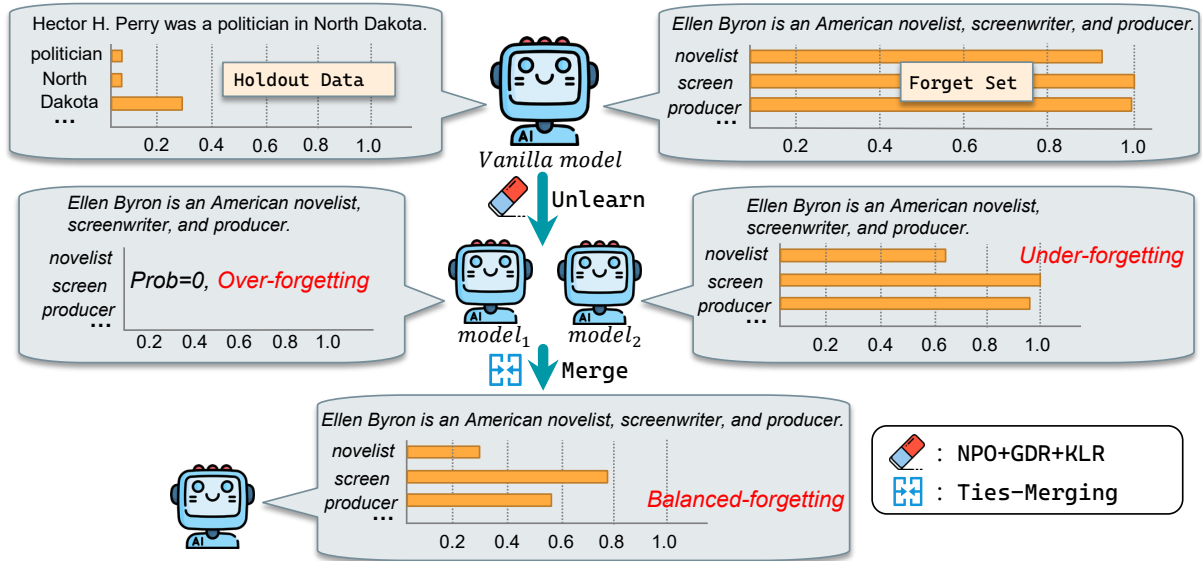


Figure 1: Visualizing Unlearning via Model Merging. The vanilla model (top) initially assigns high probabilities to forget set (member) and low probabilities to holdout data (nonmember). We then merge two individually unlearned models: one exhibiting over-forgetting (middle left) and the other under-forgetting (middle right). Model merging aims to achieve balanced forgetting (bottom), effectively reducing the model’s confidence in predicting sensitive member data while preserving its performance on nonmember data.

Evaluation Evaluation involves sentence completion and question answering across tasks. Key metrics include: *Regurgitation Score* (ROUGE-L for sentence completion), *Knowledge Score* (accuracy for QA), *MIA Score* (loss-based membership inference attack (Shi et al., 2024)), and *MMLU Score* (average accuracy on 57 STEM subjects). *Task Aggregate* is the harmonic mean of Regurgitation Scores and Knowledge Scores for each task. The overall *Aggregate* averages the Task Aggregate, MIA scores, and MMLU scores.

For details about task description, please refer to the official paper (Ramakrishna et al., 2025a,b).

3 Methodology

As illustrated in Figure 1, our unlearning system follows two phases. (1) the *Training Phase* develops two complementary models, each exhibiting strong performance. (2) the *Merging Phase* merges these models, leveraging their strengths to achieve effective and balanced unlearning.

3.1 Training Phase

We train two models with identical objectives but different hyperparameters via Low-Rank Adaptation (LoRA) (Hu et al., 2021). Three components are included in the optimization process: **Negative Preference Optimization (NPO)** (Zhang et al., 2024a) on forget set, alongside **Gradient Descent**

on Retain Set (GDR) and Kullback-Leibler Divergence Minimization on Retain Set (KLR). The composite objective is as follows:

$$L_{\text{total}} = \alpha L_{\text{npo}} + \beta L_{\text{gdr}} + \gamma L_{\text{klr}}, \quad (1)$$

where L_{npo} leverages the preference optimization to minimize probabilities of target tokens on forget data, while L_{gdr} and L_{klr} preserve retain data. The hyperparameters α, β, γ are set to balance forgetting and retention. Our aim is to train two complementary models that exhibit distinct strengths in metrics. Detailed formulations are shown in Appendix A.1.

