

On the Impossibility of Retrain Equivalence in Machine Unlearning

Jiatong Yu, Yinghui He, Anirudh Goyal, Sanjeev Arora
Princeton Language and Intelligence, Princeton University
jiatongy@princeton.edu

Abstract

Machine unlearning seeks to selectively remove the “influence” of specific training data on a model’s outputs. The ideal goal is *Retrain Equivalence*—behavior identical to a model trained from scratch on only the retained data. This goal was originally formulated for models trained on *i.i.d.* data batches, but modern pipelines often involve multi-stage training, with each stage having a distinct data distribution and objective. Examples include LLM finetuning for alignment, reasoning ability, etc. Building on prior work that has shown retrain-equivalent behavior can be ill-defined due to forgeable datasets in mini-batch training, we study a complementary source of fragility that arises in multi-stage training. In an overparameterized linear regression model, we prove an exponential lower bound on *path-dependent divergence* under a family of *local and weakly local* unlearning objectives, unifying purely local gradient ascent and weakly local algorithms with retain-set regularization. Models trained on the same stages in different orders can be driven exponentially far apart during unlearning, and we show that this pointwise divergence implies violation of the distributional (ϵ, δ) -certified unlearning definition. We then empirically measure the same kind of path dependence in LLM post-training across Llama and Qwen models (1B–14B) with gradient ascent, NPO, and SimNPO algorithms. Models finetuned via different orderings of identical training stages diverge in behavior during unlearning, with the degradation in GSM8K accuracy after unlearning varying by over 20% across paths. We also observe that some learning paths consistently produce models that unlearn slowly while preserving higher retained utility—a phenomenon we term the *recency effect*—and that the fate of probability mass during unlearning (e.g., paraphrasing vs. alternative concepts) is likewise path-dependent. Taken together, these results quantitatively demonstrate that local, path-oblivious unlearning in multi-stage pipelines can be highly sensitive to training history, adding a multi-stage perspective to the impossibility results and motivating the search for new unlearning desiderata and evaluations¹.

1 Introduction

Large generative models are trained on multi-trillion-token datasets collected from diverse, partially licensed web sources (Penedo et al., 2023; Soldaini et al., 2024). Because this training data is internalized into model parameters, adversaries can extract sensitive and unsafe information (Carlini et al., 2021; Lehman et al., 2021). Legal requirements such as the *Right to be Forgotten* (Regulation, 2016) create a pressing need to remove private data from trained models.

Machine unlearning aims to remove the influence of a specified subset of the training dataset (“forget set”) from a trained model (Cao & Yang, 2015; Bourtoule et al., 2021). A prevailing desideratum of machine unlearning is *Retrain Equivalence* (RE): an unlearned model should make predictions indistinguishable from those of a model retrained from scratch on the

¹Code available at <https://princeton-qli.github.io/impossibility-unlearning>.

remaining data (called “retain set”). Full retraining trivially satisfies RE, but it violates the desideratum of *computational efficiency*: unlearning algorithms should have runtime independent of the retain data size. Practical approaches therefore attempt *local* unlearning algorithms, whose updates depend only on the gradient information computed on the forget set. Examples of local unlearning algorithms include gradient ascent, Negative Preference Optimization (Zhang et al., 2024) and Simple NPO (Fan et al., 2024b).

A key, yet often implicit, assumption behind RE is that model behaviors are determined only by the *multiset* of training data. This is not true for contemporary LLM pipelines that use *staged training* with distinct datasets per stage, such as instruction tuning (Wei et al., 2021), safety alignment (Ouyang et al., 2022; Bai et al., 2022), and reasoning enhancement (Dubey et al., 2024). While training order matters for model quality, it also highlights a simple fact: models exposed to different orders of the same stages reach different internal states and different behaviors. Thus, performing unlearning using just the forget dataset gets problematic, since the final result of unlearning could be dependent upon the (often unknown) learning path. The current definition of retrain equivalence ignores this possibility. This mismatch is the starting point of our work.

We emphasize that, in practice, a trained model is released as a *single* artifact, but the multi-stage path that produced it is usually hidden from the downstream users, auditors, or regulators who later perform or evaluate unlearning. Comparing models trained on different orderings of the same stages should therefore be read as a controlled diagnostic of this provenance uncertainty, not as a literal scenario where multiple model variants are released simultaneously. Moreover, current local unlearning algorithms are *path-oblivious*: their updates depend only on the final parameters and the forget set, so they cannot condition on the training path even when it is known.

Main Contributions. Can a path-oblivious, local unlearning rule make models trained on the same data in different orders both match a single retrained target? Building on prior work showing RE can be ill-defined due to forgeable datasets, we focus on *multi-stage* pipelines. Throughout, our impossibility claims concern *local and weakly local* path-oblivious algorithms (Definition 2.2), the predominant class for LLM-scale unlearning; they do not extend to every conceivable unlearning method. Our contributions:

- **Theory quantifying path-dependent divergence under staged training.** In over-parameterized linear regression with staged training, applying the same gradient-based unlearning rule to models trained in different orders causes predictions to diverge *exponentially* (Theorem 3.1), so RE is guaranteed to be violated within a calculable number of iterations for at least one training history (Corollary 1). This holds across a family of objectives including gradient ascent and retain-regularized methods (Maini et al., 2024), and extends to (ϵ, δ) -certified unlearning (Remark 1; Appendix C).
- **Empirical measurements of path dependence in LLM post-training.** We finetune Llama (Dubey et al., 2024) and Qwen (Team, 2024) models (1B–14B) with four orderings of four datasets, then unlearn one dataset using GA, NPO (Zhang et al., 2024), and SimNPO (Fan et al., 2024b). Path-dependent divergence appears in all cases, along with novel phenomena: the *recency effect* (Figure 1) and path-dependent *superficial forgetting* (Section 4.3).

