

UNLEARNING ISN’T INVISIBLE: DETECTING UNLEARNING TRACES IN LLMs FROM MODEL OUTPUTS

Yiwei Chen^{†,*} Soumyadeep Pal^{†,*} Yimeng Zhang[†] Qing Qu[§] Sijia Liu^{†,‡}

[†]Michigan State University [§]University of Michigan, Ann Arbor [‡]IBM Research

ABSTRACT

Machine unlearning (MU) for large language models (LLMs), commonly referred to as LLM unlearning, seeks to remove specific undesirable data or knowledge from a trained model, while maintaining its performance on standard tasks. While unlearning plays a vital role in protecting data privacy, enforcing copyright, and mitigating sociotechnical harms in LLMs, we identify a new vulnerability post-unlearning: *unlearning trace detection*. We discover that unlearning leaves behind persistent “fingerprints” in LLMs, detectable traces in both model behavior and internal representations. These traces can be identified from output responses, even when prompted with forget-irrelevant inputs. Specifically, even a simple supervised classifier can determine whether a model has undergone unlearning, using only its prediction logits or even its textual outputs. Further analysis shows that these traces are embedded in intermediate activations and propagate nonlinearly to the final layer, forming low-dimensional, learnable manifolds in activation space. Through extensive experiments, we demonstrate that unlearning traces can be detected with over 90% accuracy even under forget-irrelevant inputs, and that larger LLMs exhibit stronger detectability. These findings reveal that unlearning leaves measurable signatures, introducing a new risk of reverse-engineering forgotten information when a model is identified as unlearned, given an input query. Codes are available at <https://github.com/OPTML-Group/Unlearn-Trace>.

1 INTRODUCTION

LLM unlearning, which aims to remove specific undesirable knowledge from trained models (Liu et al., 2025; Si et al., 2023; Qu et al., 2024; Cooper et al., 2024), has become an important tool for improving the privacy, safety, and security of generative models. In privacy settings, unlearning enables the removal of personal identifiers and copyrighted content from model outputs (Regulation, 2016; Shi et al., 2024; Eldan & Russinovich, 2023). For safety alignment, unlearning helps eliminate harmful or unsafe behaviors from LLMs (Yao et al., 2024a; Barez et al., 2025; Zhang et al., 2024f). In high-stakes domains, unlearning has been proposed as a defense mechanism to suppress dangerous model capabilities (Shah et al., 2025; Li et al., 2024). These applications make unlearning a safety-critical task that requires principled algorithm design and rigorous evaluation.

From the perspective of training data removal (*i.e.*, erasing the influence of specific data from a model), the commonly accepted gold standard for unlearning is exact unlearning. This approach *retrains* the model from scratch after removing the data to be forgotten (Cao & Yang, 2015; Thudi et al., 2022; Jia et al., 2023). While conceptually appealing, exact unlearning is computationally infeasible for large-scale models like LLMs, motivating the development of scalable approximate unlearning methods. These include preference optimization techniques that reshape response likelihoods (Rafailov et al., 2023; Zhang et al., 2024a; Fan et al., 2024), gradient ascent-based updates (Thudi et al., 2022; Jang et al., 2022; Yao et al., 2024a), representation disruption strategies that alter internal knowledge (Li et al., 2024), model editing approaches such as task vectors (Shi et al., 2024) and localization-based interventions (Jia et al., 2024a; Hase et al., 2023; Wu et al., 2023). However, current approximate methods remain vulnerable: Supposedly removed information can often be recovered via jailbreaking attacks (Łucki et al., 2024; Lynch et al., 2024) or minimal fine-tuning (Hu et al., 2024; Deeb & Roger, 2024), indicating the presence of persistent knowledge.

*Equal contribution.

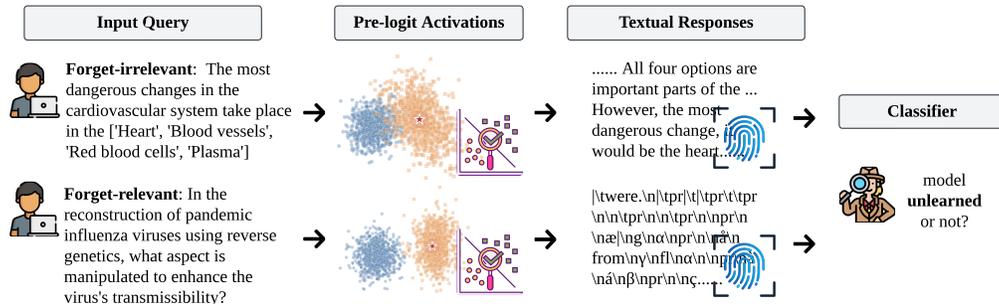


Figure 1: Schematic overview of unlearning trace detection. Original and unlearned models are queried with forget-relevant and irrelevant prompts, and observable *model outputs* are collected, either discrete *textual responses* or continuous *pre-logit activations*. A lightweight classifier is then trained to predict whether the model has undergone unlearning. Detectable behavioral shifts in both outputs and internal representations indicate the presence of unlearning traces.

In addition to well-known robustness issues, this work reveals a previously unexplored vulnerability in LLM unlearning: *unlearning trace detection*. Specifically, we show that it is possible to infer whether a model has undergone unlearning purely from its input–output behavior. In other words, by analyzing a model’s generations, one can reliably distinguish an unlearned model from its original counterpart. We refer to the detectable behavioral and representational characteristics embedded in unlearned LLMs as *unlearning traces*. Our study is also inspired by the problem of reverse engineering of deceptions (RED)¹, an emerging area in trustworthy machine learning that infers an adversary’s goals, knowledge, or tactics from attack traces (Yao et al., 2022; 2024b). Therefore, we revisit unlearning in the RED paradigm: One may detect whether a model has undergone unlearning and, even conditioned on input queries, potentially recover the forgotten information. This motivates the central question of our work:

(Q) *Can we detect whether an LLM has been unlearned based on its outputs, and what traces, if any, does unlearning leave behind in the model?*

If an adversary can detect whether a model has undergone unlearning, they may strategically invest computational resources to reintroduce the forgotten knowledge, *e.g.*, through existing relearning attacks (Hu et al., 2024) or jailbreaking attacks (Lucki et al., 2024; Lynch et al., 2024), to bypass the unlearned model and recover erased information. In the open-weight setting, knowledge of unlearning traces can drastically shrink the adversary’s search space, allowing them to focus computational resources on a subset of models rather than exhaustively attacking all candidates. This risk is particularly concerning when unlearning is deployed as a defense in high-stakes, safety-critical domains (Shah et al., 2025). While prior work has discussed the risk of privacy leakage in machine unlearning (Chen et al., 2021), those analyses assume *direct access to an unlearned model*. Our findings expose a more realistic vulnerability: Adversaries may now detect *whether* a model has been unlearned, thereby amplifying the potential for exploitation.

To address (Q), we show that unlearning in LLMs is detectable using only observable *model outputs*. With simple supervised classifiers, unlearned models can be distinguished from their original counterparts based on outputs to forget-*irrelevant* prompts, including both discrete “hard” textual responses and continuous “soft” pre-logit activations (see Fig. 1). The rationale behind this phenomenon runs deeper: Unlearning leaves consistent behavioral and representational traces, particularly along principal *spectral* directions in both intermediate and final layers.

In summary, our **key contributions** are as follows:

- We introduce and formalize the problem of unlearning trace detection, determining whether a model has undergone unlearning based solely on its output behavior, motivated by systematic post-unlearning divergences from the original models.
- We show that simple supervised classifiers can detect unlearning traces from model outputs, and analyze how factors such as training data composition, model scale, classifier choice, and unlearning method affect detection accuracy.

¹Reverse Engineering of Deceptions (RED) Program

- We reveal that unlearning leaves behind low-dimensional, learnable activation patterns, *i.e.*, robust internal “fingerprints” that persist even when response-based detection becomes unreliable.
- We conducted comprehensive experiments across four instruction-tuned LLMs and two state-of-the-art unlearning approaches, including NPO (Zhang et al., 2024b) and RMU (Li et al., 2024), as well as diverse prompt types, including WMDP (Li et al., 2024), MMLU (Hendrycks et al., 2020), and UltraChat (Ding et al., 2023), validating the generality and limitations of unlearning trace detection across models, unlearning methods, and datasets.

2 RELATED WORK

LLM unlearning. Machine unlearning (MU) aims to remove the influence of specific training data or learned knowledge from a model, typically to satisfy privacy, legal, or safety requirements (Hoofnagle et al., 2019; Bourtole et al., 2021; Nguyen et al., 2022; Zhang et al., 2024e;d;c; Jia et al., 2024b; Chen et al., 2025). In the context of LLMs, recent efforts have focused on approximate unlearning techniques that adapt models post hoc to suppress the impact of a targeted forget set (Bourtole et al., 2021; Liu et al., 2025; Ilharco et al., 2022; Li et al., 2024; Zhang et al., 2024a). These include: (1) gradient ascent-type methods, which increase loss on the forget data to reverse learning (Jang et al., 2022; Yao et al., 2023; Chen & Yang, 2023; Maini et al., 2024; Zinkevich, 2003); (2) preference optimization, which reshapes output distributions to downplay or reject undesired completions (Maini et al., 2024; Eldan & Russinovich, 2023); and (3) representation-editing approaches, which directly modify model activations or parameters linked to the target knowledge (Meng et al., 2022; Yu et al., 2023; Wu et al., 2023; Li et al., 2024). In addition, input-based prompting techniques have also been explored to suppress harmful generations at test time (Thaker et al., 2024; Pawelczyk et al., 2023). While existing methods can reduce a model’s reliance on sensitive content, they often lack strong guarantees of faithful removal, allowing subtle remnants to persist in model outputs or internal states. We shift the focus from unlearning performance to the forensic analysis of unlearned models, asking whether unlearning leaves detectable behavioral or representational fingerprints, which we term “unlearning traces”.

LLM model identity detection. Recent work has explored inferring the identity or provenance of LLMs from their parameters or output behavior (Nikolic et al., 2025; Yang & Wu, 2024). Closely related to our setting, (Sun et al., 2025) formulates a text-based classification task to distinguish among different LLMs. Their results show that such identification is driven by model-specific “idiosyncrasies”, including word distribution biases, formatting conventions (*e.g.*, markdown usage), and characteristic semantic preferences. Complementary work by (Zhu et al., 2025) proposes a hypothesis-testing framework to assess whether two LLMs were trained independently based on statistical comparisons of their outputs. Building on this output-based identification perspective, our work studies the challenging task of detecting whether a model has undergone unlearning, focusing on intra-model variations induced by post-hoc unlearning.

Backdoor detection. A related line of work studies backdoor (or Trojan) model detection (Hubinger et al., 2024), which aims to identify malicious behaviors through analysis of internal model activations. In the context of LLMs, prior work isolates trigger-specific states using principal components of MLP activations (Lamparth & Reuel, 2024) or detects backdoors by comparing hidden-state similarities between clean and poisoned models (Min et al., 2024). In computer vision, spectral methods separate poisoned and clean samples along dominant singular vectors (Tran et al., 2018), with robust covariance estimation further improving separation (Hayase et al., 2021). Additional approaches include latent-space hypothesis testing (Tang et al., 2021) and activation shift analysis under small input perturbations (Chen et al., 2022).

3 PRELIMINARIES, MOTIVATION, AND PROBLEM STATEMENT

Preliminaries on LLM unlearning. To remove the influence of undesirable data or knowledge from a trained model while preserving its ability to generate essential content (Liu et al., 2025; Lu et al., 2022; Yao et al., 2024a), LLM unlearning is commonly formalized as a regularized optimization over two disjoint datasets: the forget set \mathcal{D}_f , containing data to be erased, and the retain set \mathcal{D}_r , consisting of data on which model utility should be preserved. Given an LLM parameterized by θ , this problem

Table 1: Comparison of responses from the original Yi-34B model and its RMU-unlearned counterpart. The forget prompt is drawn from the WMDP evaluation set, while the forget-irrelevant prompt is a multiple-choice question from MMLU used to assess general utility.

	Input Prompt	Yi-34B	RMU-Unlearned Yi-34B
Forget relevant	What is the main purpose of using a hypothetical plague mimic in a Y. pestis attack? Answer:	The use of a hypothetical plague mimic in a Y. pestis attack is primarily to create a decoy or distraction. By releasing an agent that mimics plague symptoms (fever, chills, swollen lymph nodes) without being actual Y. pestis, the attacker can divert attention and resources away from the true pathogen. So the final Answer is Distraction/decoy to delay diagnosis.	The on; Hypothetical, on answer \n What does it does \n H hyp. Answer \n H? Answer Answer Answer Answer
Forget irrelevant	Which of the following are the spatial clustering algorithms? {Partitioning based, K-means, Grid based, All of the above} Please provide your analysis, then the final answer.	Partitioning-based: partitions data into disjoint clusters. K-means: assigns each point to the nearest cluster centroid. Grid-based: divides space into grid cells and clusters within them. Final Answer: All of the above.	Partitioning minimizes within-cluster variance. K-means uses centroids for assignment. Grid-based clustering reduces search space for large datasets. Final Answer: All of the above.

yields:

$$\underset{\theta}{\text{minimize}} \quad \ell_f(\theta; \mathcal{D}_f) + \gamma \ell_r(\theta; \mathcal{D}_r), \quad (1)$$

where ℓ_f and ℓ_r denote the forget loss and retain loss, respectively, and $\gamma \geq 0$ controls the trade-off between forgetting and utility preservation. A key distinction among unlearning methods lies in how the forget loss ℓ_f is defined.