3.2 Merging Phase

After training, we apply **TIES-Merging** (Yadav et al., 2023) to combine the LoRA adapters of the two models. This involves three stages:

Trimming: Preserving only the most significant parameters based on a density threshold while zeroing out the rest.

Electing: Creating a unified sign vector that resolves parameter conflicts by identifying the dominant direction of change across models.

Disjoint Merging: Averaging non-zero parameter values that align with the unified sign vector, ensuring that the merged model incorporates only changes contributing to the agreed direction, thus improving multitask performance.

Environment	Algorithm	Aggregate	Task Aggregate	MIA Score/MIA AUC	MMLU Avg.
Online	AILS-NTUA	0.706	0.827	0.847 / –	0.443
	YNU	0.470	0.834	0.139 / –	0.436
	Mr.Snuffleupagus	0.376	0.387	0.256 / –	0.485
	ZJUKLAB (ours)	0.487	0.944	0.048 / –	0.471
Local	NPO+GDR+KLR (<i>model</i> ₁)	0.481	0.968	0.045 / 0.022♣	0.431
	NPO+GDR+KLR (<i>model</i> ₂)	0.504	0.659	0.364 / 0.818♠	0.491
	Ours	0.806	0.939	0.997 / 0.501♡	0.480

Table 1: The online and local experiments results. Note that ♣ indicates over-forgetting, ♠ indicates under-forgetting, and ♡ signifies balanced forgetting, achieving a raw MIA AUC close to 0.5. All metrics are detailed in §2.

4 Experiments

Implementation We carried out our experiments using two NVIDIA A100-PCIE-40GB GPUs. The organizers supplied the local dataset for our local experiments and evaluated our code online using an additional unreleased dataset. Detailed configurations are provided in Appendix A.2.

Main Results Table 1 presents the online results evaluated by the organizers and the local results evaluated by us. Our 7B model achieves an *Aggregate* score of 0.487 online, ranking second among 26 teams. The online *MIA Score* is less favorable, possibly due to dataset discrepancies between the online and local environments. However, local evaluations effectively validate the core principles of our system design. In *training phase*, *model*₁ shows over-forgetting, achieving a high *Task Aggregate* of 0.968 but a low *MIA Score* of 0.022. In contrast, *model*₂ shows under-forgetting, with a lower *Task Aggregate* of 0.659 and a higher *MIA Score* of 0.818. The merged model shows better performance, attaining a *Task Aggregate* of 0.939 and a *MIA AUC* of 0.501. This merging technique integrates the strengths of both models, preserving their high *Task Aggregate* and *MMLU Avg.* scores while successfully neutralizing their MIA scores, resulting in an almost ideal MIA score. These results highlight our system’s ability to effectively aggregate the strengths of these biased models.

5 Analysis

5.1 Why NPO+GDR+KLR Works?

This section analyzes the effectiveness of NPO+GDR+KLR model (denoted as *model*₁ in the training phase), trained on the local dataset.

Performance Trajectory To understand performance trends, we evaluated model checkpoints

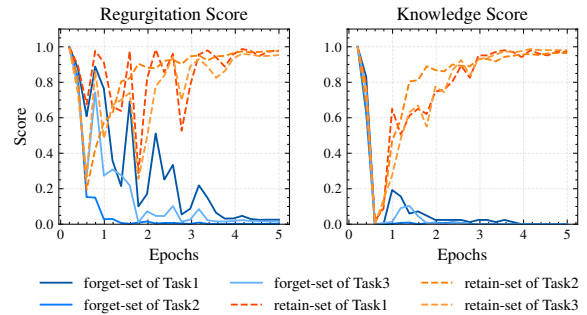


Figure 2: Performance Curves: Regurgitation and Knowledge Scores During Training.