1.1 Related Work

Certified Unlearning Certified methods provide formal (ϵ, δ) -guarantees (Sekhari et al., 2021), under assumptions such as convexity (Basaran et al., 2025; Warnecke et al., 2021; Guo et al., 2019; Allouah et al., 2024), smooth non-convex losses with noisy gradient steps (Chourasia & Shah, 2023; Chien et al., 2024; Mu & Klabjan, 2024), or gradient/model clipping (Koloskova et al., 2025); see also Wei et al. (2024). These methods are generally inapplicable to efficient LLM unlearning.




Figure 1: **History dependence of gradient ascent unlearning.** Each panel shows the unlearning process for four models finetuned from the same base LLM. Each of the four curves corresponds to a base model fine-tuned on the same four datasets, but with the unlearn set introduced at a different position ($p \in \{1, 2, 3, 4\}$) in the training sequence, as detailed in Section 4.1. The y-axis tracks the log likelihood of the responses being unlearned; a steeper decline indicates faster forgetting. Different values of p lead to very different outcomes. The red curve ($p = 4$) represents the case where unlearning immediately follows learning of the forget set, and here unlearning is slowest (**recency effect**, see Section 4.2).

Local Unlearning Algorithms for LLMs The most common efficient baseline is gradient ascent (GA) on the forget set (Maini et al., 2024; Jang et al., 2022; Zhang et al., 2024), often augmented with retain-set regularization (GA+Retain) (Yao et al., 2024; Liu et al., 2022; Maini et al., 2024; Li et al., 2024) or modified losses such as NPO (Zhang et al., 2024) and SimNPO (Fan et al., 2024b). LLM unlearning is broader than GA-style objectives: other families include representation-steering methods such as RMU (Li et al., 2024), model-editing and distillation-based removal (Meng et al., 2022; Eldan & Russinovich, 2023), and preference-based suppression (Rafailov et al., 2023). Our strongest claims concern local and weakly local path-oblivious methods and do not automatically extend to these families. We show that a unified family covering GA and GA+Retain exhibits exponential path-dependent divergence in a linear model (Section 3); NPO, SimNPO, and GA+KL also diverge empirically (Section 4; Appendix E).

Ambiguities in Unlearning Evaluation and Definition GA-based methods fail under adversarial attacks (Hayes et al., 2025; Pawelczyk et al., 2024; Hu et al., 2024), and forgetting outcomes depend on properties of the forget data (Baluta et al., 2024; Fan et al., 2024a; Barbulescu & Triantafillou, 2024; Sepahvand et al., 2025), complicating evaluation. Works adopting RE as the gold standard measure distributional distance to the retrained ideal (Triantafillou et al., 2024; Pawelczyk et al., 2023; Hayes et al., 2025), while Thudi et al. (2022) show RE can be trivially satisfied when distinct datasets yield identical models (Kong et al., 2022). The community has consequently been moving beyond RE as the sole metric (Shi et al., 2024; Maini et al., 2024); modern behavioral benchmarks such as MUSE (Shi et al., 2024) and WMDP (Li et al., 2024) are natural settings for future evaluations of path dependence. Our work adds to this line by quantifying how local, path-oblivious unlearning amplifies history dependence in multi-stage pipelines.

Data Ordering and Training Trajectory Effects Training order can significantly affect model behavior and can be exploited adversarially (Shumailov et al., 2021); variability has also been noted in *i.i.d.* batch settings. We study the multi-stage case, where stage orderings induce variability more broadly.

2 Preliminaries

Consider a model θ trained on dataset $D = D_f \cup D_r$ (forget and retain sets). Retrain Equivalence asks an unlearning algorithm $\mathcal{U}(\cdot)$ to produce θ_u that behaves identically to the retrained model θ_r on any test set X_{test} .

Definition 2.1 (Retrain Equivalence (RE)). Let $\text{Pred}(\theta, x) \in \mathbb{R}^h$ denote the predictions of model θ on a test point drawn from $X_{\text{test}} = \{x_i\}_{i=1}^m$. For a pair of models (θ_u, θ_r) , define RE distance as the average distance between the predictions of the two models measured on X_{test} , *i.e.*,

$$d(\theta_u, \theta_r) := \frac{1}{m} \sum_{i=1}^m \|\text{Pred}(\theta_u, x_i) - \text{Pred}(\theta_r, x_i)\|_2^2$$

Let θ_u be the outcome of an unlearning algorithm on forget set D_f , and θ_r be the model retrained from scratch on D_r . For some $\varepsilon > 0$, Retrain Equivalence is satisfied when $d(\theta_u, \theta_r) \leq \varepsilon$.

Another desideratum is efficiency: $T_{\text{unlearn}} = o(T_{\text{retrain}})$. For LLMs, even revisiting a small fraction of the retain set is prohibitive, motivating local unlearning.

Definition 2.2 (Local Unlearning). An unlearning algorithm $\mathcal{U}(\cdot, D_f)$ is considered local if it only requires gradient information computed on the forget set D_f .

Examples of local unlearning algorithms include gradient ascent on the forget set, Negative Preference Optimization (Zhang et al., 2024), and SimNPO (Fan et al., 2024b). We call an algorithm *weakly local* if it additionally uses limited retain-set gradient information, as in GA+Retain or GA+KL (Maini et al., 2024; Yao et al., 2024); *non-local* methods instead require substantial retain-set access (e.g., retraining, distillation, or certified unlearning). Locality matters in practice because the retain corpus is often unavailable, and even revisiting a fraction of it is prohibitive at LLM scale.

3 Impossibility Theorem in Overparametrized Linear Models

To cleanly isolate the mechanism behind history-dependent unlearning, we first study a simpler setting of overparametrized linear regression with staged training and local unlearning. Overparameterization admits many directions with negligible effect on predictions (Bartlett et al., 2020), giving multi-stage training room to steer solutions into order-dependent internal states. Although simplified, such models capture aspects of implicit bias in deep networks (Soudry et al., 2018; Belkin et al., 2019). Insights from our linear analysis may also be transferable and complementary to existing bounds for SGD-like training in more general settings. In this section, we show that applying the same local-unlearning rule to models trained in different orders causes predictions to diverge exponentially. Consequently, they cannot all simultaneously satisfy Retrain Equivalence: at most one can make predictions close to the retrained baseline.