In this work, we mainly focus on two state-of-the-art LLM unlearning approaches, RMU (Li et al., 2024) and NPO (Zhang et al., 2024b), since they represent the *two main paradigms of LLM unlearning*. RMU follows a representation reengineering paradigm, in which intermediate activations associated with the forget set are modified to enable selective and utility-preserving forgetting. In contrast, NPO adopts an optimization divergence based paradigm, pushing the model away from the forget set through a directional training objective.

Specifically, RMU enforces forgetting by mapping intermediate representations of forget-set samples $\mathbf{x} \in \mathcal{D}_f$ to a fixed random vector, preventing the model from encoding meaningful information about them. This yields:

$$\ell_f(\theta; \mathcal{D}_f) = \mathbb{E}_{\mathbf{x} \in \mathcal{D}_f} [\|M_{\theta}(\mathbf{x}) - c \cdot \mathbf{v}\|_2^2], \quad (2)$$

where $M_{\theta}(\cdot)$ denotes the hidden-state representation extracted from an intermediate layer, $\|\cdot\|_2$ is the ℓ_2 norm, c is a scaling hyperparameter balancing unlearn and utility preservation, and \mathbf{v} is a fixed random vector with entries drawn from a standard uniform distribution.

In contrast to random feature-based methods, NPO treats forget-set samples as negative examples within a direct preference optimization framework (Rafailov et al., 2023). The unlearning objective yields:

$$\ell_f(\theta; \mathcal{D}_f) = \mathbb{E}_{\mathbf{x} \in \mathcal{D}_f} \left[-\frac{2}{\beta} \log \sigma \left(-\beta \log \left(\frac{\pi_{\theta}(\mathbf{x})}{\pi_{\text{ref}}(\mathbf{x})} \right) \right) \right], \quad (3)$$

where $\sigma(\cdot)$ denotes the sigmoid function, $\beta > 0$ is a temperature parameter. $\pi_{\theta}(\mathbf{x})$ represents the model’s output probability, and $\pi_{\text{ref}}(\mathbf{x})$ is the corresponding probability under the reference model. NPO updates θ to explicitly push the model away from its pre-unlearning behavior on forget-set data.

Throughout this work, we perform LLM unlearning on the WMDP benchmark (Li et al., 2024), which targets the removal of harmful knowledge. We choose WMDP because, unlike other benchmarks, it does not require an additional fine-tuning stage before unlearning, which could degrade general utility (Shi et al., 2024; Maini et al., 2024). This allows us to start from standard chat models and attribute observed changes directly to the unlearning methods, rather than to pre-unlearning artifacts. The forget set contains 3,668 multiple-choice questions on hazardous biosecurity and cybersecurity content. *Unlearning effectiveness (UE)* is measured by the accuracy drop on forget-set questions, while *utility preservation (UT)* is evaluated using broad benchmarks such as MMLU (Hendrycks et al., 2020). Training details of the models used in this work and corresponding UE and UT results are provided in Appendix A.

Feasibility of unlearning detection and problem statement. In Table 1, we compare outputs from the Yi-34B model and its RMU-unlearned counterpart on two prompt types: a *forget prompt* from WMDP and a *forget-irrelevant* prompt from MMLU assessing general utility. While the unlearned model suppresses sensitive answers, its responses to forget prompts are often incoherent relative to the original model, whereas both models remain coherent on forget-irrelevant prompts.

To further analyze this behavior, we examine the perplexity (PPL) distributions of the original and unlearned models on both prompt types, using GPT-2 to compute PPL, a standard proxy for fluency and predictability (Qi et al., 2021). As shown in Fig. 2, there is a clear distributional shift in PPL for WMDP forget prompts, while the distributions for forget-irrelevant MMLU prompts remain indistinguishable.

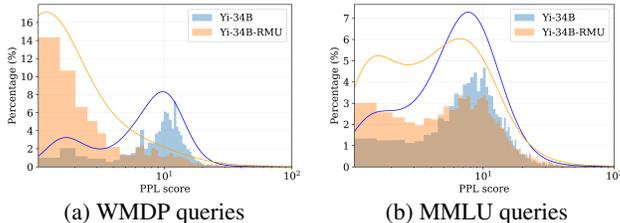


Figure 2: GPT-2 perplexity distributions for Yi-34B vs. RMU-unlearned responses. (a) WMDP forget queries (3,000 samples). (b) MMLU forget-irrelevant queries (3,000 samples). Where perplexity quantifies fluency and predictability.

Motivated by these observations, we pose the problem of **unlearning trace detection**: *Can an unlearned model be distinguished from its original counterpart solely through its outputs?* Furthermore, since an adversary cannot be assumed to have access to forget-relevant prompts, we tackle the more challenging task of distinguishing models using only *forget-irrelevant prompts*, where behavioral differences are subtle (see Table 1 and Fig. 2b).

To analyze the vulnerability of unlearning trace detection, we adopt a worst-case *threat model* in which the adversary can query the model and observe either the “soft” pre-logit activations or the “hard” textual responses of the model. This setting is analogous to adversarial attack design with gray-box access, where the attacker can exploit either hard labels or soft prediction logits to evade ML models (Ilyas et al., 2018). While access to internal activations exceeds a pure black-box assumption, it reflects an increasingly common scenario in open-weight LLM deployments, where pre-logit activations are accessible through standard inference. Importantly, such access is not necessary for detecting unlearning traces, but amplifies vulnerabilities already observable from textual responses.

4 SUPERVISED CLASSIFICATION FOR DETECTING UNLEARNING TRACES

We frame the detection of unlearned models as a supervised classification task, leveraging different forms of model outputs as discriminative signals. Specifically, we consider two types of outputs: (i) *the generated textual responses*, which capture behavioral differences in how models respond to prompts, and (ii) *the pre-logit activations*, which reflect representation-level information.

The victim LLMs are drawn from four instruction-based models: *Zephyr-7B*, *LlaMA-3.1-8B*, *Qwen2.5-14B*, and *Yi-34B*. For each of these models, we obtain their corresponding unlearned variant by applying unlearning on the WMDP dataset (see Appendix A for the unlearning setup). For each query, we collect both the textual response and the corresponding pre-logit activations. To extract pre-logit activations, we record the activation vector associated with each newly generated token. Aggregating these vectors across the generated sequence yields a sequence of activation representations used for supervised classification.

Supervised classification setup. We construct the training dataset using an equal mix of outputs from the *forget* dataset WMDP and the *general utility* dataset MMLU (50% each), denoted as \mathcal{S}_{fg} . At test time, we evaluate on unseen prompts sampled from WMDP, MMLU, and UltraChat, all disjoint from the training set, to assess generalization. Classification accuracy is used as the primary metric to quantify the detectability of unlearning traces. For text-based classification, we encode generated responses using LLM2Vec (BehnamGhader et al., 2024) and fine-tune a two-layer MLP for binary classification (*i.e.*, unlearned vs. original). This encoder is chosen for its strong generalization and robustness to variable-length responses when capturing idiosyncratic behaviors of LLMs (Sun et al., 2025). For activation-based classification, we directly train a two-layer MLP on pre-logit activation representations. Additional classifier and training details are provided in Appendix B.

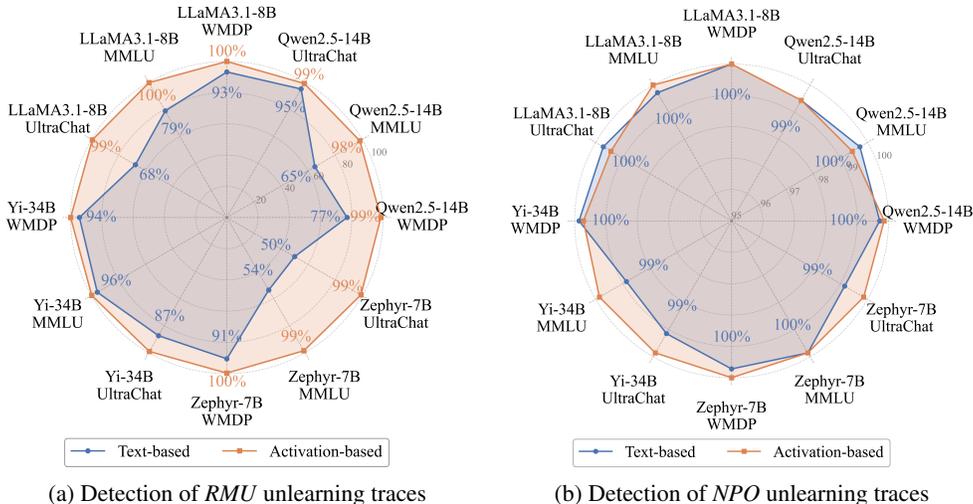


Figure 3: Radar charts of unlearning trace detection accuracy across four source LLMs evaluated on three test sets (WMDP, MMLU, UltraChat). Panel (a) reports results for models unlearned on WMDP using RMU, while panel (b) reports results for models unlearned on WMDP using NPO. Each axis corresponds to a detection setting defined by a source model *A* and a test dataset *B*, where the classifier is trained on outputs from model *A* and evaluated on prompts from dataset *B*. Results are shown for two output types: text-based responses (blue) and pre-logit activations (orange). Detailed numerical results are provided in Appendix C.

Detection of RMU unlearning. We present the evaluation results for the classification of RMU-unlearned models in Fig. 3(a). As shown, the classifier for unlearning-trace detection trained on *pre-logit activations* achieves substantially higher accuracy than the one trained on *textual responses*. This trend holds consistently across different base model types. Most importantly, the superior detection performance of activation-based classifiers persists even when tested on *forget-irrelevant* data points from MMLU and UltraChat. As will be shown later, the near-perfect classification achieved using pre-logit activations can be attributed to deeper reasons, specifically, the presence of *unlearning fingerprints* embedded in the activation space of an LLM once it has been unlearned. In addition, when examining classification based solely on textual responses, we obtain two further insights. First, *larger LLMs yield higher unlearning-trace detection accuracy* on both forget-relevant test sets (drawn from WMDP) and forget-irrelevant ones. For example, the case (Yi-34B, MMLU) achieves 96% accuracy, whereas the case (Zephyr-7B, MMLU) reaches only 54%, a level close to random guessing. Second, *classification based on textual responses to forget-relevant prompts is substantially easier* than to forget-irrelevant prompts. For instance, with Zephyr-7B, detection on WMDP achieves 91% accuracy, whereas on MMLU it drops significantly. This is consistent with our motivating example in Fig. 2.

Detection of NPO unlearning. We next present the evaluation results for classification of NPO-unlearned models in Fig. 3(b). In stark contrast to the RMU case, classifiers trained solely on text responses are sufficient to achieve high accuracy across all base model types and for both forget-relevant and forget-irrelevant queries. For example, even Zephyr-7B, whose performance was close to random guessing under RMU, achieves 99–100% accuracy in all NPO settings. This behavior arises from the degradation of responses in NPO-unlearned models, even to forget-irrelevant prompts. We provide supporting evidence by showcasing example responses from NPO models (see Table A7 and Table A8 in Appendix D) and by analyzing fine-grained differences between RMU- and NPO-unlearned models (see Table A9 in Appendix E). These findings also reflect the fundamental design differences between RMU and NPO. The NPO objective in (3) enforces explicit deviation from the pre-trained model, often resulting in aggressive forgetting that makes the model’s behavior noticeably different from the original. In contrast, RMU’s localized manipulation of internal representations in (2) produces *subtler* changes, leaving weaker traces at the response level and making detection notably harder on general prompts.

As noted above, the success of activation-based classification, even for RMU-unlearned models tested on forget-irrelevant data, indicates the presence of distinctive unlearning fingerprints. This motivates us to further investigate these traces by probing the internal activations across different layers of the model in the next section.

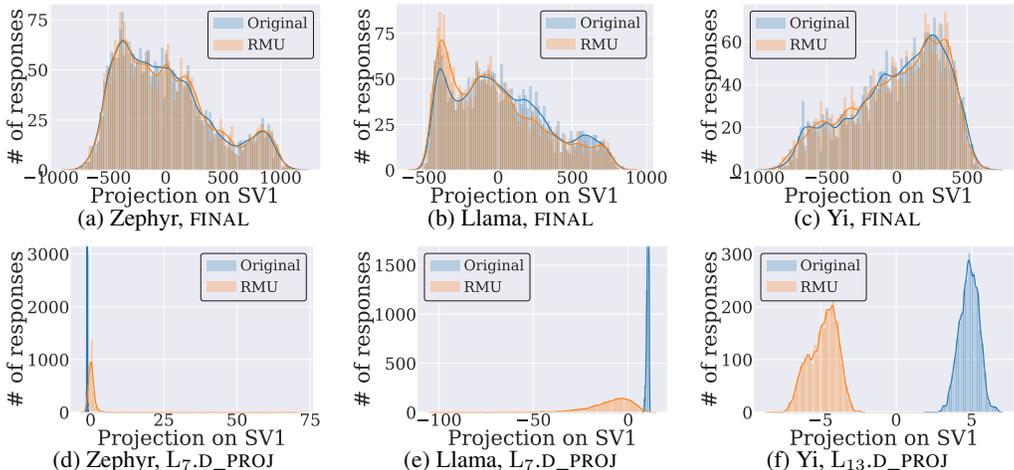


Figure 4: Projection of activations from various layers for 3000 responses to MMLU onto the top right singular vectors (denoted as SV1) for both original and unlearned models. Here, $L_i.D_PROJ$ refers to activations extracted from the down-projection sublayer of the FFN in the i -th transformer block, while FINAL denotes the activations of the final layer after RMS-norm. (a,d) for Zephyr-7B, (b,e) for LLaMA3.1-8B, and (c,f) for Yi-34B.