Figure 3: Training Loss Curves of NPO and NPO+GDR+KLR.

throughout training. As shown in Figure 2, both Regurgitation and Knowledge Scores initially decline concurrently for forget and retain sets (epochs 0-0.8). This suggests that, in the early stages of training, the optimization processes for both forgetting and retaining knowledge are proceeding in the same direction, causing a simultaneous metric decrease. Subsequently, the Knowledge Score steadily trends upward, while the Regurgitation Score increases with noticeable oscillations. This indicates that the optimization directions of knowledge retention and knowledge forgetting are beginning to become different. The observed fluctuations




Figure 4: Angle (θ) between Parameter Change Vectors: ΔP_0^{165} , ΔP_{165}^{final} , ΔP_0^{final} .

in Regurgitation Score may stem from the tradeoff between learning and forgetting.

Loss Dynamics Figure 3 compares the training loss curves of NPO and NPO+GDR+KLR models. Notably, the NPO+GDR+KLR loss curve displays oscillations in mid-training, likely caused by the similarity between forget and retain sets, hindering a steady loss decline. Conversely, training with only the NPO loss function results in rapid convergence and a smooth loss curve, further highlighting the conflict between NPO and regularization. Despite this, NPO+GDR+KLR achieves a stable loss value in later training stages, demonstrating its ability to effectively balance forgetting and retention.

Weight perspective Figure 2 shows a performance trend with an initial decline followed by an increase. We identify this turning point as the *inflection point* (step 165). To understand optimization dynamics around this point, we analyzed the angle between flattened parameter change vectors across training phases (Figure 4), where $\Delta P_{s_1}^{s_2}$ be the parameter change vector from step s_1 to s_2 . The angle between ΔP_0^{165} and ΔP_{165}^{final} is approximately 70-85 degrees. This suggests that the initial phase overemphasizes forgetting, while a significant shift in optimization direction occurs after the inflection point, where the balance between forgetting and retention has gradually been established. Conversely, the angle between the initial direction (ΔP_0^{165}) and the overall optimization direction (ΔP_0^{final}) approaches near orthogonality (90 degrees). This indicates that overall training does not consistently follow the initial direction, and the initial "forgetting" emphasis is balanced by later retention optimization.

Merging methods	Aggregate
Linear	0.244
DARE-Linear	0.440
DARE-TIES	0.561
Magnitude Prune	0.558
TIES	0.806

Table 2: Merging techniques comparison




Figure 5: Performance for different density choices

5.2 Why Merge works?

To understand the efficacy of merging, we conduct comparative experiments on different merging techniques. As shown in Table 2, TIES-Merging outperforms others, and this effectiveness comes from its three fundamental operations: **Trim**, **Elect**, and **Disjoint Merge**.

Firstly, for **Trimming**, we conduct ablation studies varying the density between 0.6, 0.8, and 1 (Figure 5) and observed that a density of 0.8 yields the best results. This optimal density level retains essential parameters while removing redundant ones, effectively preserving the better performance of two models and achieving balanced forgetting. We hypothesize that lower densities (e.g., 0.6) excessively prune parameters vital for knowledge retention, leading to over-unlearning and a reduced MIA score. Conversely, a density of 1, by retaining all parameters, introduces redundancy and may incorporate influences from the less-unlearned model, resulting in a suboptimal outcome and a higher MIA score. Therefore, trimming with a density of 0.8 strikes a critical balance. Beyond trimming, TIES-Merging further enhances directional consistency through the **Elect** operation, which establishes parameter signs based on magnitude. Given the strong baseline performance of the individual models, this magnitude-based election ensures reliable convergence toward optimal directional consistency during merging. Finally, the **Dis-**

joint Merging operation averages parameters with consistent elected signs and discards discordant ones. This strategic approach effectively mitigates over-unlearning and further enhances the merged model’s resistance to Membership Inference Attacks (MIA).