Two-phase staged training. Let $\theta \in \mathbb{R}^d$ be the parameter vector. We assume two training datasets $S_A = (X_A, y_A)$ and $S_B = (X_B, y_B)$, where X_A and X_B are independently drawn from different, continuous distributions. Each training stage fits a different dataset with ridge regression loss regularized towards the previous iterate

$$\mathcal{L}(\theta; X, y, \theta_{\text{prev}}, \mu) = \|X\theta - y\|_2^2 + \mu \|\theta - \theta_{\text{prev}}\|_2^2$$

Consider models θ_{AB} and θ_{BA} initialized at $\theta_0 = 0$ and trained from opposite stage orders. Fix $\mu_A, \mu_B > 0$, the proximal updates yield

$$\begin{aligned} \theta_{AB} &= (X_B^\top X_B + \mu_B I)^{-1} (X_B^\top y_B + \mu_B \theta_A) \\ \theta_{BA} &= (X_A^\top X_A + \mu_A I)^{-1} (X_A^\top y_A + \mu_A \theta_B) \end{aligned} \tag{1}$$

Forget set and parameterized unlearning. Fix a forget set $S_F = (X_F, y_F)$ with $X_F \in \mathbb{R}^{k_F \times d}$ and a retain set $S_R = (X_R, y_R)$ with $X_R \in \mathbb{R}^{k_R \times d}$. For a retain strength $\alpha \geq 0$, define the parameterized unlearning objective

$$L_\alpha(\theta) := -\ell(S_F, \theta) + \alpha \ell(S_R, \theta), \quad \ell(S, \theta) := \frac{1}{k} \|X\theta - y\|_2^2, \quad (2)$$

and unlearn by gradient descent on L_α with step size $\eta > 0$:

$$\theta^{(t+1)} = \theta^{(t)} - \eta \nabla_\theta L_\alpha(\theta^{(t)}). \quad (3)$$

When $\alpha = 0$ this reduces to gradient ascent on the forget loss (purely local unlearning); when $\alpha > 0$ it adds a retain-set regularizer, recovering the widely used GA+Retain (Gradient Difference) objective (Maini et al., 2024; Yao et al., 2024) at $\alpha = 1$.

Define the empirical second-moment matrices and their α -weighted difference:

$$\Sigma_F := \frac{1}{k_F} X_F^\top X_F, \quad \Sigma_R := \frac{1}{k_R} X_R^\top X_R, \quad \Sigma_\alpha := \Sigma_F - \alpha \Sigma_R. \quad (4)$$

Let $M_\alpha := 2\eta \Sigma_\alpha$. Denote the eigendecomposition of Σ_α by $\Sigma_\alpha = \sum_{j=1}^d \lambda_j v_j v_j^\top$. Define the *expanding subspace* $E_+^\alpha := \text{span}\{v_j : \lambda_j > 0\}$ with dimension $r_\alpha := \dim(E_+^\alpha)$ and orthogonal projector P_α .

3.1 Main Theorem

Our main result shows that the parameterized unlearning objective (2) amplifies the initial discrepancy between two models trained on the same data but in different orders. During unlearning, the behaviors of the two models on a generic test set diverge exponentially fast with the number of unlearning steps. Consequently, a single retrain-equivalent target (Definition 2.1) cannot be reached from both histories.

Theorem 3.1 (Main Result). *Let $\theta_{AB}, \theta_{BA} \in \mathbb{R}^d$ be the outcome of two-stage ridge-regression training, given by Equation (1), and set $\Delta\theta_0 := \theta_{AB} - \theta_{BA}$. Fix forget and retain sets S_F, S_R as above, and perform unlearning via (3) with retain strength $\alpha \geq 0$ and step size $\eta > 0$. Assume*

- (i) Σ_α has at least one positive eigenvalue (i.e., $r_\alpha \geq 1$); and
- (ii) the step size satisfies $\eta \leq -1/\lambda_j$ for every negative eigenvalue λ_j of Σ_α ; equivalently, if $\lambda_{\min}(\Sigma_\alpha) < 0$ then $\eta \leq 1/|\lambda_{\min}(\Sigma_\alpha)|$ (and the condition is vacuous when $\Sigma_\alpha \succeq 0$).

Fix a test set $X_{\text{test}} \in \mathbb{R}^{m \times d}$ with i.i.d. rows from a continuous distribution, independent of the training data, with $m \geq r_\alpha$. Define

$$\rho_\star := \frac{\langle P_\alpha \Delta\theta_0, M_\alpha P_\alpha \Delta\theta_0 \rangle}{\|P_\alpha \Delta\theta_0\|_2^2}, \quad \sigma_\alpha := \sigma_{\min}(X_{\text{test}} P_\alpha). \quad (5)$$

Then $\rho_\star > 0$ and $\sigma_\alpha > 0$ almost surely, and there exists a constant $t_0 \in \mathbb{Z}$ such that for all $t \geq \max(t_0, 0)$, the RE distance in Def. 2.1 measured on X_{test} satisfies

$$d(\theta_{AB}^{(t)}, \theta_{BA}^{(t)}) \geq \frac{\sigma_\alpha^2 \|P_\alpha \Delta\theta_0\|_2^2}{4m} (1 + \rho_\star)^{2t}. \quad (6)$$

The next corollary makes the RE violation explicit. In practice, t_0 is negative (Appendix D), so the exponential regime dominates from the first unlearning step.

Corollary 1 (Violation of Retrain Equivalence). *Theorem 3.1 gives an upper bound on the number of unlearning iterations before Retrain Equivalence is guaranteed to be violated. Fix a target RE tolerance $\varepsilon > 0$ as in Definition 2.1, then for any unlearning iteration $t^\star > 0$ that satisfies*

$$t^\star \geq \max \left(\left\lceil \frac{\log(16m\varepsilon) - \log(\sigma_\alpha^2 \|P_\alpha \Delta\theta_0\|_2^2)}{2 \log(1 + \rho_\star)} \right\rceil, t_0 \right),$$

Theorem 3.1 implies that either $d(\theta_{AB}^{(t^\star)}, \theta_r) > \varepsilon$ or $d(\theta_{BA}^{(t^\star)}, \theta_r) > \varepsilon$, therefore the two models cannot both satisfy retrain equivalence.