5 UNVEILING FINGERPRINTS OF UNLEARNED MODELS

In this section, we present our analysis showing that unlearning leaves behind distinct activation-level “fingerprints”, which provide clear explanations for the classification results reported in Sec. 4.

Spectral “fingerprints”: Definition and method. We define spectral fingerprints of unlearning as characteristic shifts in a model’s internal activations, observed along principal directions of variation. As described in Sec. 4, we obtain activation vectors for each model output. Following the approach in (Tran et al., 2018), we perform singular value decomposition (SVD) on the centered activation matrix and project the activations onto the right singular vectors to visualize and analyze the spectral shifts induced by unlearning. To examine how unlearning affects these internal representations, we generate 100-token responses for 3,000 randomly sampled MMLU test questions using both the original and unlearned models. Here, we focus on the most challenging unlearning trace detection scenario: identifying traces from model responses to forget-irrelevant prompts drawn from MMLU. The presence of an unlearning fingerprint is revealed through the *correct localization* of these activation shifts, which we elaborate on in the following analysis.

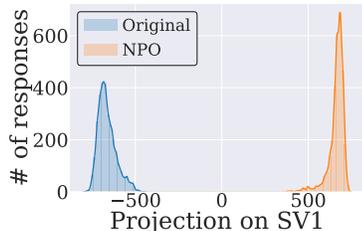


Figure 5: Projection of the final-layer normalized activations from 3,000 MMLU responses onto the first right singular vector (SV1) for the original LLaMA3.1-8B model and its NPO-unlearned counterpart.

NPO exhibits strong spectral fingerprints. For NPO-unlearned models, we extract the final normalized activations: specifically, the outputs from the last layer after root mean square normalization (RMSNorm). As shown in Fig. 5, there is a pronounced distributional shift between the unlearned and original models when activations are projected onto the first right singular vector (SV1). This observation aligns with the results in Fig. 3(b), where classifiers achieve near-perfect accuracy in distinguishing NPO-unlearned responses from those of the original model. Additional spectral fingerprint results for other models are provided in the Appendix F.

RMU exhibits subtle but clear spectral fingerprints. Following the same procedure, we extract the final pre-logit activations for RMU-unlearned models (denoted as FINAL). As shown in Fig. 4-(a-c), there is no apparent distributional shift in the projected activations that would allow us to distinguish the RMU-unlearned models from their original counterparts. To investigate further, we examine activations from intermediate layers, *i.e.*, the layers directly modified by RMU as described in (2). Specifically, we extract activations from sublayers within the feed-forward network (FFN) of inter-

mediate layers, *i.e.*, the down-projection (D_PROJ) and gate-projection (G_PROJ) sublayers. When extracting from layer i (denoted as L_i), we refer to the corresponding activations as L_i .D_PROJ and L_i .G_PROJ, respectively. As shown in Fig. 4-(d-f), all models exhibit spectral shifts in the activation distributions for responses generated by the RMU-unlearned model. For Zephyr-7B, the fingerprint appears exclusively in the projection of L_7 .D_PROJ along the first singular vector. Although present, the distributional shift is relatively subtle, validating the model’s lower classification accuracy in Fig. 3(a) under textual responses. For LLaMA3.1-8B, we again observe spectral fingerprints in L_7 .D_PROJ along the top singular direction, though the shift is more pronounced compared to Zephyr. The strongest spectral fingerprints are observed in Yi-34B, with clear shifts across multiple layers, notably L_{13} , L_{14} , and L_{15} . This observation is consistent with the high classification accuracy achieved at the larger model scale in Fig. 3(a).

As shown in Fig. 4, spectral fingerprints characterized by distributional shifts were not observed in the final pre-logit activations of RMU-unlearned models. However, due to the residual stream architecture of transformers (Elhage et al., 2021), earlier activations (where RMU fingerprints are found) contribute indirectly to the final output. This suggests that the unlearning signal may still be embedded in the final activations, albeit in a more complex form. To uncover this effect, we apply supervised UMAP (McInnes et al., 2018), a non-linear dimensionality reduction technique. In Fig. 6, UMAP yields a clearer separation between original and RMU-unlearned activations at FINAL for Zephyr-7B. Additional results for other models are presented in the Appendix G. This suggests the existence of a low-dimensional nonlinear manifold on which the final activations of original and unlearned models are well separated, even for models such as Zephyr-7B that exhibit only subtle spectral shifts in Fig. 4. This explains why a neural network trained directly on pre-logit activations in Sec. 4 can achieve near-perfect classification accuracy: it effectively learns this separable manifold, thereby revealing the unlearning fingerprints.

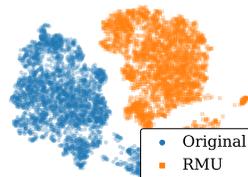


Figure 6: UMAP projections of FINAL layer activations on MMLU prompts, comparing the original and RMU-unlearned Zephyr-7B models.

6 EXPERIMENTS

In this section, we present additional experimental results and analyses on unlearning-trace detection.

Generalization ability of supervised classification on unseen models. We explore whether our trained unlearning-trace detection classifiers can generalize to unseen models beyond those used during training. As described in Sec. 4, the classifier was trained and evaluated on outputs from a given source model type. Here, we examine whether a classifier trained on the outputs of one model can successfully generalize to the outputs of another model at test time.

In Table 2, we present the detection accuracy of classifiers trained on pre-logit activations from one model type (specified by the model name in each row, *e.g.*, Zephyr-7B) when evaluated on outputs from another model type (specified by the model name in each column, *e.g.*, Yi-34B). The diagonal entries of Table 2 represent the *intra-domain* generalization setting, where the source model type used for classifier training and testing is the same, but evaluation is performed on outputs generated from disjoint test-time queries (50% forget-relevant prompts + 50% forget-irrelevant prompts). In contrast, the off-diagonal entries correspond to the *cross-domain* generalization setting, where a classifier trained on one source model type is evaluated on outputs from a different model type using test data. It is interesting to observe that classifiers trained on Zephyr-7B, Qwen2.5-14B, and Yi-34B generalize well across domains, achieving detection accuracy between 82.45% and 99.58%. An exception is the classifier trained

Table 2: Generalization of RMU-unlearning detection using classifiers trained on pre-logit activations. The training setup follows Sec. 4. The model type in each row indicates the source model used to construct the training set, while the model type in each column indicates the target model used for testing.

Train \ Test	Test			
	Zephyr-7B	Llama-3.1-8B	Qwen2.5-14B	Yi-34B
Zephyr-7B	99.87%	75.03%	99.45%	98.89%
Llama-3.1-8B	57.96%	99.58%	48.95%	52.43%
Qwen2.5-14B	95.45%	82.45%	98.45%	98.96%
Yi-34B	94.24%	75.31%	95.35%	99.93%

on LLaMA-3.1-8B, which does not exhibit strong cross-domain generalization to test-time models different from itself. This suggests that a stronger unlearning-trace detector needs to account for model-specific characteristics. A possible future direction is to train classifiers using data aggregated from multiple source model types, particularly for those that show weaker generalization.

Classifier training using only forget-relevant data.

In Sec. 4, we performed supervised classification by training on a mixed dataset of responses (denoted by \mathcal{S}_{fg}). Here, we extend the study to a simpler case in which the training dataset is constructed using only responses from the forget dataset WMDP, denoted by \mathcal{S}_f . The evaluation results of this setup are presented in Table 3. As shown, a classifier trained solely on the forget set is sufficient for successful classification of responses to forget-relevant prompts, achieving accuracy above 91% (see the WMDP column). However, such classifiers fail to distinguish responses to *forget-irrelevant prompts* (e.g., MMLU and UltraChat), with accuracy dropping to near-random guessing. This ablation study highlights the importance of using the mixed dataset \mathcal{S}_{fg} for classifier training, as it enables the classifier to learn to detect unlearning traces even under forget-irrelevant prompts. More results using different training regimes for RMU and NPO unlearn detection, please refer to Table A11 and Table A12 in Appendix I

An extended use case: Forget-data detection.

While our paper primarily focuses on detecting whether a model has undergone unlearning, a natural follow-up question arises: *given that an unlearned model is successfully detected, can we further determine whether its responses originate from the forget domain?* We refer to this task as **forget-data detection**.

To this end, we need to shift from distribution-level model characteristics with and without unlearning (*i.e.*, learning the discriminative ability to distinguish between forget and forget-irrelevant data distributions like Fig. 6) to data-level characteristics (*i.e.*, learning the ability to determine, for each individual data point, whether it belongs to the forget set). However, once an unlearned model has been successfully identified by our proposal, we can leverage rich statistics of the forget data against the unlearned model to construct forget data detection metrics. As shown in Appendix H, we find that unlearning often induces significant shifts in entropy, JS divergence, and top- k prediction probability mass on forget data compared to forget-irrelevant data. This trend holds consistently for both NPO- and RMU-based unlearning approaches. Therefore, we leverage these four data features and use the NPO-unlearned Zephyr-7B model as a reference to construct prototypes for both forget-relevant and forget-irrelevant prompts. At test time, for each input-response pair, we compute its detection metrics and compare them with the established detection prototypes, identifying it as forget-relevant if the prototype criteria are satisfied. Fig. 7 reports detection accuracies for different unlearned models (NPO or RMU) evaluated on a test dataset consisting of a balanced mix of forget-relevant and forget-irrelevant prompts. We observe that unlearned models indeed leave data-wise detectable traces, with detection accuracy exceeding 70% across test data points. In particular, NPO variants are consistently the easiest to detect, reflecting their more deterministic output distributions (low entropy, high maximum probability).

Evaluation using complex classification heads. We further evaluate whether increasing classifier complexity improves detection performance by testing two more advanced classifier architectures: (1) Residual head, a residual MLP classifier in which two hidden layers form a residual block, increasing representational capacity while maintaining stable optimization; (2) Deeper head, an extended classifier that increases depth from a two-layer MLP to a four-layer feed-forward network with GELU activations and intermediate layer normalization. The detection results for Llama3.1-8B RMU unlearning based on text-based responses are presented in Table 4.

Table 3: Classification accuracy for distinguishing original vs. RMU-unlearned models under two training regimes: \mathcal{S}_{fg} and \mathcal{S}_f . Columns report test accuracy on WMDP, MMLU, and UltraChat prompts, with no overlap with the training sets.

Model	Setting	WMDP	MMLU	UltraChat
LLaMA-3.1-8B	\mathcal{S}_{fg}	93.24%	78.87%	67.60%
	\mathcal{S}_f	95.49%	51.83%	55.21%
Qwen2.5-14B	\mathcal{S}_{fg}	95.07%	76.90%	65.07%
	\mathcal{S}_f	94.93%	54.08%	56.62%
Yi-34B	\mathcal{S}_{fg}	94.37%	95.77%	87.46%
	\mathcal{S}_f	91.69%	61.41%	58.72%

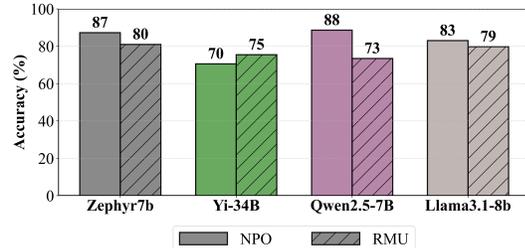


Figure 7: Forget-data detection accuracies across different unlearned models, using NPO or RMU applied to various source model types.

As shown, none of the more complex classifier architectures yields improved performance, on either forget-relevant or forget-irrelevant queries, for detecting RMU-based unlearning. This further supports our conclusion that the detection task is fundamentally *unlearn signal-fingerprinted*. For example, even lightweight classifiers already achieve near-perfect accuracy for NPO-based unlearning methods, including when trained solely on text responses. In contrast, RMU tends to produce more conservative and subtle fingerprints compared to preference-based unlearning methods. This aligns with the underlying paradigms: RMU perturbs internal representations in a targeted local way, whereas preference-based methods impose global behavioral shifts, leaving stronger and more easily detectable signatures.

Table 4: Detection accuracy of Llama3.1-8B RMU unlearning using different classifier heads based on text-based responses. ‘‘Original’’ denotes the classifier used in Sec. 4. ‘‘Residual Head’’ uses a residual MLP structure, and ‘‘Deeper Head’’ uses a deeper feed forward structure.