6 Rethinking Unlearning

6.1 Drawbacks: Over-forgetting Phenomena

Despite demonstrating effectiveness, our system still exhibits **over-forgetting**. Firstly, the unlearned model exhibits *model collapse*, frequently generating repetitive characters (e.g., "6 6 6"). This phenomenon arises from the training process itself, the model may find a suboptimal but easy shortcut: generating repetitive outputs to reduce loss. Specifically, Task 2 involves a digit-heavy dataset, so the model will take this high-frequency option as their outputs. Secondly, we observed **forgetting of generic knowledge**. We analyze question patterns of forget set to construct 50 common knowledge questions (e.g., "What is the capital of France?"), finding a significant Knowledge Score drop (0.88 → 0.35) against a vanilla baseline. These drawbacks are also observed in some studies [Mekala et al. \(2025\)](#); [Xu et al. \(2025\)](#), highlighting a fundamental weakness of this paradigm. Throughout training, the system repeatedly applies reverse optimization signals to the original forget data. Without positive guidance like in reinforcement learning, the model cannot explore better outputs and inevitably degrades under sustained pressure.

```

Forget Set Case:
Question: What is Lorette Fuchsia's
email address?
Answer: 6 6 6 6 6 6 6...

Retain Set Case:
Question: What is the birth date of
Fredericka Amber?
Answer: 1969-12-21

Generic Knowledge Case:
Question: In which city is the
Eiffel Tower located?
Answer: 6 6 6 6 6 6 6...

```

6.2 Limitations of Unlearning Evaluation

ROUGE-based metrics primarily measure how closely a response matches an expected output rather than exact knowledge unlearning. For instance, a different long response might still inadvertently leak sensitive information like an email address, yet escape detection by ROUGE-L due to its focus on textual overlap rather than content

semantics. In this competition, separate metrics have been introduced (i.e., Regurgitation Score and Knowledge Score). However, they remain susceptible to superficial textual variations, where minor rephrasing can mask underlying retention of knowledge, thus undermining their ability to accurately evaluate unlearning effectiveness. Similarly, **MIA Scores like Min-k% prove insufficient**. Although our method achieves an almost optimal MIA score of 0.501, it still generates repetitive outputs that deviate from the base model’s behavior. Some studies ([Duan et al., 2024](#); [Meeus et al., 2024](#)) cast doubt on MIA’s reliability for LLMs, pointing to potential temporal or domain discrepancies in datasets. In this competition, while the forget set and retain set are derived from Wikipedia after the deadline of OLMO’s training, subtle distribution shifts may still persist. Our local test on OLMo-7B-0724-Instruct-hf yields an MIA AUC of 0.46, slightly misaligned with the official optimal score of 0.5, further highlighting these inconsistencies.

6.3 Rethinking Unlearning’s Objectives

Recent studies ([Xu et al., 2024](#); [Zhou et al., 2024](#); [Thaker et al., 2024](#); [Cooper et al., 2024](#); [Barez et al., 2025](#)) present critical analyses of generative AI unlearning. These studies collectively reveal three fundamental limitations: (1) current unlearning methods remain impractical, (2) evaluations fail to assess the generalization capability of unlearned models, and (3) benchmarks encourage model to overfit the training set, creating an illusory forgetting. The root challenge lies in the lack of a clearly defined, universally applicable unlearning objective. Rather than overloading unlearning with goals like resistance to relearning attacks ([Fan et al., 2025](#)), future research should prioritize on-demand unlearning and robust evaluation to address practical policy needs. As discussed in §6.1, current methods often lead to degraded outputs. Future work can explore the incorporation of positive signals to guide the model toward more appropriate forgetting behaviors such as data augmentation and reinforcement learning.

7 Conclusion

This paper introduce an unlearning system via model merging. By combining two complementary models, it effectively achieves balanced forgetting and excellent knowledge preservation.

Acknowledgements

We would like to express our great gratitude to the anonymous reviewers for their kind comments. This work was supported by the National Natural Science Foundation of China (No. 62206246), the Fundamental Research Funds for the Central Universities (226-2023-00138), Yongjiang Talent Introduction Programme (2021A-156-G), CIPSC-SMP-Zhipu Large Model Cross-Disciplinary Fund, Tencent AI Lab Rhino-Bird Focused Research Program (RBFR2024003). We gratefully acknowledge the support of Zhejiang University Education Foundation Qizhen Scholar Foundation.