Figure 2: **Illustration of the training setup in Section 4.1.** Each base model is fine-tuned four times, varying only the safety stage’s position ($p \in \{1, 2, 3, 4\}$); the resulting models are then unlearned on the same safety dataset.

3.2 Proof Sketch

Full proof is in Appendix C. Let $\Delta\theta_t := \theta_{AB}^{(t)} - \theta_{BA}^{(t)}$ denote the weight difference after t unlearning steps.

Lemma 3.2 (*Evolution of weight difference*). *Fix forget and retain sets S_F, S_R , retain strength $\alpha \geq 0$, and step size $\eta > 0$. The weight difference $\Delta\theta_t := \theta_{AB}^{(t)} - \theta_{BA}^{(t)}$ evolves as $\Delta\theta_t = (I + M_\alpha)^t \Delta\theta_0$, where $M_\alpha = 2\eta \Sigma_\alpha = 2\eta (\Sigma_F - \alpha \Sigma_R)$.*

Label-dependent gradient terms cancel in the difference, leaving a linear recurrence. Decomposing $\Delta\theta_0 = P_\alpha \Delta\theta_0 + (I - P_\alpha) \Delta\theta_0$, the component in E_+^α grows exponentially. The complement stays bounded precisely because of the step-size condition (ii): on a negative eigendirection $\lambda_j < 0$ the factor is $1 + 2\eta\lambda_j$, and $\eta \leq -1/\lambda_j$ ensures $|1 + 2\eta\lambda_j| \leq 1$ (otherwise that direction would expand too). This gives

$$\|X_{\text{test}} \Delta\theta_t\|_2 \geq \|X_{\text{test}} (I + M_\alpha)^t P_\alpha \Delta\theta_0\|_2 - C_0,$$

where C_0 is a constant. Lower-bounding the growing term via the Rayleigh quotient ρ_* , for $t \geq t_0$ (chosen so the constant is at most half the growing term):

$$\frac{1}{m} \|X_{\text{test}} \Delta\theta_t\|_2 \geq \frac{\sigma_\alpha \|P_\alpha \Delta\theta_0\|_2}{2m} (1 + \rho_*)^t \quad (7)$$

Lemmas C.2 and C.4 (Appendix) verify $P_\alpha \Delta\theta_0 \neq 0$ and $\sigma_\alpha > 0$ almost surely. Squaring gives Theorem 3.1.

Appendix E validates that GA+Retain and GA+KL ($\alpha > 0$) exhibit the same phenomena in LLM experiments.

Remark 1 (Extension to Distributional RE). *Theorem 3.1 also violates the distributional (ϵ, δ) -certified unlearning framework (Sekhari et al., 2021): generalizing to N stages, a counting argument shows (ϵ, δ) -RE requires $\delta \geq 1 - 1/N$ (Corollary 2, Appendix C).*

4 Experiments

Section 3 indicates why Retrain Equivalence (RE) is in general impossible for gradient ascent on simple linear models subject to staged training. However, it is hard to perform analogous theoretical analysis for LLMs powered by deep neural networks. In addition, a practical concern is whether this path-dependent behavior emerges within the first few gradient steps, as long-running unlearning often leads to model collapse and is avoided in practice.

This section closes this gap by empirically showing that even a small number of unlearning updates produce path-dependent divergence across common algorithms. This indicates that Retrain Equivalence is also infeasible in more complex settings. We highlight that the purpose of the following experiments is *not* to identify which learning order produces models closer to the retrained ideal after unlearning. Rather, by demonstrating that models trained on the same data diverge within just a few unlearning steps, we argue the Retrain Equivalence ideal is inherently ill-posed: as long as unlearning algorithms remain path-oblivious, their success depends on factors outside of their algorithmic design. We do not

Table 1: **Forget score and utility scores after one epoch of unlearning with different methods on models of different sizes from two families.** See Figure 3 for definition of p . Definitions of scores are: (1) forget score is the amount of decrease of the average log probability of explicitly unlearned safe response “Sorry, I can not assist you” and 20 similar rephrasings such as “I’m afraid I can’t” and “I’m unable to assist”. A higher forget score means larger decrease in probability mass of safe responses. (2) The TOFU utility score is defined as $1 - |\Delta\pi_\theta(y|x)|$, where $\Delta\pi_\theta(y|x)$ represents the difference in the average log-likelihood of ground-truth TOFU answers before and after unlearning. Score closer to 1 indicates better utility preservation. (3) Math utility score is measured by change in GSM8K test accuracy before and after unlearning. Cell colors are a visual aid only: within each column, greener shades mark values closer to the best and redder shades closer to the worst, following the direction of each metric’s arrow; no statistical threshold is implied. The $p = 4$ data show slower unlearning (“recency effect”).