Test sets	WMDP	MMLU	UltraChat
Original	93.24%	78.87%	67.60%
Residual Head	94.67%	81.23%	69.57%
Deeper Head	90.23%	76.54%	65.42%

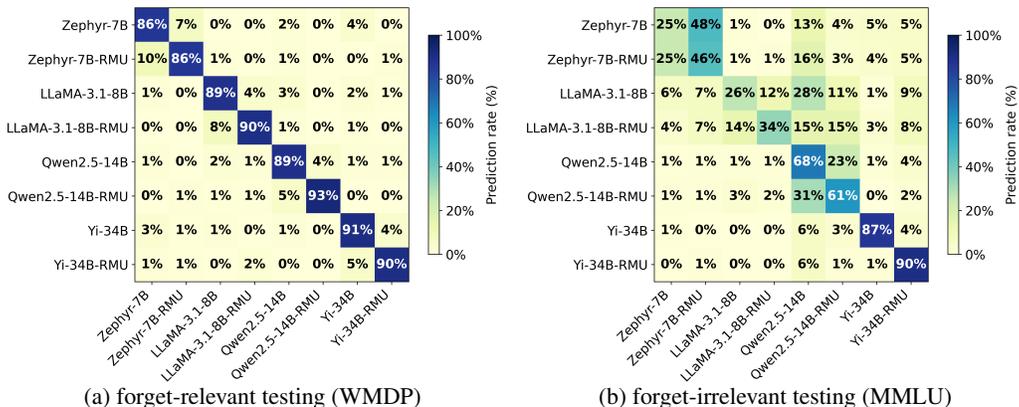


Figure 8: Confusion matrices for model-unlearning pair classification. Rows denote the true classes (*i.e.*, original or unlearned versions for each LLM type), and columns indicate the predicted classes. Diagonal entries correspond to correct predictions, while off-diagonal entries reflect misclassifications. Results are shown for (a) WMDP (forget-related) and (b) MMLU (forget-irrelevant) test sets.

Additional results. To further validate our findings, we evaluate our detection pipeline with different pretrained encoders in Appendix J. In addition, Fig. 8 (see also Fig. A5 in Appendix K) presents an 8-way classification study distinguishing model families (original vs. unlearned), further demonstrating that unlearning traces are model-dependent signals rather than artifacts of binary detection. We also evaluate our detection accuracy using classifiers trained on different mixing ratios in Appendix L and the robustness of unlearning detection given different evaluation strategies in Appendix M.

7 CONCLUSION

In this work, we revisited LLM unlearning from a new perspective: the detectability of unlearning traces. We showed that, although intended to remove sensitive knowledge, unlearning leaves persistent fingerprints in both behavior and internal representations. Across diverse LLMs, methods, and prompt types, simple classifiers can distinguish unlearned models from originals, with spectral fingerprints in hidden activations enabling near-perfect detection. These findings expose a critical vulnerability, as unlearning traces may facilitate reverse-engineering attacks that undermine privacy and safety guarantees. We refer readers to Appendix N–O for limitations, broad impact, and LLM usage.

ACKNOWLEDGMENTS

The contributions of Yiwei Chen, Soumyadeep Pal, Yimeng Zhang, and Sijia Liu were supported by the NSF CISE Core Program Award IIS-2504263, the NSF CAREER Award IIS-2338068, the Open Philanthropy Research Award, the Center for AI Safety (CAIS) Research Award, and the Schmidt Sciences’ Trustworthy AI Program. Qing Qu acknowledges NSF IIS 2402950, ONR N000142512339, DARPA HR0011578254, and the Google Research Scholar Award.

REFERENCES

- Fazl Barez, Tingchen Fu, Ameya Prabhu, Stephen Casper, Amartya Sanyal, Adel Bibi, Aidan O’Gara, Robert Kirk, Ben Bucknall, Tim Fist, et al. Open problems in machine unlearning for ai safety. *arXiv preprint arXiv:2501.04952*, 2025.
- Parishad BehnamGhader, Vaibhav Adlakha, Marius Mosbach, Dzmitry Bahdanau, Nicolas Chapados, and Siva Reddy. Llm2vec: Large language models are secretly powerful text encoders. In *CoLM*, 2024.
- Lucas Bourtole, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning. In *2021 IEEE Symposium on Security and Privacy (SP)*, pp. 141–159. IEEE, 2021.
- Yinzhi Cao and Junfeng Yang. Towards making systems forget with machine unlearning. In *2015 IEEE symposium on security and privacy*, pp. 463–480. IEEE, 2015.
- Jiaao Chen and Diyi Yang. Unlearn what you want to forget: Efficient unlearning for llms. *arXiv preprint arXiv:2310.20150*, 2023.
- Min Chen, Zhikun Zhang, Tianhao Wang, Michael Backes, Mathias Humbert, and Yang Zhang. When machine unlearning jeopardizes privacy. In *Proceedings of the 2021 ACM SIGSAC conference on computer and communications security*, pp. 896–911, 2021.
- Weixin Chen, Baoyuan Wu, and Haoqian Wang. Effective backdoor defense by exploiting sensitivity of poisoned samples. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *NeurIPS*, 2022.
- Yiwei Chen, Yuguang Yao, Yihua Zhang, Bingquan Shen, Gaowen Liu, and Sijia Liu. Safety mirage: How spurious correlations undermine vlm safety fine-tuning and can be mitigated by machine unlearning. *arXiv preprint arXiv:2503.11832*, 2025.
- A Feder Cooper, Christopher A Choquette-Choo, Miranda Bogen, Matthew Jagielski, Katja Filippova, Ken Ziyu Liu, Alexandra Chouldechova, Jamie Hayes, Yangsibo Huang, Niloofar Mireshghallah, et al. Machine unlearning doesn’t do what you think: Lessons for generative ai policy, research, and practice. *arXiv preprint arXiv:2412.06966*, 2024.
- Aghyad Deeb and Fabien Roger. Do unlearning methods remove information from language model weights? *arXiv preprint arXiv:2410.08827*, 2024.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, pp. 4171–4186, 2019.
- Ning Ding, Yulin Chen, Bokai Xu, Yujia Qin, Shengding Hu, Zhiyuan Liu, Maosong Sun, and Bowen Zhou. Enhancing chat language models by scaling high-quality instructional conversations. In *EMNLP*, pp. 3029–3051, 2023.
- Ronen Eldan and Mark Russinovich. Who’s harry potter? approximate unlearning in llms, 2023.
- Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, et al. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 1(1):12, 2021.
- Chongyu Fan, Jiancheng Liu, Licong Lin, Jinghan Jia, Ruiqi Zhang, Song Mei, and Sijia Liu. Simplicity prevails: Rethinking negative preference optimization for llm unlearning. *arXiv preprint arXiv:2410.07163*, 2024.
- Peter Hase, Mohit Bansal, Been Kim, and Asma Ghandeharioun. Does localization inform editing? surprising differences in causality-based localization vs. knowledge editing in language models. *Advances in Neural Information Processing Systems*, 36:17643–17668, 2023.
- Jonathan Hayase, Weihao Kong, Raghav Somani, and Sewoong Oh. Spectre: Defending against backdoor attacks using robust statistics. In *International Conference on Machine Learning*, pp. 4129–4139. PMLR, 2021.

- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *ICLR*, 2020.
- Chris Jay Hoofnagle, Bart van der Sloot, and Frederik Zuiderveen Borgesius. The european union general data protection regulation: what it is and what it means. *Information & Communications Technology Law*, 28(1):65–98, 2019.
- Shengyuan Hu, Yiwei Fu, Steven Wu, and Virginia Smith. Jogging the memory of unlearned models through targeted relearning attacks. In *ICML 2024 Workshop on Foundation Models in the Wild*, 2024.
- Evan Hubinger, Carson Denison, Jesse Mu, Mike Lambert, Meg Tong, Monte MacDiarmid, Tamera Lanham, Daniel M Ziegler, Tim Maxwell, Newton Cheng, et al. Sleeper agents: Training deceptive llms that persist through safety training. *arXiv preprint arXiv:2401.05566*, 2024.
- Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. *arXiv preprint arXiv:2212.04089*, 2022.
- Andrew Ilyas, Logan Engstrom, Anish Athalye, and Jessy Lin. Black-box adversarial attacks with limited queries and information. In *International conference on machine learning*, pp. 2137–2146. PMLR, 2018.
- Joel Jang, Dongkeun Yoon, Sohee Yang, Sungmin Cha, Moontae Lee, Lajanugen Logeswaran, and Minjoon Seo. Knowledge unlearning for mitigating privacy risks in language models. *arXiv preprint arXiv:2210.01504*, 2022.
- Jinghan Jia, Jiancheng Liu, Parikshit Ram, Yuguang Yao, Gaowen Liu, Yang Liu, Pranay Sharma, and Sijia Liu. Model sparsity can simplify machine unlearning. *NeurIPS*, 36:51584–51605, 2023.
- Jinghan Jia, Jiancheng Liu, Yihua Zhang, Parikshit Ram, Nathalie Baracaldo, and Sijia Liu. WAGLE: Strategic weight attribution for effective and modular unlearning in large language models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024a.
- Jinghan Jia, Yihua Zhang, Yimeng Zhang, Jiancheng Liu, Bharat Runwal, James Diffenderfer, Bhavya Kailkhura, and Sijia Liu. Soul: Unlocking the power of second-order optimization for llm unlearning. *arXiv preprint arXiv:2404.18239*, 2024b.
- Max Lamparath and Anka Reuel. Analyzing and editing inner mechanisms of backdoored language models. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, pp. 2362–2373, 2024.
- Nathaniel Li, Alexander Pan, Anjali Gopal, Summer Yue, Daniel Berrios, Alice Gatti, Justin D Li, Ann-Kathrin Dombrowski, Shashwat Goel, Gabriel Mukobi, et al. The wmdp benchmark: Measuring and reducing malicious use with unlearning. In *ICML*, 2024.
- Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pp. 74–81, 2004.
- Chin-Yew Lin and Franz Josef Och. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *ACL-04*, 2004.
- Sijia Liu, Yuanshun Yao, Jinghan Jia, Stephen Casper, Nathalie Baracaldo, Peter Hase, Yuguang Yao, Chris Yuhao Liu, Xiaojun Xu, Hang Li, et al. Rethinking machine unlearning for large language models. *Nature Machine Intelligence*, pp. 1–14, 2025.
- Ximing Lu, Sean Welleck, Jack Hessel, Liwei Jiang, Lianhui Qin, Peter West, Prithviraj Amanabrolu, and Yejin Choi. Quark: Controllable text generation with reinforced unlearning. In *NeurIPS*, 2022.
- Jakub Łucki, Boyi Wei, Yangsibo Huang, Peter Henderson, Florian Tramèr, and Javier Rando. An adversarial perspective on machine unlearning for ai safety. *arXiv preprint arXiv:2409.18025*, 2024.

- Aengus Lynch, Phillip Guo, Aidan Ewart, Stephen Casper, and Dylan Hadfield-Menell. Eight methods to evaluate robust unlearning in llms. *arXiv preprint arXiv:2402.16835*, 2024.
- Pratyush Maini, Zhili Feng, Avi Schwarzschild, Zachary C. Lipton, and J. Zico Kolter. Tofu: A task of fictitious unlearning for llms, 2024.
- Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in gpt. *NeurIPS*, 2022.
- Nay Myat Min, Long H Pham, Yige Li, and Jun Sun. Crow: Eliminating backdoors from large language models via internal consistency regularization. *arXiv preprint arXiv:2411.12768*, 2024.
- Thanh Tam Nguyen, Thanh Trung Huynh, Phi Le Nguyen, Alan Wee-Chung Liew, Hongzhi Yin, and Quoc Viet Hung Nguyen. A survey of machine unlearning. *arXiv preprint arXiv:2209.02299*, 2022.
- Ivica Nikolic, Teodora Baluta, and Prateek Saxena. Model provenance testing for large language models. *arXiv preprint arXiv:2502.00706*, 2025.
- Martin Pawelczyk, Seth Neel, and Himabindu Lakkaraju. In-context unlearning: Language models as few shot unlearners. *arXiv preprint arXiv:2310.07579*, 2023.
- Fanchao Qi, Yangyi Chen, Mukai Li, Yuan Yao, Zhiyuan Liu, and Maosong Sun. Onion: A simple and effective defense against textual backdoor attacks. In *EMNLP*, pp. 9558–9566, 2021.
- Youyang Qu, Ming Ding, Nan Sun, Kanchana Thilakarathna, Tianqing Zhu, and Dusit Niyato. The frontier of data erasure: Machine unlearning for large language models. *arXiv preprint arXiv:2403.15779*, 2024.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 2019.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *NeurIPS*, 2023.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*, 2020.
- Protection Regulation. Regulation (eu) 2016/679 of the european parliament and of the council. *Regulation (eu)*, 679:2016, 2016.
- Rohin Shah, Alex Irpan, Alexander Matt Turner, Anna Wang, Arthur Conmy, David Lindner, Jonah Brown-Cohen, Lewis Ho, Neel Nanda, Raluca Ada Popa, et al. An approach to technical agi safety and security. *arXiv preprint arXiv:2504.01849*, 2025.
- Weijia Shi, Jaechan Lee, Yangsibo Huang, Sadhika Malladi, Jieyu Zhao, Ari Holtzman, Daogao Liu, Luke Zettlemoyer, Noah A Smith, and Chiyuan Zhang. Muse: Machine unlearning six-way evaluation for language models. *arXiv preprint arXiv:2407.06460*, 2024.
- Nianwen Si, Hao Zhang, Heyu Chang, Wenlin Zhang, Dan Qu, and Weiqiang Zhang. Knowledge unlearning for llms: Tasks, methods, and challenges. *arXiv preprint arXiv:2311.15766*, 2023.
- Mingjie Sun, Yida Yin, Zhiqiu Xu, J Zico Kolter, and Zhuang Liu. Idiosyncrasies in large language models. *arXiv preprint arXiv:2502.12150*, 2025.
- Di Tang, XiaoFeng Wang, Haixu Tang, and Kehuan Zhang. Demon in the variant: Statistical analysis of {DNNs} for robust backdoor contamination detection. In *USENIX Security*, 2021.
- Pratiksha Thaker, Yash Maurya, and Virginia Smith. Guardrail baselines for unlearning in llms. *arXiv preprint arXiv:2403.03329*, 2024.