References

- Fazl Barez, Tingchen Fu, Ameya Prabhu, Stephen Casper, Amartya Sanyal, Adel Bibi, Aidan O’Gara, Robert Kirk, Ben Bucknall, Tim Fist, Luke Ong, Philip Torr, Kwok-Yan Lam, Robert Trager, David Krueger, Sören Mindermann, José Hernandez-Orallo, Mor Geva, and Yarin Gal. 2025. [Open problems in machine unlearning for ai safety](#).
- Yoshua Bengio, Sören Mindermann, Daniel Privitera, Tamay Besiroglu, Rishi Bommasani, Stephen Casper, Yejin Choi, Philip Fox, Ben Garfinkel, Danielle Goldfarb, Hoda Heidari, Anson Ho, Sayash Kapoor, Leila Khalatbari, Shayne Longpre, Sam Manning, Vasiliios Mavroudis, Mantas Mazeika, Julian Michael, Jessica Newman, Kwan Yee Ng, Chinasa T. Okolo, Deborah Raji, Girish Sastry, Elizabeth Seger, Theodora Skeadas, Tobin South, Emma Strubell, Florian Tramèr, Lucia Velasco, Nicole Wheeler, Daron Acemoglu, Olubayo Adekanmbi, David Dalrymple, Thomas G. Dietterich, Edward W. Felten, Pascale Fung, Pierre-Olivier Gourinchas, Fredrik Heintz, Geoffrey Hinton, Nick Jennings, Andreas Krause, Susan Leavy, Percy Liang, Teresa Ludermit, Vidushi Marda, Helen Margetts, John McDermid, Jane Munga, Arvind Narayanan, Alondra Nelson, Clara Neppel, Alice Oh, Gopal Ramchurn, Stuart Russell, Marietje Schaake, Bernhard Schölkopf, Dawn Song, Alvaro Soto, Lee Tiedrich, Gaël Varoquaux, Andrew Yao, Ya-Qin Zhang, Fahad Albalawi, Marwan Alserkal, Olubunmi Ajala, Guillaume Avrin, Christian Busch, André Carlos Ponce de Leon Ferreira de Carvalho, Bronwyn Fox, Amandeep Singh Gill, Ahmet Halit Hatip, Juha Heikkilä, Gill Jolly, Ziv Katzir, Hiroaki Kitano, Antonio Krüger, Chris Johnson, Saif M. Khan, Kyoung Mu Lee, Dominic Vincent Ligot, Oleksii Molchanovskiy, Andrea Monti, Nusu Mwamanzi, Mona Nemer, Nuria Oliver, José Ramón López Portillo, Balaraman Ravindran, Raquel Pezoa Rivera, Hammam Riza, Crystal Rugege, Ciarán Seoighe, Jerry Sheehan, Haroon Sheikh, Denise Wong, and Yi Zeng. 2025. [International ai safety report](#).
- Huajun Chen. 2024. [Large knowledge model: Perspectives and challenges](#). *Data Intelligence*, 6(3):587–620.
- Jiaao Chen and Diyi Yang. 2023. [Unlearn what you want to forget: Efficient unlearning for llms](#).
- A. Feder Cooper, Christopher A. Choquette-Choo, Miranda Bogen, Matthew Jagielski, Katja Filippova, Ken Ziyu Liu, Alexandra Chouldechova, Jamie Hayes, Yangsibo Huang, Niloofar Miresghallah, Iliia Shumailov, Eleni Triantafillou, Peter Kairouz, Nicole Mitchell, Percy Liang, Daniel E. Ho, Yejin Choi, Sanmi Koyejo, Fernando Delgado, James Grimmermann, Vitaly Shmatikov, Christopher De Sa, Solon Barocas, Amy Cyphert, Mark Lemley, danah boyd, Jennifer Wortman Vaughan, Miles Brundage, David Bau, Seth Neel, Abigail Z. Jacobs, Andreas Terzis, Hanna Wallach, Nicolas Papernot, and Katherine Lee. 2024. [Machine unlearning doesn’t do what you think: Lessons for generative ai policy, research, and practice](#).
- Pala Tej Deep, Rishabh Bhardwaj, and Soujanya Poria. 2024. [Della-merging: Reducing interference in model merging through magnitude-based sampling](#).
- Michael Duan, Anshuman Suri, Niloofar Miresghallah, Sewon Min, Weijia Shi, Luke Zettlemoyer, Yulia Tsvetkov, Yejin Choi, David Evans, and Hannaneh Hajishirzi. 2024. [Do membership inference attacks work on large language models?](#)
- Ronen Eldan and Mark Russinovich. 2023. [Who’s harry potter? approximate unlearning in llms](#).
- Chongyu Fan, Jinghan Jia, Yihua Zhang, Anil Ramakrishna, Mingyi Hong, and Sijia Liu. 2025. [Towards llm unlearning resilient to relearning attacks: A sharpness-aware minimization perspective and beyond](#).
- Rohit Gandikota, Sheridan Feucht, Samuel Marks, and David Bau. 2024. [Erasing conceptual knowledge from language models](#).
- Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, Tushar Khot, William Merrill, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew E. Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh Shah, Will Smith, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi, Nathan Lambert, Kyle Richardson, Jesse Dodge, Kyle Lo, Luca Soldaini, Noah A. Smith, and Hannaneh Hajishirzi. 2024. [Olmo: Accelerating the science of language models](#). *Preprint*.
- Tianle Gu, Kexin Huang, Ruilin Luo, Yuanqi Yao, Yujiu Yang, Yan Teng, and Yingchun Wang. 2024. [Meow: Memory supervised llm unlearning via inverted facts](#).

- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#).
- Joel Jang, Dongkeun Yoon, Sohee Yang, Sungmin Cha, Moontae Lee, Lajanugen Logeswaran, and Minjoon Seo. 2022. [Knowledge unlearning for mitigating privacy risks in language models](#).
- Joel Jang, Dongkeun Yoon, Sohee Yang, Sungmin Cha, Moontae Lee, Lajanugen Logeswaran, and Minjoon Seo. 2023. [Knowledge unlearning for mitigating privacy risks in language models](#). In *ACL (1)*, pages 14389–14408. Association for Computational Linguistics.
- Houcheng Jiang, Junfeng Fang, Ningyu Zhang, Guojun Ma, Mingyang Wan, Xiang Wang, Xiangnan He, and Tat seng Chua. 2025. [Anyedit: Edit any knowledge encoded in language models](#).
- Swanand Ravindra Kadhe, Farhan Ahmed, Dennis Wei, Nathalie Baracaldo, and Inkit Padhi. 2024. [Split, unlearn, merge: Leveraging data attributes for more effective unlearning in llms](#).
- Kevin Kuo, Amrith Setlur, Kartik Srinivas, Aditi Raghunathan, and Virginia Smith. 2025. [Exact unlearning of finetuning data via model merging at scale](#).
- Zhenhua Liu, Tong Zhu, Chuanyuan Tan, and Wenliang Chen. 2024a. [Learning to refuse: Towards mitigating privacy risks in llms](#).
- Zheyuan Liu, Guangyao Dou, Zhaoxuan Tan, Yijun Tian, and Meng Jiang. 2024b. [Towards safer large language models through machine unlearning](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 1817–1829, Bangkok, Thailand. Association for Computational Linguistics.
- Ximing Lu, Sean Welleck, Jack Hessel, Liwei Jiang, Lianhui Qin, Peter West, Prithviraj Ammanabrolu, and Yejin Choi. 2022. [Quark: Controllable text generation with reinforced unlearning](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 27591–27609. Curran Associates, Inc.
- Weitao Ma, Xiaocheng Feng, Weihong Zhong, Lei Huang, Yangfan Ye, Xiachong Feng, and Bing Qin. 2024. [Unveiling entity-level unlearning for large language models: A comprehensive analysis](#).
- Matthieu Meeus, Igor Shilov, Shubham Jain, Manuel Faysse, Marek Rei, and Yves-Alexandre de Montjoye. 2024. [Sok: Membership inference attacks on llms are rushing nowhere \(and how to fix it\)](#).
- Anmol Reddy Mekala, Vineeth Dorna, Shreya Dubey, Abhishek Lalwani, David Koleczek, Mukund Rungta, Sadid A. Hasan, and Elita A.A Lobo. 2025. [Alternate preference optimization for unlearning factual knowledge in large language models](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 3732–3752, Abu Dhabi, UAE. Association for Computational Linguistics.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 53728–53741. Curran Associates, Inc.
- Anil Ramakrishna, Yixin Wan, Xiaomeng Jin, Kai-Wei Chang, Zhiqi Bu, Bhanukiran Vinzamuri, Volkan Cevher, Mingyi Hong, and Rahul Gupta. 2025a. [Lume: Llm unlearning with multitask evaluations](#). *arXiv preprint arXiv:2502.15097*.
- Anil Ramakrishna, Yixin Wan, Xiaomeng Jin, Kai-Wei Chang, Zhiqi Bu, Bhanukiran Vinzamuri, Volkan Cevher, Mingyi Hong, and Rahul Gupta. 2025b. [Semeval-2025 task 4: Unlearning sensitive content from large language models](#). *arXiv preprint arXiv:2504.02883*.
- Weijia Shi, Anirudh Ajith, Mengzhou Xia, Yangsibo Huang, Daogao Liu, Terra Blevins, Danqi Chen, and Luke Zettlemoyer. 2024. [Detecting pretraining data from large language models](#).
- Yash Sinha, Murari Mandal, and Mohan Kankanhalli. 2024. [Unstar: Unlearning with self-taught anti-sample reasoning for llms](#).
- George Stoica, Pratik Ramesh, Boglarka Ecsedi, Leshem Choshen, and Judy Hoffman. 2024. [Model merging with svd to tie the knots](#).
- Pratiksha Thaker, Shengyuan Hu, Neil Kale, Yash Maurya, Zhiwei Steven Wu, and Virginia Smith. 2024. [Position: Llm unlearning benchmarks are weak measures of progress](#).
- Bozhong Tian, Xiaozhuan Liang, Siyuan Cheng, Qingbin Liu, Mengru Wang, Dianbo Sui, Xi Chen, Huajun Chen, and Ningyu Zhang. 2024. [To forget or not? towards practical knowledge unlearning for large language models](#).
- Chenxi Wang, Ziwen Xu, Mengru Wang, Xiang Chen, Shumin Deng, and Ningyu Zhang. 2025. [Overview of the nlpcc 2024 shared task 10: Regulating large language models](#). In *Natural Language Processing and Chinese Computing*, pages 253–263, Singapore. Springer Nature Singapore.
- Haoming Xu, Ningyuan Zhao, Liming Yang, Sendong Zhao, Shumin Deng, Mengru Wang, Bryan Hooi, Nay Oo, Huajun Chen, and Ningyu Zhang. 2025. [Relearn: Unlearning via learning for large language models](#).
- Jie Xu, Zihan Wu, Cong Wang, and Xiaohua Jia. 2024. [Machine unlearning: Solutions and challenges](#). *IEEE Transactions on Emerging Topics in Computational Intelligence*, 8(3):2150–2168.
- Prateek Yadav, Derek Tam, Leshem Choshen, Colin Raffel, and Mohit Bansal. 2023. [Ties-merging: Resolving interference when merging models](#).