Path (p)	Llama1B			Llama8B			Llama13B			Qwen1B			Qwen14B		
	GA	NPO	SimNPO	GA	NPO	SimNPO	GA	NPO	SimNPO	GA	NPO	SimNPO	GA	NPO	SimNPO
<i>Forget Score</i> ↑															
1	7.893	1.934	1.05	14.851	2.184	1.788	0.475	0.469	0.329	5.33	1.761	0.738	0.119	0.117	0.24
2	5.341	0.66	1.471	7.203	2.212	1.652	0.33	0.328	0.343	10.082	3.997	3.559	1.368	1.021	0.64
3	6.346	1.215	1.465	8.217	3.445	1.734	0.621	0.616	0.531	8.143	3.066	2.943	2.704	2.313	1.968
4	0.178	0.046	0.293	2.92	0.683	0.668	0.146	0.148	0.196	0.547	0.644	1.893	-0.855	-0.855	0.991
<i>Utility Score (TOFU)</i> ↑															
1	-0.705	0.567	0.562	-1.496	0.451	0.585	0.995	0.996	0.993	-0.559	0.408	0.827	0.999	0.998	0.697
2	-0.27	0.499	0.326	-0.148	0.446	0.581	0.986	0.986	0.977	-1.924	-0.301	-0.22	0.547	0.631	0.749
3	-1.916	0.18	0.201	-2.073	0.223	0.488	0.915	0.914	0.911	-2.366	0.148	0.31	0.042	0.207	0.304
4	0.622	0.742	0.651	0.333	0.932	0.741	0.991	0.991	0.978	-0.566	0.405	-0.629	0.445	0.671	0.796
<i>Utility Score (Math)</i> ↑															
1	0.387	0.394	-0.07	-0.052	0.02	-0.028	0.121	0.133	0.116	0.018	-0.161	-0.32	0.059	0.034	-0.099
2	-0.079	-0.198	-0.221	-0.091	-0.056	0.01	0.193	0.188	0.165	-0.006	-0.326	-0.39	-0.075	-0.046	-0.149
3	0	0	0.041	-0.143	-0.143	0.029	0.187	0.22	0.215	-0.181	0.292	-0.113	-0.225	-0.235	-0.268
4	0.001	-0.053	-0.035	-0.022	0.062	0.019	0.362	0.371	0.373	-0.554	-0.54	-0.084	0.007	0.035	0.172

claim that data ordering is the only—or even the dominant—source of error in approximate unlearning; the experiments isolate path dependence as a distinct failure mode arising from multi-stage training.

4.1 Setup

Training Stages. To emulate LLM post-training, we design four finetuning stages:

- **Instruction Tuning** S_{inst} . 4k synthetic instruction–response pairs from INSTRUCT-SKILLMIX (Kaur et al., 2025); 10 epochs.
- **TOFU Fictitious Knowledge** S_{tofu} . 4k question–answer pairs about fictitious authors from the TOFU dataset (Maini et al., 2024); 4 epochs.
- **Mathematical Reasoning** S_{math} . 8k GSM8K problems (Cobbe et al., 2021) rewritten by GPT-4o with step-by-step reasoning traces; 2 epochs.
- **Safety Behavior (Unlearn Set)** S_U . 4.5k GPT-4o-generated refusal responses to unsafe questions from SORRY-BENCH (Xie et al., 2024); 2 epochs. This is the dataset unlearned after all four stages complete².

Finetuning and Unlearning Setup. For each pretrained base model we create four finetuned models, each with a unique ordering of the above training stages (see Figure 2). The relative order of the first three training stages (*i.e.*, S_{inst} , S_{tofu} , and S_{math}) is fixed, and the safety dataset S_U occurs at four different positions. The finetuned models then go through an identical unlearning stage on the safety dataset S_U , using following *local* unlearning algorithms. Given a question–answer pair (x, y) , let $\pi_\theta(y|x)$ denote a language model’s prediction probability. We use:

²Unlearning S_U is conceptually a form of jailbreak tuning, but our goal is to study history dependence, not jailbreaking. See Appendix A for Ethics Statement.




Figure 3: **Change in forget quality and retained utility in three models during unlearning of safe responses “Sorry, I can not assist you...” on the safety dataset S_U .** Each of the four curves corresponds to a base model fine-tuned on the same four datasets, but with the unlearn set S_U introduced at a different position ($p \in \{1, 2, 3, 4\}$) in the training sequence, as detailed in Section 4.1. For each base model, the left column reports its four finetuned models’ average log likelihood of the label “Sorry, I can not assist you” on questions from S_U . The right column reports average retained utility, measured by the average log probability of ground truth responses in the TOFU dataset S_{tofu} , which were seen during training. As shown, local unlearning is fundamentally path-dependent. Models with different training histories—even those starting from similar performance points—diverge in unpredictable ways within a few iterations.

- **Gradient Ascent (GA)** The GA loss is given by $\mathcal{L}_{\text{GA}} = \mathbb{E}_{(x,y) \in D_f} [\log \pi_{\theta}(y|x)]$.
- **Negative Preference Optimization (NPO)** For some hyperparameter $\beta > 0$, with π_{ref} denoting the reference model (initial parameters prior to unlearning),

$$\mathcal{L}_{\text{NPO}} = \mathbb{E}_{(x,y) \in D_f} \left[-\frac{2}{\beta} \log \sigma \left(-\beta \log \frac{\pi_{\theta}(y|x)}{\pi_{\text{ref}}(y|x)} \right) \right],$$

where $\sigma(t) = 1/(1 + e^{-t})$ is the sigmoid function.

- **Simple NPO (SimNPO)** SimNPO removes the reliance on the reference model:

$$\mathcal{L}_{\text{SimNPO}} = \mathbb{E}_{(x,y) \in D_f} \left[-\frac{2}{\beta} \log \sigma(-\beta \log \pi_{\theta}(y|x)) \right].$$

We experiment with the following pretrained base models: Qwen2.5-1.5B, Qwen2.5-14B, Llama3.2-1B, Llama3.1-8B, and Llama2-13B (Team, 2024; Dubey et al., 2024). We use a learning rate of 1×10^{-5} to finetune and unlearn all models, except for Llama2-13B and Qwen2.5-14B where we use $\text{lr} = 5 \times 10^{-6}$ for unlearning. During finetuning, LR schedulers and optimizers are *re-initialized* in each of the four training stages. Dataset examples and the full training configurations are in Appendix F.

4.2 Results

Local unlearning is path-dependent across algorithms and models. Figure 3 and Table 1 show that unlearning outcomes are consistently sensitive to training order. Different paths lead to GSM8K accuracy differences of over 20% for Llama13B and Qwen14B.

Prior work attributes unlearning dynamics primarily to the target’s initial likelihood (Ren & Sutherland, 2024), but learning path is a confounding variable: for Llama3.2-1B and Llama3.1-8B, models with near-identical initial predictions diverge during unlearning.