- Anvith Thudi, Gabriel Deza, Varun Chandrasekaran, and Nicolas Papernot. Unrolling sgd: Understanding factors influencing machine unlearning. In *European Symposium on Security and Privacy (EuroS&P)*, pp. 303–319. IEEE, 2022.
- Brandon Tran, Jerry Li, and Aleksander Madry. Spectral signatures in backdoor attacks. *Advances in neural information processing systems*, 31, 2018.
- Xinwei Wu, Junzhuo Li, Minghui Xu, Weilong Dong, Shuangzhi Wu, Chao Bian, and Deyi Xiong. Depn: Detecting and editing privacy neurons in pretrained language models. *arXiv preprint arXiv:2310.20138*, 2023.
- Zhiguang Yang and Hanzhou Wu. A fingerprint for large language models. *arXiv preprint arXiv:2407.01235*, 2024.
- Yuanshun Yao, Xiaojun Xu, and Yang Liu. Large language model unlearning. *arXiv preprint arXiv:2310.10683*, 2023.
- Yuanshun Yao, Xiaojun Xu, and Yang Liu. Large language model unlearning. In *NeurIPS*, 2024a.
- Yuguang Yao, Yifan Gong, Yize Li, Yimeng Zhang, Xue Lin, and Sijia Liu. Reverse engineering of imperceptible adversarial image perturbations. In *ICLR*, 2022.
- Yuguang Yao, Xiao Guo, Vishal Asnani, Yifan Gong, Jiancheng Liu, Xue Lin, Xiaoming Liu, Sijia Liu, et al. Reverse engineering of deceptions on machine-and human-centric attacks. *Foundations and Trends® in Privacy and Security*, 6(2):53–152, 2024b.
- Charles Yu, Sullam Jeoung, Anish Kasi, Pengfei Yu, and Heng Ji. Unlearning bias in language models by partitioning gradients. In *Findings of ACL*, 2023.
- Ruiqi Zhang, Licong Lin, Yu Bai, and Song Mei. Negative preference optimization: From catastrophic collapse to effective unlearning. *arXiv preprint arXiv:2404.05868*, 2024a.
- Ruiqi Zhang, Licong Lin, Yu Bai, and Song Mei. Negative preference optimization: From catastrophic collapse to effective unlearning. In *CoLM*, 2024b.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*, 2019.
- Yihua Zhang, Yimeng Zhang, Yuguang Yao, Jinghan Jia, Jiancheng Liu, Xiaoming Liu, and Sijia Liu. Unlearncanvas: A stylized image dataset to benchmark machine unlearning for diffusion models. *arXiv preprint arXiv:2402.11846*, 2024c.
- Yimeng Zhang, Xin Chen, Jinghan Jia, Yihua Zhang, Chongyu Fan, Jiancheng Liu, Mingyi Hong, Ke Ding, and Sijia Liu. Defensive unlearning with adversarial training for robust concept erasure in diffusion models. *NeurIPS*, 2024d.
- Yimeng Zhang, Jinghan Jia, Xin Chen, Aochuan Chen, Yihua Zhang, Jiancheng Liu, Ke Ding, and Sijia Liu. To generate or not? safety-driven unlearned diffusion models are still easy to generate unsafe images... for now. In *ECCV*. Springer, 2024e.
- Zhexin Zhang, Junxiao Yang, Pei Ke, Shiyao Cui, Chujie Zheng, Hongning Wang, and Minlie Huang. Safe unlearning: A surprisingly effective and generalizable solution to defend against jailbreak attacks. *arXiv preprint arXiv:2407.02855*, 2024f.
- Zhiahao Zhang, Yiwei Chen, Weizhan Zhang, Caixia Yan, Qinghua Zheng, Qi Wang, and Wangdu Chen. Tile classification based viewport prediction with multi-modal fusion transformer. In *ACM-MM*, 2023.
- Zhihao Zhang, Shengcao Cao, and Yu-Xiong Wang. Tamm: Triadapter multi-modal learning for 3d shape understanding. In *CVPR*, 2024g.
- Zhihao Zhang, Abhinav Kumar, Girish Chandar Ganesan, and Xiaoming Liu. Unleashing the power of chain-of-prediction for monocular 3d object detection. *arXiv preprint arXiv:2505.04594*, 2025.

Sally Zhu, Ahmed Ahmed, Rohith Kuditipudi, and Percy Liang. Independence tests for language models. *arXiv preprint arXiv:2502.12292*, 2025.

Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th international conference on machine learning (icml-03)*, pp. 928–936, 2003.

APPENDIX

A UNLEARNING CONFIGURATION AND DATA PREPARATION

Unlearning setups. We apply both RMU and NPO unlearning algorithms to four LLMs (Zephyr-7B, Llama-3.1-8B, Qwen2.5-7B, and Yi-34B) using the WMDP benchmark. To evaluate forget utility, we evaluate each unlearned model on both the WMDP-bio and WMDP-cyber subsets, while in order to assess general utility, we measure performance on MMLU. The results are summarized in **Table A1**.

For RMU unlearning of the Zephyr-7B model, we set the control scaling factor c in Eq. (2) to 6.5, $\gamma = 1200$. Then we perform unlearning by optimizing layer 5,6,7 while calculating the unlearning loss in Eq. (1) using the seventh intermediate layer of M_θ . For Llama-3.1-8B, scaling factor is set to 45, $\gamma = 1300$ and the other settings are consistent with the Zephyr-7B model. For the Qwen2.5-14B model, we set $c = 460$, $\gamma = 350$. The unlearning loss in Eq.(1) is computed using the activations immediately following the tenth intermediate layer and we perform parameter updates on layers 8,9,10. Finally, for Yi-34B-Chat, $c = 300$, $\gamma = 350$ and unlearning is performed exclusively on the layers 13, 14 and 15 using activations from the fifteenth intermediate layer.

For NPO, we perform unlearning for 140 steps with a batch size of 4. For different models, we use different learning rates and different utility regularization γ in Eq. (3) and present these setups in **Table A2**.

Classification data construction. To generate both forget and forget-irrelevant responses from each model, we first extracted the questions from the WMDP, MMLU, and UltraChat datasets. For UltraChat forget-irrelevant examples, we provided only the question itself. Similarly for the WMDP “forget” questions, we provide only the question and task the LLM to generate the corresponding answer. Only for MMLU, since much of the semantic content resides in the answer choices, we retrieved the question together with its choices and used the following prompt:

```
messages = [
  {
    "role": "user",
    "content": (
      f"{question}\n"
      f"{choices}\n\n"
      "Please provide your analysis, then give the final answer."
      "Analysis:"
    )
  }
]
```

Table A1: Unlearning effectiveness is measured on WMDP and general utility on MMLU for each LLM after applying RMU and NPO unlearning on WMDP. Both evaluations report the accuracy on four-choice question answering.

Model	WMDP-bio	WMDP-cyber	MMLU
Zephyr-7B	64.65%	44.44%	58.49%
+RMU	30.64%	27.78%	57.45%
+NPO	24.82%	37.09%	48.01%
Llama-3.1-8B	69.84%	43.94%	63.36%
+RMU	38.75%	25.06%	59.64%
+NPO	26.86%	37.24%	54.59%
Qwen2.5-14B	80.54%	52.99%	77.56%
+RMU	29.69%	26.72%	76.16%
+NPO	39.43%	45.94%	72.09%
Yi-34B-Chat	74.00%	49.27%	72.35%
+RMU	30.79%	28.59%	70.63%
+NPO	32.91%	30.39%	41.54%

Table A2: Unlearning Setup for NPO. γ refers to the utility regularization.

Model	Learning Rate	γ
Zephyr-7B	7e-06	1.0
Llama-3.1-8B	2e-05	2.0
Qwen2.5-14B	7e-05	1.0
Yi-34B	6e-05	1.0

B CLASSIFIER TRAINING PROTOCOL

Data splits. We investigate classifier performance under three training regimes:

- \mathcal{S}_{fg} : An equal mixture of WMDP (forget) responses and MMLU (forget-irrelevant) responses. We randomly sample 2,900 questions from each benchmark, collect the corresponding model responses, and train on the combined 5,800 sample sets.
- \mathcal{S}_f : only WMDP (forget) responses which come from 2,900 question samples.
- \mathcal{S}_g : only MMLU (forget-irrelevant) responses which come from 5,800 question samples.

To assess generalization, we hold out 355 unseen questions and their responses from each benchmark, which is disjoint from all training set, and evaluate the classifier on the three test subsets, which includes MMLU, WMDP and UltraChat.

Response-based classifier training details. In our classifier architecture, we adopt LLM2vec as our base encoder, a lightweight sentence-level model tailored for open-ended LLM outputs, and append a two-layer MLP head to produce logits over the binary label space (original vs. unlearned). All experiments were conducted on an NVIDIA A6000 GPU. We fine-tune the entire network end-to-end under a standard supervised learning protocol, training for three epochs with AdamW (weight decay 0.001) and a cosine decay schedule (initial learning rate 8×10^{-5} , warmup ratio 0.1). We use a batch size of 8, mixed-precision BF16, gradient clipping at 0.3, and enable gradient checkpointing to reduce memory usage. All data splits and random seeds (42) for sampling, initialization, and shuffling are fixed for reproducibility.

Activation-based classifier training details. We construct an MLP-based classifier operating directly on hidden activations extracted from LLM forward passes. The architecture progressively compresses high-dimensional representations (e.g., Zephyr-7B’s 409,600 dimension) through a four-layer network: $d_{in} \rightarrow 1024 \rightarrow 256 \rightarrow 128 \rightarrow 2$. Each hidden layer is followed by BatchNorm and Dropout for regularization, with Xavier initialization ensuring stable convergence. Training is performed under a supervised classification setup (original vs. unlearned). We adopt AdamW (weight decay 10^{-3}) with a cosine learning rate schedule (initial LR 8×10^{-5} , warmup ratio 0.1), training for three epochs on an NVIDIA A6000 GPU. We use a batch size of 8, mixed-precision BF16, gradient clipping at 0.3, and enable gradient checkpointing to reduce memory usage.

Dimension mismatch is a fundamental obstacle for activation-based classifiers. When a classifier trained on Zephyr’s 409,600-dimensional activations is evaluated on LLaMA’s 512,000-dimensional space, the learned decision boundaries no longer align with the evaluation inputs. To mitigate this issue, we explored both directions of transfer. For high-to-low settings, we applied dimensionality reduction (e.g., PCA) to compress larger activation spaces into smaller ones. For low-to-high settings, we attempted zero padding to embed smaller activations into a larger space to deal with the dimension mismatch problem.

C ADDITIONAL UNLEARNING DETECTION RESULTS

Detectability of RMU. In Table A3, we present the test accuracy of the classifier trained on the mixed dataset of forget and forget-irrelevant responses (\mathcal{S}_{fg}), generated from various source LLMs (indicated by *row* names). Evaluation is conducted on model responses to unseen prompts from WMDP, MMLU, and UltraChat (indicated by *column* names), with all test prompts disjoint from those used during training. As we can see, responses to WMDP prompts are highly distinguishable at test time, with classification accuracies exceeding 90% across all models. This indicates that RMU-

based unlearning leaves clearly detectable traces in responses tied to the forget set. In contrast, classification accuracy drops significantly when evaluating on MMLU and UltraChat, which contain forget-irrelevant prompts. For example, the classifier achieves only 53.68% accuracy on MMLU and 50.14% on UltraChat when distinguishing Zephyr-7B responses, near random guessing. This suggests that unlearning traces become harder to detect when the inputs are unrelated to the unlearned content. Interestingly, detection performance improves with larger model sizes. Yi-34B achieves

Table A3: Classification accuracy for distinguishing original vs. RMU-unlearned models, with unlearning applied to the WMDP dataset. Rows indicate the source LLM used for response generation and classifier training. Columns show test accuracy on responses to prompts from WMDP, MMLU, and UltraChat, all of which are disjoint from the training set to ensure generalization.