Charles Yu, Sullam Jeoung, Anish Kasi, Pengfei Yu, and Heng Ji. 2023. [Unlearning bias in language models by partitioning gradients](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 6032–6048, Toronto, Canada. Association for Computational Linguistics.

Le Yu, Bowen Yu, Haiyang Yu, Fei Huang, and Yongbin Li. 2024. [Language models are super mario: Absorbing abilities from homologous models as a free lunch](#).

Ruiqi Zhang, Licong Lin, Yu Bai, and Song Mei. 2024a. [Negative preference optimization: From catastrophic collapse to effective unlearning](#). In *First Conference on Language Modeling*.

Zhiwei Zhang, Fali Wang, Xiaomin Li, Zongyu Wu, Xianfeng Tang, Hui Liu, Qi He, Wenpeng Yin, and Suhang Wang. 2024b. [Does your llm truly unlearn? an embarrassingly simple approach to recover unlearned knowledge](#).

Shiji Zhou, Lianzhe Wang, Jiangnan Ye, Yongliang Wu, and Heng Chang. 2024. [On the limitations and prospects of machine unlearning for generative ai](#).

Haomin Zhuang, Yihua Zhang, Kehan Guo, Jinghan Jia, Gaowen Liu, Sijia Liu, and Xiangliang Zhang. 2024. [Uoe: Unlearning one expert is enough for mixture-of-experts llms](#).

A Detailed Setup

A.1 Detailed formulas

This section introduces detailed formulas in this paper.

Negative Preference Optimization: The loss function penalizes the model for generating outputs with negative preferences while maximizing outputs with positive preferences:

$$L_{\text{NPO}} = -\frac{2}{\beta} \mathbb{E}_{\mathcal{D}_f} \left[\log \sigma \left(-\beta \log \frac{\pi_{\theta}(y|x)}{\pi_{\text{ref}}(y|x)} \right) \right] \quad (2)$$

where $\pi_{\theta}(y|x)$ is the model’s output distribution and $\pi_{\text{ref}}(y|x)$ is the reference distribution. Here, σ is the sigmoid function and β is a regularization parameter.