Figure 4: **Unlearning outcomes on Llama-3B under different learning sequences.** Each panel shows the evolution of log probabilities during unlearning, where U (red) is the compliance (unsafe) phrasing selected as the unlearning target, C (blue) is the alternative compliance phrasing, and R (orange) is the refusal (safe) response. *The type of forgetting during unlearning is path-dependent:* sequences θ_1 , θ_2 , θ_5 , and θ_6 yield superficial forgetting (only U drops), while θ_3 and θ_4 yield deep forgetting (both C and U drop below R).

A common hope is that fewer steps or smaller learning rates suffice to approximate RE. Figure 3 shows otherwise: divergence appears within a few steps even at $\text{lr} = 5\text{e-}6$ for Qwen14B and Llama13B. This persists across learning rates and schedulers (Appendix J).

Recency Effect: Unlearning is hardest when information is fresh. In all except one cases, we find that *unlearning proceeds slowest when it immediately follows the corresponding learning* (see Table 1). Even in the exception case of Qwen2.5-1.5B paired with SimNPO unlearning, we see that the forget score of $p = 4$ model is still lower than average. We call it a *recency effect*, because for this path there is no intermediate finetuning of retained data between the learning and unlearning of the forget set S_U . Figure 1 further shows that recency effect occurs throughout the unlearning process, including the very first few gradient updates.

Table 1 shows that slower unlearning is often accompanied by higher retained utilities. While this alludes to a fundamental trade-off between forget quality and retained utilities, we highlight that path-dependence dynamics is often more complex and unpredictable. In Appendix J, we show that recency effect occurs across learning rates and LR schedulers—but increasing the learning rate for Qwen2.5-1.5B models causes the $p = 4$ model to suffer the most severe utility degradation while being slowest to unlearn.

4.3 Investigating the history dependency of superficial unlearning

A central question in machine unlearning is whether forgetting is superficial or deep (Wu et al., 2024; Jang et al., 2025; Kim et al., 2025; Yamashita et al., 2025). In this work, we define *superficial forgetting* as the suppression of predictions for one specific phrasing of an undesired response, while leaving semantically equivalent alternatives intact. In contrast, *deep forgetting* in our work refers to the broader suppression of all reasonable paraphrases. Superficial forgetting is thus often undesirable.

Since Section 4.2 establishes that unlearning outcomes are path-dependent, a natural next question is whether this also holds for the depth of forgetting. To investigate this question, we take Llama-3.2-3B as a case study, and conduct a more controlled fine-tuning→unlearning experiment. We curate a synthetic dataset with 500 unsafe prompts,

each paired with two compliance responses (i.e., unsafe responses) with different phrasings, denoted as C and U , and one refusal response (i.e., safe response), denoted as R .

During the initial learning phase, models are trained on all three response types under 6 different permutations of training stages (θ_1 – θ_6 , as defined in Section I.2). We then unlearn one unsafe phrasing U via gradient ascent, and track the log probabilities of all three responses (i.e., R , C , and U) across unlearning epochs. We concentrate on whether unlearning the specific compliance response U reduces the likelihood of its semantically similar counterpart C . If the log probability of C falls in tandem with U , we interpret this as evidence of deep forgetting. If instead the probability of C does not decrease with U , the model exhibits superficial forgetting.

In Figure 4, the type of forgetting diverges after epoch 5 and is clear by epoch 10: sequences θ_1 , θ_2 , θ_5 , and θ_6 show *superficial forgetting* (only U declines), while θ_3 and θ_4 exhibit *deep forgetting*, with both U and C dropping below R . These results show that the depth of forgetting is also path-dependent, implying that the depth of unlearning is not determined by the algorithm alone, making it difficult to guarantee the complete removal of unsafe content in practical safety training.

5 Discussion and Conclusions

Our results add a multi-stage perspective to the growing body of work showing that Retrain Equivalence can be a fragile target for approximate unlearning. When models are trained in stages, local, path-oblivious unlearning rules amplify small differences in training history into exponentially diverging behaviors, both in our linear analysis and empirically in LLM post-training. This creates a practical tension between three desiderata for unlearning algorithms: (1) *path-independence*, (2) *retrain equivalence*, and (3) *locality*—in staged pipelines, insisting on all three simultaneously is unrealistic.

Is RE ill-posed, or merely hard to achieve? For local, path-oblivious unlearning, we posit that RE is ill-posed rather than merely difficult: models trained on different orderings of the same stages diverge under identical unlearning updates, so at most one of the combinatorially many training histories can approach the single retrained target (Corollary 1). Since the standard problem specification omits the training path—the algorithm sees only the final parameters and the forget set—success is decided by information the algorithm can never access.

Alternative targets. This tension forces a pragmatic choice: forgo strict locality and accept the cost of non-local methods, or relax Retrain Equivalence as the sole gold standard in favor of richer, behavioral, task-specific evaluations (Shi et al., 2024; Thudi et al., 2022; Maini et al., 2024); benchmarks such as MUSE (Shi et al., 2024) and WMDP (Li et al., 2024) are promising examples.

Open questions and future work. The recency effect persists across learning rates and LR schedulers (Appendix J), but whether it is primarily an optimization or a representation phenomenon remains open. Our four-stage pipeline is one controlled post-training design: a homogeneous-data control—splitting a single dataset into per-stage subsets and unlearning one split—would isolate ordering effects from task heterogeneity; replacing the unlearning stage with an enhancing task would test whether path dependence is specific to unlearning; and messier pipelines deserve study. Finally, can path-induced behavior be distinguished from algorithm-induced behavior, and does the recency effect arise in RL post-training, which likewise applies gradient ascent to non-desired answers (Rafailov et al., 2023; Shao et al., 2024)?

6 Reproducibility Statement

We provide the training configurations, hyperparameters, training cost in Appendix F. We provide the prompts used to curate synthetic datasets in Section 4 in Appendix G. We provide details on the evaluation metrics for experiments in Section 4.1 in Appendix H.