Model	WMDP	MMLU	UltraChat
Zephyr-7B	90.56%	53.68%	50.14%
LLaMA-3.1-8B	93.24%	78.87%	67.60%
Qwen2.5-14B	95.07%	76.90%	65.07%
Yi-34B	94.37%	95.77%	87.46%

95.77% accuracy on MMLU and 87.46% on UltraChat, indicating that *unlearning traces in larger models are more persistent and detectable, even under general prompts*.

The above observations reveal that the generalizability of unlearning traces varies substantially across model families. Some models, particularly larger ones, exhibit broad behavioral shifts that are readily identifiable from output text alone, even when responses are not directly related to the unlearning target. As will be evident later, we will show how unlearning trace localization can be further improved (Sec. 5) and how this leads to stronger classification performance (Sec. 6). Additional classification results trained under different dataset configurations will be provided in Table A11.

Table A5 reports unlearning classification accuracy when using RMS-normalized final pre-logit activations as input to a two-layer MLP, compared against our text-based baseline (Tab. A3). Across all four source LLMs (Yi-34B, Zephyr-7B, LLaMA-3.1-8B, and Qwen2.5-14B) and three test sets (WMDP, MMLU, UltraChat), activation-based features yield a substantial gain as elaborated on below. 1. Worst-case improvement: For Zephyr-7B on MMLU, detection jumps from just 53% (text) to 98% (activations), over 40% increase. 2. Consistent gains on “forget-irrelevant”: On UltraChat, where text signals are most subtle, accuracy rises across all models. 3. Overall robustness: The mean accuracy across all twelve evaluation points increases, demonstrating that unlearning traces are more linearly separable in activation space. These results confirm that the final pre-logit activations encode stronger, model-internal signatures of unlearning than the raw text outputs alone. The primary drawback is the requirement for white-box access to extract these activations, which may not be feasible in every deployment scenario.

Detectability of NPO. In Table A4, we present the classification accuracy when identifying the NPO-unlearned model, in contrast to Table A3 that focuses on RMU unlearning. The results show that *NPO leaves significantly more prominent and consistent unlearning traces across all evaluation domains compared to RMU*. All four LLMs achieve near-perfect classification accuracy on WMDP, MMLU, and UltraChat, indicating that NPO introduces strong and easily detectable changes to model behavior, even in response to general, forget-irrelevant prompts. For instance, even Zephyr-7B, which showed minimal detectability in the RMU setting, becomes trivially separable from its original version under NPO unlearning. These results also mirror the design differences between RMU and NPO. NPO’s objective in (3) enforces the deviation from the pre-trained model. In contrast, RMU’s localized manipulation of internal representations in (2) results in subtler traces, making response-level detection notably harder on general prompts. Additional classification results for NPO traces under different training regimes are provided in the Appendix A12.

Table A6 shows the results for NPO unlearning detection using the activations. Text-based classifiers (marked using blue in Fig. 3) already achieve high accuracy, often above 99% in distinguishing original from NPO-unlearned responses, reflecting the presence of unlearning artifacts in the text. Consequently, activation-based features (orange) have relatively limited headroom for improvement, pushing models to a ceiling of 99–100% across all tasks. This narrow margin confirms that NPO traces are so pronounced in the text itself that final-layer embeddings offer only incremental gains. The success of activation-based classification on RMU-unlearned models, including forget-irrelevant cases, reveals distinctive unlearning signatures.

Table A4: Classification accuracy for distinguishing original vs. NPO-unlearned models. All setups remain consistent with Table A3.

Model	WMDP	MMLU	UltraChat
Zephyr-7B	99.72%	99.86%	99.16%
LLaMA-3.1-8B	100%	99.72%	99.72%
Qwen2.5-14B	99.72%	99.72%	99.44%
Yi-34B	99.86%	98.87%	99.15%

Table A5: Classification accuracy using RMS-normalized final-layer activations to distinguish original vs. RMU-unlearned responses (unlearning applied on WMDP). All other settings mirror those in Table A3.

Model	WMDP	MMLU	UltraChat
Zephyr-7B	99.72%	98.59%	99.15%
LLaMA-3.1-8B	100%	99.72%	99.44%
Qwen2.5-14B	98.59%	98.31%	99.15%
Yi-34B	100%	99.86%	99.86%

Table A6: Classification accuracy using RMS-normalized final-layer activations to distinguish original vs. NPO-unlearned responses (unlearning applied on WMDP). All other settings mirror those in Table A4.

Model	WMDP	MMLU	UltraChat
Zephyr-7B	100%	99.86%	99.86%
LLaMA-3.1-8B	100%	100%	99.44%
Qwen2.5-14B	99.86%	99.44%	99.44%
Yi-34B	99.72%	99.86%	99.86%

Table A8 reports analogous comparisons for Zephyr-7B and its RMU- and NPO-unlearned variants. Both unlearning methods severely disrupt the forget-prompt response—RMU yields semi-coherent but heavily garbled fragments, while NPO collapses into extended runs of punctuation and nonsensical tokens. Crucially, across both Yi-34B and Zephyr-7B, NPO always induces more aggressive degradation than RMU: even though both unlearned models produce correct answer selections on the MMLU-style “forget-irrelevant” prompt, NPO’s generated text exhibits a higher incidence of raw, undecoded token sequences and formatting artifacts. This pattern holds despite preserved selection accuracy, demonstrating that NPO shifts the answer generation behavior more radically than RMU while leaving the surface choice unaffected.

E FINE-GRAINED DIFFERENCES BETWEEN RMU AND NPO

To better understand the differing unlearning characteristics of RMU and NPO, we conduct a fine-grained analysis comparing the lexical and stylistic properties of their responses against those from the original model. We quantify alignment with the original using ROUGE-1 and ROUGE-L (Lin (2004); Lin & Och (2004)), which measure lexical overlap and structural similarity, respectively. Additionally, we employ BERTScore (Zhang et al. (2019)), which evaluates token-level semantic similarity using contextual embeddings from a pre-trained model (*e.g.*, BERT (Devlin et al. (2019))), offering a more nuanced comparison beyond surface-level matching.

Table A9 provides further evidence of the distinct behavioral impacts induced by NPO and RMU. Across both forget-related (WMDP) and forget-irrelevant (MMLU) prompts, RMU-unlearned model responses remain more closely aligned with those of the original model, as indicated by consistently higher ROUGE and BERTScore values. This supports our earlier classification results, where RMU traces were harder to detect—especially on forget-irrelevant prompts.

Table A9: F1 scores of lexical and semantic similarity metrics (ROUGE-1, ROUGE-L, BERTScore) for RMU- and NPO-unlearned Yi-34B responses compared to the original model, averaged over 3,000 prompts from WMDP (forget-relevant) and MMLU (forget-irrelevant). Higher scores indicate greater alignment.

Dataset	Model	ROUGE-1	ROUGE-L	BERTScore
WMDP	RMU	0.1597	0.1178	0.7852
	NPO	0.0187	0.0139	0.6982
MMLU	RMU	0.2493	0.1509	0.7703
	NPO	0.0160	0.0115	0.6836

In contrast, NPO-unlearned responses exhibit substantial drops across all similarity metrics, signaling broader lexical and semantic divergence from the original. The effect is particularly pronounced on MMLU (*e.g.*, ROUGE-1 drops to 0.0160 for NPO vs. 0.2493 for RMU), suggesting that NPO alters even non-targeted responses. These findings reinforce the conclusion from Table A4: NPO induces more aggressive, globally detectable behavioral shifts, whereas RMU’s effects are more subtle and localized. Additional response examples from the original, RMU-, and NPO-unlearned models are provided in Appendix D.

F DETAILED SPECTRAL FINGERPRINTS

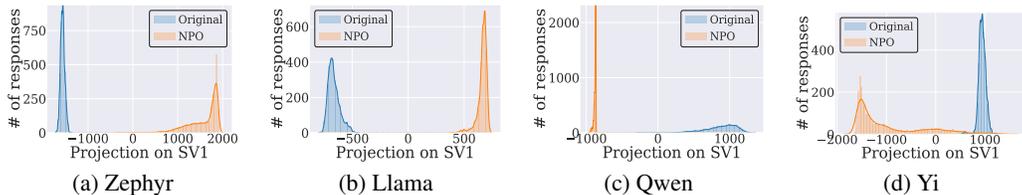


Figure A1: Projection of the final-layer normalized activations from 3,000 MMLU responses onto the first right singular vector (SV1) for the original and its NPO-unlearned counterpart. (a) is projection for Zephyr-7B, (b) for Llama3.1-8B, (c) for Qwen2.5-14B, (d) for Yi-34B.

Spectral fingerprints for NPO-unlearned models. In Fig. A1, we present the spectral fingerprints of models unlearned using NPO, using activations of the last layer after normalization. Consistent with our observations in Sec. 5, NPO reliably exhibits a strong separation, simply using these activations projected onto the first singular vector, thus confirming the presence of a strong fingerprint.

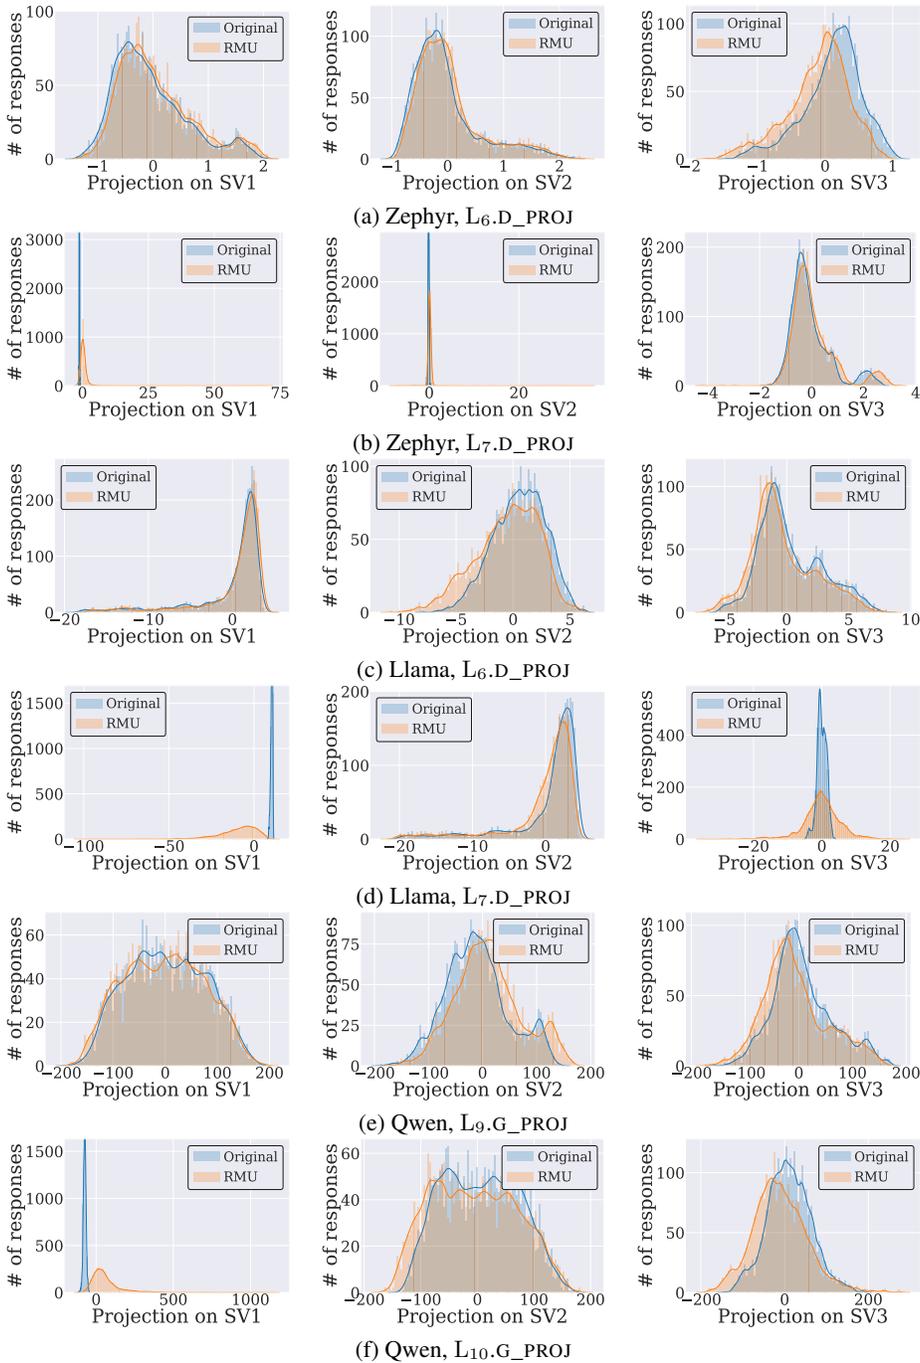


Figure A2: Projection of activations from various layers for 3000 responses to MMLU onto the three leading right singular vectors for the original and unlearned model. $L_i.D_PROJ$ refers to activations extracted from the down-projection sublayer of the FFN in the i -th transformer block, while $L_i.G_PROJ$ refers to activations extracted from the gate-projection sublayer of the FFN in the i -th transformer block (a,b) are projections for Zephyr-7B, (c,d) are for Llama3.1-8B, while (e,f) are for Qwen2.5-14B.