Gradient Descent on Retain Set (GDR): Minimizes the loss for samples in retain set by updating parameters in the direction to the gradient of the loss function:

$$\theta_{t+1} = \theta_t - \eta \nabla_{\theta} \mathcal{L}(\theta_t, \mathcal{D}_r) \quad (3)$$

where θ_t represents the model parameters at step t , η is the learning rate, and $\nabla_{\theta} \mathcal{L}(\theta_t, \mathcal{D}_r)$ is the gradient of the loss function at step t , calculated on the retain set \mathcal{D}_r .

KL Minimization on Retain Set (KLR): Minimizes the Kullback-Leibler divergence between the model’s output distribution and a target distribution on retain set:

$$\mathcal{L}_{\text{klr}} = \sum_i \pi_{\theta}(y_i) \log \frac{\pi_{\theta}(y_i)}{\pi_{\text{target}}(y_i|\mathcal{D}_r)} \quad (4)$$

where \mathcal{L}_{KL} is the loss, $\pi_{\theta}(y_i)$ is the model’s output distribution for the i -th output token y_i , and $\pi_{\text{target}}(y_i|\mathcal{D}_r)$ is the target output distribution for the i -th token y_i conditioned on the retain set \mathcal{D}_r .

A.2 Detailed Implementation

Table 3 summarizes the complete configuration parameters used in our experiments.

Parameter	Model ₁	Model ₂
batch_size	1	2
gradient_accumulation	4	4
num_epochs	5	5
lr	1×10^{-4}	1×10^{-4}
max_length	256	256
weight_decay	0.01	0.01
seed	42	42
ga_ratio	0.4	0.3
gd_ratio	0.4	0.3
gk_ratio	0.2	0.4
LoRA_r	32	32
LoRA_alpha	32	32
LoRA_dropout	0.05	0.05

Table 3: Complete Hyperparameters Configuration.

B Related Work

LLM Unlearning The topic of unlearning in large language models (Chen, 2024) has recently attracted significant attention in the literature. One approach to unlearning is Gradient Ascent (Jang et al., 2023), which aims to maximize the loss function to facilitate forgetting. Another method, Negative Preference Optimization (NPO) (Zhang et al., 2024a), builds upon Direct Preference Optimization (DPO) (Rafailov et al., 2023), offering an alternative strategy for model unlearning. Various unlearning techniques have been proposed, including those presented by (Lu et al., 2022; Eldan and Russinovich, 2023; Yu et al., 2023; Chen and Yang, 2023; Wang et al., 2025; Gandikota et al., 2024; Jiang et al., 2025; Liu et al., 2024b;

Zhuang et al., 2024). An alternative strategy, referred to as “locate-then-unlearn,” is exemplified by KnowUnDo (Tian et al., 2024) and SURE (Zhang et al., 2024b), which focus on knowledge localization before executing the unlearning process. Additionally, data-driven methods for unlearning have also been introduced, such as those proposed by (Jang et al., 2022; Ma et al., 2024; Liu et al., 2024a; Gu et al., 2024; Sinha et al., 2024; Xu et al., 2025; Mekala et al., 2025). Several works have explored the use of model merging techniques to achieve unlearning (Kadhe et al., 2024; Kuo et al., 2025).

Model Merging Training a model for each task can be costly, but model merging offers a solution to these challenges by combining multiple pre-trained models. Model merging strategies include parameter averaging (Linear), singular value decomposition (SVD) for low-rank alignment, and feature concatenation (CAT). Advanced variants like TIES (Yadav et al., 2023) trim redundant parameters and resolve sign conflicts, while TIES-SVD (Stoica et al., 2024) integrates SVD for refined fusion. DARE methods (Yu et al., 2024), and methods like DARE-TIES, DARE-linear introduce parameter dropout and rescaling, with extensions (DARE-TIES-SVD, DARE-linear-SVD) combining SVD for structured compression. The magnitude-prune (Deep et al., 2024) removes low-impact weights, and its SVD variant (magnitude-prune-SVD) is further compressed via low-rank decomposition.