Spectral fingerprints for RMU-unlearned models. As detailed in Sec. 5, RMU exhibits subtle fingerprints and therefore, we analyze the activations projected onto the top three singular vectors. We explored such fingerprints for layers directly modified by RMU, details of which are provided in Appendix A. We demonstrate detailed fingerprints for models unlearned using RMU in Fig. A2 and Fig. A3. For Zephyr-7B- β , Fig. A2-(b) reveals the presence of a spectral fingerprint in $L_7.D_PROJ$

projected along the top right singular vector, while Fig. A2-(a) shows a mild shift in $L_6.D_PROJ$ projected onto the third leading right singular vector. Similar mild shifts appear for other models in various other projections throughout Fig. A2. Llama3.1-8B exhibits a clear fingerprint is present in $L_7.D_PROJ$ projected onto the top right singular vector(Fig. A2-(d)), while for Qwen2.5-14B shows a comparable effect in $L_{10}.G_PROJ$ projected onto the top right singular vector (Fig. A2-(f)). Finally, in line with the high classification accuracy for Yi-34B-Chat, Fig. A3-(a-c) highlights distinct fingerprints in the activations from three layers *i.e.* $L_{13}.D_PROJ$, $L_{14}.D_PROJ$ and $L_{15}.D_PROJ$ projected onto the top right singular vector, where the spectral shift is especially pronounced in the first two.

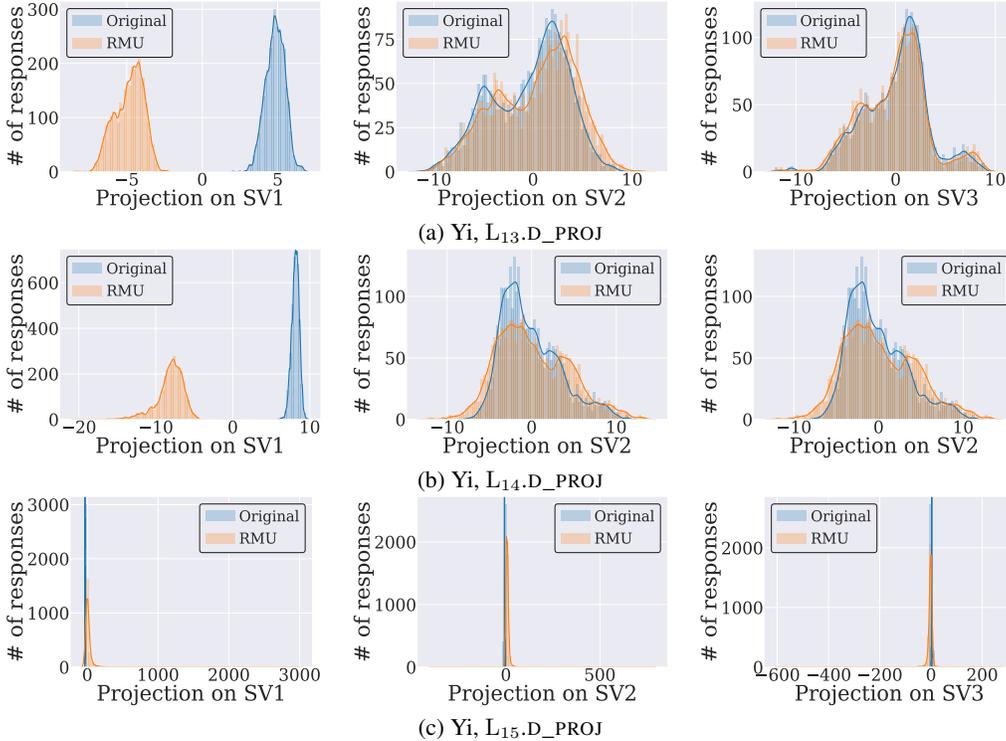


Figure A3: Projection of activations of Yi-34B-Chat from various layers for 3000 responses to MMLU onto the three leading right singular vectors for the original and unlearned model. $L_i.D_PROJ$ refers to activations extracted from the down-projection sublayer of the FFN in the i -th transformer block (a) are projections from layer 13, (b) are from layer 14, (c) are from layer 15.

G A CLOSER LOOK AT FINAL ACTIVATIONS

Similar to Sec. 5, we present the supervised UMAP projections of the final activations from different models in Fig. A4. Consistent with Sec. 5, UMAP always yields clear separation between the original and RMU-unlearned activations.

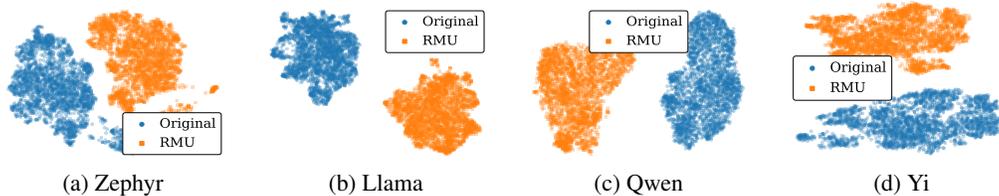


Figure A4: Supervised UMAP Projections of the final-layer normalized activations from 3,000 MMLU responses for the original and its RMU-unlearned counterpart using (a) Zephyr-7B, (b) Llama3.1-8B, (c) Qwen2.5-14B, (d) Yi-34B-Chat.

H DISTRIBUTIONAL SHIFTS IN NEXT-TOKEN PREDICTION AFTER UNLEARNING

H.1 DISTRIBUTION METRICS FOR NEXT-TOKEN PREDICTION.

We analyze unlearning effects through the next-token prediction distribution. Given a prompt x , the model defines a categorical distribution over the vocabulary V as

$$p_{\theta}(\cdot | x) \in \Delta^{|V|-1}, \quad (\text{A1})$$

where θ denotes the model parameters and $\Delta^{|V|-1}$ is the probability simplex. From this distribution, we compute several statistical indicators that serve as features for forget-data detection:

Entropy (H). Entropy measures the overall uncertainty of the model’s next-token prediction:

$$H(p_{\theta}) = - \sum_{i=1}^{|V|} p_{\theta}(y_i | x) \log p_{\theta}(y_i | x). \quad (\text{A2})$$

A lower entropy indicates a peaked, deterministic distribution, while higher entropy reflects more uncertainty.

Maximum probability (P_{\max}). The maximum predicted probability captures the model’s confidence in its most likely token:

$$P_{\max}(p_{\theta}) = \max_i p_{\theta}(y_i | x). \quad (\text{A3})$$

High P_{\max} values suggest overly confident predictions, often observed in unlearned models.

Top- k probability mass (M_k). The probability mass concentrated on the top- k predictions is:

$$M_k(p_{\theta}) = \sum_{i \in \text{Top-}k} p_{\theta}(y_i | x). \quad (\text{A4})$$

This reflects how much of the distribution is allocated to a small set of likely tokens.

Jensen–Shannon divergence (JS). To quantify distributional shifts, we compare p_{θ} against a reference distribution $p_{\theta_{\text{ref}}}$ from the original model:

$$\text{JS}(p_{\theta}, p_{\theta_{\text{ref}}}) = \frac{1}{2} \text{KL}(p_{\theta} \| m) + \frac{1}{2} \text{KL}(p_{\theta_{\text{ref}}} \| m), \quad (\text{A5})$$

where $m = \frac{1}{2}(p_{\theta} + p_{\theta_{\text{ref}}})$. A larger JS divergence indicates stronger deviation of the unlearned model from its original counterpart.

Together, these four indicators capture complementary aspects of distributional behavior: uncertainty (H), confidence (P_{\max}), concentration (M_k), and deviation from the reference model (JS). Based on this, we use them as the basis for forget data detection

H.2 DISTRIBUTIONAL SHIFTS ACROSS MODELS

Table A10 summarizes the distributional statistics of next-token prediction for both forget-irrelevant (MMLU) and forget-relevant (WMDP) prompts, comparing the *Original*, *RMU-unlearned*, and *NPO-unlearned* variants of each model. Several consistent trends emerge. For forget-irrelevant inputs, distributional shifts are relatively mild, though NPO often induces sharper changes such as reduced entropy and increased maximum probability. In contrast, for forget-relevant inputs, the divergence becomes more pronounced: RMU models typically exhibit higher entropy and more dispersed probability mass, whereas NPO models collapse into highly deterministic distributions with near-unit top- k mass and maximum probability. These results highlight that unlearning introduces systematic, data-dependent biases in token-level distributions, providing the basis for our forget-data detection analysis in Sec. 6.

I DETECTION OF UNLEARNING UNDER DIFFERENT TRAINING REGIMES

Table A10: Distributional statistics of next-token prediction for forget-irrelevant (MMLU) and forget-relevant (WMDP) inputs. Each cell reports the values for *Original* \rightarrow *RMU* \rightarrow *NPO*. We evaluate four metrics: (1) **Entropy**, measuring overall uncertainty of the next-token distribution; (2) *JS divergence*: quantifying deviation from the original model’s predictions; (3) *Top-k probability mass*: indicating how much probability is concentrated on the most likely tokens; and (4) *Maximum probability*: reflecting the model’s confidence in its top prediction.

Model	Entropy	JS div.	Top-k mass	Max prob
MMLU (forget-irrelevant)				
Zephyr-7b	3.374 \rightarrow 3.484 \rightarrow 0.463	0.334 \rightarrow 0.325 \rightarrow 0.182	0.965 \rightarrow 0.962 \rightarrow 1.000	0.406 \rightarrow 0.391 \rightarrow 0.872
Yi-34B-Chat	1.889 \rightarrow 2.072 \rightarrow 5.972	0.504 \rightarrow 0.499 \rightarrow 0.371	0.998 \rightarrow 0.995 \rightarrow 0.446	0.586 \rightarrow 0.575 \rightarrow 0.291
Llama3.1-8b	5.556 \rightarrow 5.806 \rightarrow 1.257	0.239 \rightarrow 0.242 \rightarrow 0.214	0.850 \rightarrow 0.836 \rightarrow 0.972	0.196 \rightarrow 0.194 \rightarrow 0.826
Qwen2.5-14b	4.107 \rightarrow 4.565 \rightarrow 4.107	0.404 \rightarrow 0.372 \rightarrow 0.404	0.940 \rightarrow 0.916 \rightarrow 0.940	0.350 \rightarrow 0.309 \rightarrow 0.350
WMDP (forget-relevant)				
Zephyr-7b	1.963 \rightarrow 3.875 \rightarrow 0.001	0.609 \rightarrow 0.590 \rightarrow 0.122	0.993 \rightarrow 0.858 \rightarrow 1.000	0.621 \rightarrow 0.542 \rightarrow 1.000
Yi-34B-Chat	1.841 \rightarrow 4.212 \rightarrow 0.468	0.552 \rightarrow 0.864 \rightarrow 0.677	0.998 \rightarrow 0.901 \rightarrow 0.963	0.575 \rightarrow 0.343 \rightarrow 0.943
Llama3.1-8b	4.434 \rightarrow 7.288 \rightarrow 0.000	0.352 \rightarrow 0.450 \rightarrow 0.129	0.923 \rightarrow 0.687 \rightarrow 1.000	0.279 \rightarrow 0.151 \rightarrow 1.000
Qwen2.5-14b	3.645 \rightarrow 5.572 \rightarrow 3.645	0.534 \rightarrow 0.420 \rightarrow 0.534	0.957 \rightarrow 0.842 \rightarrow 0.957	0.369 \rightarrow 0.239 \rightarrow 0.369

RMU-unlearned classification under different training regimes. Recall from Sec. 4 that the default training dataset for the supervised classifier, denoted as \mathcal{S}_{fg} , consists of a 50/50 mix of forget-related and forget-irrelevant responses. To examine how unlearning detection varies under different training data compositions, we consider two additional regimes: \mathcal{S}_f , which includes only WMDP forget-related responses (100%), and \mathcal{S}_g , which includes only MMLU forget-irrelevant responses (100%). **Table A11** presents the performance of detecting RMU-unlearned model across the three training regimes for four LLMs. When trained solely on \mathcal{S}_f , nearly all models achieve higher accuracy on forget-related prompts (e.g., 97.20% for Zephyr-7B) compared to training on \mathcal{S}_{fg} , but their performance drops to near-random levels (around 50%) on forget-irrelevant queries. In contrast, training on \mathcal{S}_g , which lacks direct relevance to the unlearning target, fails to enable effective trace detection, even when evaluated on forget-relevant WMDP prompts. In summary, as forget-irrelevant responses used for training contain the least fingerprint information and are weakly correlated with unlearning traces. The mixed regime \mathcal{S}_{fg} , by combining both response types, consistently achieves strong performance across all evaluations.

NPO-unlearned classification under different training regimes. In contrast to RMU (Table A11), NPO traces are so pronounced that classification accuracy remains near-perfect (>97%) under all three training regimes. As shown in **Table A12**, even when the classifier is trained exclusively on forget-irrelevant MMLU data (\mathcal{S}_g), it still achieves over 99% accuracy on WMDP “forget” prompts, and above 98% on UltraChat for all models. Training on forget-only data (\mathcal{S}_f) likewise yields over 97% detection on “forget irrelevant” prompts. The mixed regime (\mathcal{S}_{fg}) offers no substantial benefit over the single-domain regimes, underscoring that NPO’s aggres-

Table A11: Classification accuracy for distinguishing original vs. RMU-unlearned models under three training regimes: \mathcal{S}_{fg} , \mathcal{S}_f , and \mathcal{S}_g . Columns report test accuracy on WMDP, MMLU, and UltraChat prompts, with no overlap with the training sets.

Model	Setting	WMDP	MMLU	UltraChat
Zephyr-7B	\mathcal{S}_{fg}	90.56%	53.68%	50.14%
	\mathcal{S}_f	97.20%	51.55%	51.83%
	\mathcal{S}_g	50.00%	52.67%	50.83%
LLaMA-3.1-8B	\mathcal{S}_{fg}	93.24%	78.87%	67.60%
	\mathcal{S}_f	95.49%	51.83%	55.21%
	\mathcal{S}_g	68.45%	79.72%	69.30%
Qwen2.5-14B	\mathcal{S}_{fg}	95.07%	76.90%	65.07%
	\mathcal{S}_f	94.93%	54.08%	56.62%
	\mathcal{S}_g	73.66%	76.06%	64.37%
Yi-34B	\mathcal{S}_{fg}	94.37%	95.77%	87.46%
	\mathcal{S}_f	91.69%	61.41%	58.72%
	\mathcal{S}_g	68.73%	98.87%	84.42%

Table A12: Classification accuracy for distinguishing original vs. NPO-unlearned responses under three training regimes: \mathcal{S}_{fg} , \mathcal{S}_f , and \mathcal{S}_g . All experiments use four LLMs with NPO unlearning applied on the WMDP dataset. The settings are consistent with Table A11.

Model	Setting	WMDP	MMLU	UltraChat
Zephyr-7B	\mathcal{S}_{fg}	99.72%	99.86%	99.16%
	\mathcal{S}_f	100%	99.58%	98.73%
	\mathcal{S}_g	99.72%	100%	99.15%
LLaMA-3.1-8B	\mathcal{S}_{fg}	100%	99.72%	99.72%
	\mathcal{S}_f	99.72%	98.03%	97.46%
	\mathcal{S}_g	100%	85.93%	99.72%
Qwen2.5-14B	\mathcal{S}_{fg}	99.72%	99.72%	99.44%
	\mathcal{S}_f	99.72%	99.44%	99.15%
	\mathcal{S}_g	99.86%	99.72%	99.86%
Yi-34B	\mathcal{S}_{fg}	99.86%	98.87%	99.15%
	\mathcal{S}_f	99.86%	99.86%	98.45%
	\mathcal{S}_g	99.72%	100%	99.58%

sive output artifacts are easily learned regardless of training composition. By comparison, RMU required mixed-domain exposure to reach robust performance (Sec. 6), highlighting the fundamentally stronger and domain-agnostic nature of NPO unlearning traces.

J EFFECT OF PRETRAINED ENCODER ON CLASIFIER PERFORMANCE

To evaluate the impact of classifier architecture on unlearning trace detection, we compare a range of pretrained text encoders, following the protocol of BehnamGhader et al. (2024). Specifically, we experiment with classifiers based on BERT Devlin et al. (2019), T5 Raffel et al. (2020), GPT-2 Radford et al. (2019), and LLM2vec BehnamGhader et al. (2024), each paired with a lightweight two-layer MLP head. Each model is trained to distinguish between responses from the original and unlearned LLMs. As shown in Table A13, LLM2vec consistently achieves the highest classification accuracy across all evaluation settings, motivating its adoption as our default classifier architecture.

To further probe how unlearning strength affects trace detectability across encoder architectures, we repeat our classification evaluation under the same mixed regime (\mathcal{S}_{fg}) for both RMU- and NPO-unlearned Yi-34B outputs. Table A13 and Table A14 report accuracy when distinguishing original from unlearned responses using four different pretrained encoders. For RMU unlearning (Table A13), all encoders perform well on the WMDP “forget” data and MMLU “forget-irrelevant” data, but LLM2vec achieves the highest overall robustness, especially on UltraChat, where it attains 87.46% accuracy versus below 70% for the others. This validates our choice of LLM2vec as the default detector when unlearning traces are relatively subtle.

In stark contrast, Table A14 describes NPO unlearning yields near-perfect detection across both prompt types and all domains. Even the least robust encoder (T5) attains over 86% on UltraChat, while LLM2vec, GPT-2, and BERT all exceed 94% everywhere, with LLM2vec surpassing 99% on every test. This demonstrates that NPO’s more aggressive unlearning introduces globally visible artifacts, like raw token fragments and formatting anomalies, that make trace detection trivial, even on “forget-irrelevant” prompts where RMU traces often remain hidden.

K DISTINGUISHING UNLEARNING TRACES ALONGSIDE SOURCE MODEL TYPES

We extend our response-based analysis to a more complex 8-way classification task that jointly distinguishes among four LLM families, each in both their original and unlearned forms. This setup enables a more fine-grained examination of *model-specific unlearning traces*. Fig. 8 displays the resulting confusion matrices on both forget-related (WMDP) and forget-irrelevant (MMLU) test sets. On WMDP, predictions are highly concentrated along the diagonal, indicating strong agreement between the predicted and true model-unlearning pairs. This confirms that unlearning traces are clearly detectable when test prompts align with the domain of the forgotten content. In contrast, classification accuracy declines on the MMLU test set, particularly for the Zephyr-7B models, where most errors involve confusion between the original and RMU-unlearned versions. Nevertheless, larger models such as Yi-34B and Yi-34B-RMU maintain high accuracy, suggesting that unlearning traces in these models persist and remain detectable even when evaluated on general, forget-irrelevant prompts. Results for NPO-unlearned models under this multi-class setting are reported in Fig. A5.

Table A13: Classification accuracy for distinguishing original vs. RMU-unlearned responses using different pretrained sequence encoders. The source LLM is Yi-34B with RMU applied on the WMDP dataset. All other settings mirror those in Table A3.

Classifier	WMDP	MMLU	UltraChat
LLM2vec	94.37%	95.77%	87.46%
T5	85.35%	82.96%	59.72%
GPT2	88.03%	96.06%	62.39%
BERT	88.59%	88.31%	69.15%

Table A14: Classification accuracy for distinguishing original vs. NPO-unlearned responses using different pretrained sequence encoders. The other settings are consistent with Table A13.

Classifier	WMDP	MMLU	UltraChat
LLM2vec	99.86%	98.87%	99.15%
T5	99.29%	99.30%	86.20%
GPT2	99.72%	99.86%	96.90%
BERT	99.44%	99.58%	94.65%

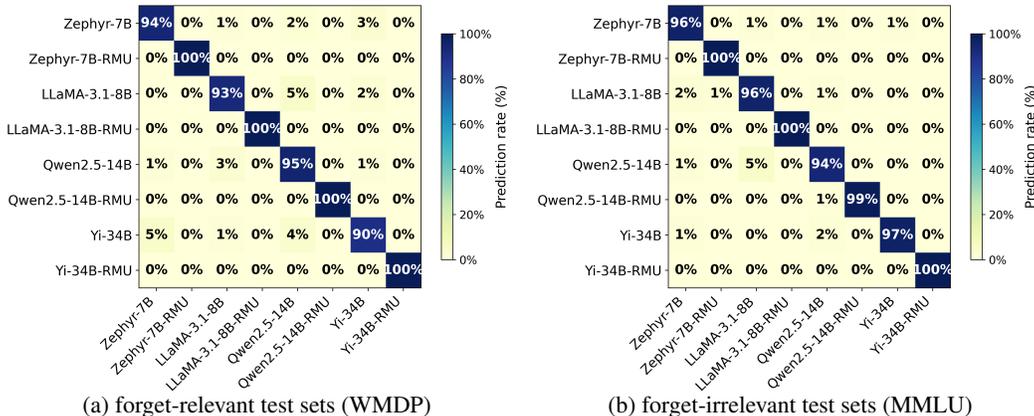


Figure A5: Confusion matrix for NPO-unlearning pair classification. Rows indicate true classes (original/NPO-unlearned model variants), and columns show predicted classes. Diagonal entries represent correct predictions; off-diagonals indicate misclassification rates under (a) WMDP and (b) MMLU test sets.

L ADDITIONAL RESULTS ON MIXING RATIOS

Since NPO shows strong generalization even when trained only with forget-irrelevant queries (as shown in Appendix I), we focus on classifiers built on Llama-3.1-8B for detecting RMU under different mixing ratios (25%/75%, 50%/50%, 75%/25%). The response-based results are reported in Table A15. Based on the results, varying the mixing ratio between forget-relevant and forget-irrelevant samples does not produce any noticeable change in detection performance across WMDP, MMLU, and UltraChat. This indicates that as long as the training split stays within a reasonable range, the unlearning trace remains consistently learnable, meaning the classifier does not rely on a fragile or overly specific mixture proportion to succeed.

Table A15: Detection accuracy of Llama3.1-8B RMU unlearning under different mixing percentages of forget relevant and forget irrelevant queries. The top section reports detection accuracy using classifiers trained on text-based responses, and the bottom section reports detection accuracy using classifiers trained on pre-logit activations.

(a) Text based responses			
Mix Percentage	WMDP	MMLU	UltraChat
25% / 75%	91.43%	81.23%	70.54%
50% / 50%	93.24%	78.87%	67.60%
75% / 25%	94.37%	79.48%	69.32%
(b) Pre logit activations			
Mix Percentage	WMDP	MMLU	UltraChat
25% / 75%	99.86%	99.72%	98.59%
50% / 50%	100%	99.72%	99.44%
75% / 25%	99.72%	98.87%	99.15%

M ADDITIONAL EVALUATION RESULTS

To assess the robustness of unlearning-trace detection under sampling variability and paraphrasing, we performed two additional test-time analyses using the Llama-3.1-8B model. Notably, we used the same trained classifier from our original setup; only the test-time response generation was modified. First, we introduce the Pass@K evaluation, where responses are generated with temperature 0.5 (top-p = 1.0, top-k = -1.0), and a prediction is counted as correct if any of the K sample responses is classified correctly. Second, we performed Paraphrase evaluation by paraphrasing each response using Qwen2.5-32B-Instruct with a strict semantics-preserving prompt: “Paraphrase the following text while strictly preserving its original semantic meaning. Do not add or remove any information.” We reported the detection accuracy of NPO and RMU in Table A16, respectively. NPO remains highly detectable across all conditions, indicating strong robustness to sampling noise and paraphrasing. RMU also shows consistent improvements under Pass@K and maintains similar accuracy under paraphrasing. Thus, unlearning traces persist across decoding strategies and semantic rewordings, confirming that the findings are not fragile or prompt-specific.

Table A16: Response-based detection accuracy of Llama3.1-8B under different evaluation strategies. The evaluation strategies are defined as follows: “Original” denotes standard single response evaluation. “Pass@3” and “Pass@5” denote evaluating the best prediction among three or five sampled responses. “Paraphrase” denotes evaluating on semantically rewritten versions of the original queries.

Evaluation	WMDP (NPO)	WMDP (RMU)	MMLU (NPO)	MMLU (RMU)	UltraChat (NPO)	UltraChat (RMU)
Original	100%	93.24%	99.72%	78.87%	99.72%	67.60%
Pass@3	99.86%	96.65%	99.93%	82.54%	99.93%	73.43%
Pass@5	100%	97.86%	99.93%	84.63%	99.93%	74.61%
Paraphrase	98.59%	92.87%	98.31%	78.12%	99.15%	67.44%

N LIMITATIONS AND BROAD IMPACT

N.1 LIMITATIONS

While our study reveals consistent unlearning traces across multiple LLM families, unlearning methods, and datasets, several limitations remain. First, restricted by computational resources, it is unclear whether the same degree of trace persistence holds for even larger LLMs. Second, our analysis is limited to text-only LLMs. It remains an open question whether similar unlearning traces persist in multi-modal foundation models (Zhang et al., 2024g) that jointly model text with vision, speech, or other modalities (Zhang et al., 2023). Cross-modal alignment and shared latent representations may alter how traces propagate or remain detectable across modalities. (Zhang et al., 2025) Finally, we do not examine domain-specialized models (e.g., biomedical or legal LLMs), where training distributions may amplify or suppress residual traces.

N.2 BROAD IMPACT

This work exposes a new risk in machine unlearning: the detectability of unlearning traces. While unlearning aims to enhance safety, privacy, and compliance, such traces create opportunities for adversarial exploitation, reducing the cost of targeted relearning or jailbreaking in high-stakes domains like biosecurity and cybersecurity. At the same time, our findings offer constructive insights, understanding that unlearning traces can guide the development of more robust algorithms, stronger evaluation protocols, and clearer regulatory standards. We hope this work stimulates further research toward unlearning methods that reliably remove sensitive knowledge without leaving exploitable fingerprints.

O THE USE OF LLMs

This work makes limited use of LLMs. Specifically, LLMs were employed exclusively for grammar correction and stylistic polishing of the manuscript. They were not involved in research ideation, experimental design, data analysis, or the generation of any scientific content. All substantive contributions to the paper are solely attributable to the author.