References

- Youssef Allouah, Joshua Kazdan, Rachid Guerraoui, and Sanmi Koyejo. The utility and complexity of in-and out-of-distribution machine unlearning. *arXiv preprint arXiv:2412.09119*, 2024.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.
- Teodora Baluta, Pascal Lamblin, Daniel Tarlow, Fabian Pedregosa, and Gintare Karolina Dziugaite. Unlearning in-vs. out-of-distribution data in llms under gradient-based method. *arXiv preprint arXiv:2411.04388*, 2024.
- George-Octavian Barbulescu and Peter Triantafillou. To each (textual sequence) its own: Improving memorized-data unlearning in large language models. *arXiv preprint arXiv:2405.03097*, 2024.
- Peter L Bartlett, Philip M Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070, 2020.
- Umit Yigit Basaran, Sk Miraj Ahmed, Amit Roy-Chowdhury, and Basak Guler. A certified unlearning approach without access to source data. *arXiv preprint arXiv:2506.06486*, 2025.
- Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019.
- 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.
- Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data from large language models. In *30th USENIX security symposium (USENIX Security 21)*, pp. 2633–2650, 2021.
- Eli Chien, Haoyu Wang, Ziang Chen, and Pan Li. Langevin unlearning: A new perspective of noisy gradient descent for machine unlearning. *Advances in neural information processing systems*, 37:79666–79703, 2024.
- Rishav Chourasia and Neil Shah. Forget unlearning: Towards true data-deletion in machine learning. In *International conference on machine learning*, pp. 6028–6073. PMLR, 2023.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukas Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv e-prints*, pp. arXiv–2407, 2024.
- Ronen Eldan and Mark Russinovich. Who’s harry potter? approximate unlearning in LLMs. *arXiv preprint arXiv:2310.02238*, 2023.
- Chongyu Fan, Jiancheng Liu, Alfred Hero, and Sijia Liu. Challenging forgets: Unveiling the worst-case forget sets in machine unlearning. In *European Conference on Computer Vision*, pp. 278–297. Springer, 2024a.

- 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*, 2024b.
- Government of Canada. Bill c-27: Digital charter implementation act, 2022 – consumer privacy protection act (CPPA), 2022. URL <https://www.parl.ca/DocumentViewer/en/44-1/bill/C-27/first-reading>. Proposed legislation to modernize Canadian privacy law, including rights to erasure and AI oversight.
- Chuan Guo, Tom Goldstein, Awni Hannun, and Laurens Van Der Maaten. Certified data removal from machine learning models. *arXiv preprint arXiv:1911.03030*, 2019.
- Jamie Hayes, Ilia Shumailov, Eleni Triantafillou, Amr Khalifa, and Nicolas Papernot. Inexact unlearning needs more careful evaluations to avoid a false sense of privacy. In *2025 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, pp. 497–519. IEEE, 2025.
- Shengyuan Hu, Yiwei Fu, Zhiwei Steven Wu, and Virginia Smith. Unlearning or obfuscating? jogging the memory of unlearned llms via benign relearning. *arXiv preprint arXiv:2406.13356*, 2024.
- 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.
- Yeonwoo Jang, Shariqah Hossain, Ashwin Sreevatsa, and Diogo Cruz. Prompt attacks reveal superficial knowledge removal in unlearning methods, 2025. URL <https://arxiv.org/abs/2506.10236>.
- Simran Kaur, Simon Park, Anirudh Goyal, and Sanjeev Arora. Instruct-skillmix: A powerful pipeline for LLM instruction tuning. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=44z7HL4mfx>.
- Yongwoo Kim, Sungmin Cha, and Donghyun Kim. Are we truly forgetting? a critical re-examination of machine unlearning evaluation protocols, 2025. URL <https://arxiv.org/abs/2503.06991>.
- Anastasia Koloskova, Youssef Allouah, Animesh Jha, Rachid Guerraoui, and Sanmi Koyejo. Certified unlearning for neural networks. *arXiv preprint arXiv:2506.06985*, 2025.
- Zhifeng Kong, Amrita Roy Chowdhury, and Kamalika Chaudhuri. Forgeability and membership inference attacks. In *Proceedings of the 15th ACM workshop on artificial intelligence and security*, pp. 25–31, 2022.
- Eric Lehman, Sarthak Jain, Karl Pichotta, Yoav Goldberg, and Byron C Wallace. Does bert pretrained on clinical notes reveal sensitive data? *arXiv preprint arXiv:2104.07762*, 2021.
- Nathaniel Li, Alexander Pan, Anjali Gopal, Summer Yue, Daniel Berrios, Alice Gatti, Justin D Li, Ann-Kathrin Dombrowski, Shashwat Goel, Long Phan, et al. The wmdp benchmark: Measuring and reducing malicious use with unlearning. *arXiv preprint arXiv:2403.03218*, 2024.
- Bo Liu, Qiang Liu, and Peter Stone. Continual learning and private unlearning. In *Conference on Lifelong Learning Agents*, pp. 243–254. PMLR, 2022.
- Pratyush Maini, Zhili Feng, Avi Schwarzschild, Zachary C Lipton, and J Zico Kolter. Tofu: A task of fictitious unlearning for llms. *arXiv preprint arXiv:2401.06121*, 2024.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in GPT. *Advances in Neural Information Processing Systems*, 35:17359–17372, 2022.
- Siqiao Mu and Diego Klabjan. Rewind-to-delete: Certified machine unlearning for nonconvex functions. *arXiv preprint arXiv:2409.09778*, 2024.

- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- Martin Pawelczyk, Seth Neel, and Himabindu Lakkaraju. In-context unlearning: Language models as few shot unlearners. *arXiv preprint arXiv:2310.07579*, 2023.
- Martin Pawelczyk, Jimmy Z Di, Yiwei Lu, Ayush Sekhari, Gautam Kamath, and Seth Neel. Machine unlearning fails to remove data poisoning attacks. *arXiv preprint arXiv:2406.17216*, 2024.
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. The refinedweb dataset for falcon llm: outperforming curated corpora with web data, and web data only. *arXiv preprint arXiv:2306.01116*, 2023.
- 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. *Advances in neural information processing systems*, 36:53728–53741, 2023.
- Protection Regulation. Regulation (eu) 2016/679 of the european parliament and of the council. *Regulation (eu)*, 679(2016):10–13, 2016.
- Yi Ren and Danica J Sutherland. Learning dynamics of llm finetuning. *arXiv preprint arXiv:2407.10490*, 2024.
- Ayush Sekhari, Jayadev Acharya, Gautam Kamath, and Ananda Theertha Suresh. Remember what you want to forget: Algorithms for machine unlearning. *Advances in Neural Information Processing Systems*, 34:18075–18086, 2021.
- Nazanin Mohammadi Sepahvand, Anvith Thudi, Berivan Isik, Ashmita Bhattacharyya, Nicolas Papernot, Eleni Triantafillou, Daniel M Roy, and Gintare Karolina Dziugaite. Leveraging per-instance privacy for machine unlearning. *arXiv preprint arXiv:2505.18786*, 2025.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- 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.
- Ilia Shumailov, Zakhar Shumaylov, Dmitry Kazhdan, Yiren Zhao, Nicolas Papernot, Murat A Erdogdu, and Ross J Anderson. Manipulating sgd with data ordering attacks. *Advances in Neural Information Processing Systems*, 34:18021–18032, 2021.
- Luca Soldaini, Rodney Kinney, Akshita Bhagia, Dustin Schwenk, David Atkinson, Russell Authur, Ben Bogin, Khyathi Chandu, Jennifer Dumas, Yanai Elazar, et al. Dolma: An open corpus of three trillion tokens for language model pretraining research. *arXiv preprint arXiv:2402.00159*, 2024.
- Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *Journal of Machine Learning Research*, 19(70):1–57, 2018.
- Qwen Team. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024.
- Anvith Thudi, Hengrui Jia, Ilia Shumailov, and Nicolas Papernot. On the necessity of auditable algorithmic definitions for machine unlearning. In *31st USENIX security symposium (USENIX Security 22)*, pp. 4007–4022, 2022.

- Eleni Triantafillou, Peter Kairouz, Fabian Pedregosa, Jamie Hayes, Meghdad Kurmanji, Kairan Zhao, Vincent Dumoulin, Julio Jacques Junior, Ioannis Mitliagkas, Jun Wan, et al. Are we making progress in unlearning? findings from the first neurips unlearning competition. *arXiv preprint arXiv:2406.09073*, 2024.
- UK Government. UK general data protection regulation (UK GDPR), 2021. URL <https://www.gov.uk/data-protection>. Post-Brexit adaptation of the EU GDPR, enforced by the UK Information Commissioner’s Office.
- Joaquin Vanschoren, Jan N. van Rijn, Bernd Bischl, and Luis Torgo. Openml: Networked science in machine learning. *SIGKDD Explorations*, 15(2):49–60, 2013. doi: 10.1145/2641190.2641198. URL <http://doi.acm.org/10.1145/2641190.2641198>.
- Alexander Warnecke, Lukas Pirch, Christian Wressnegger, and Konrad Rieck. Machine unlearning of features and labels. *arXiv preprint arXiv:2108.11577*, 2021.
- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*, 2021.
- Stanley Wei, Sadhika Malladi, Sanjeev Arora, and Amartya Sanyal. Provable unlearning in topic modeling and downstream tasks. *arXiv preprint arXiv:2411.12600*, 2024.
- Ruihan Wu, Chhavi Yadav, Russ Salakhutdinov, and Kamalika Chaudhuri. Evaluating deep unlearning in large language models, 2024. URL <https://arxiv.org/abs/2410.15153>.
- Tinghao Xie, Xiangyu Qi, Yi Zeng, Yangsibo Huang, Udari Madhushani Sehwag, Kaixuan Huang, Luxi He, Boyi Wei, Dacheng Li, Ying Sheng, et al. Sorry-bench: Systematically evaluating large language model safety refusal. *arXiv preprint arXiv:2406.14598*, 2024.
- Tomoya Yamashita, Akira Ito, Yuuki Yamanaka, Masanori Yamada, Takayuki Miura, and Toshiki Shibahara. Sparse-autoencoder-guided internal representation unlearning for large language models, 2025. URL <https://arxiv.org/abs/2509.15631>.
- Yuanshun Yao, Xiaojun Xu, and Yang Liu. Large language model unlearning. *Advances in Neural Information Processing Systems*, 37:105425–105475, 2024.
- Ruiqi Zhang, Licong Lin, Yu Bai, and Song Mei. Negative preference optimization: From catastrophic collapse to effective unlearning. *arXiv preprint arXiv:2404.05868*, 2024.

A Ethics Statement

This work studies machine unlearning and its limitations (*i.e.*, hardness to achieve Retrain Equivalence) when applied large scale neural networks including LLMs. Unlearning algorithms for LLMs are usually used for the removal of private, licensed, or unsafe information, mandated by legislative efforts such as General Data Protection Regulation (GDPR) in EU and UK (UK Government, 2021; Regulation, 2016) and Consumer Privacy Protection Act (CPPA) from Canada (Government of Canada, 2022). The thesis of this paper—that it is infeasible for computationally efficient, local unlearning algorithms to achieve Retrain Equivalence—should be useful for developing evaluations, benchmarks, and societal regulations of LLM safety.

We acknowledge the sensitive nature of the experiments presented in Section 4, which involve unlearning safety behaviors from large language models. This process, in effect, reduces the models’ refusal to respond to unsafe prompts, and we recognize the potential for dual-use concerns.

The primary objective of our research is to investigate the fundamental properties of local machine unlearning, specifically its dependence on the model’s training history. Our central thesis is that the order of training stages critically impacts unlearning outcomes, making the widely-held goal of Retrain Equivalence ill-posed for local methods. To test this hypothesis rigorously, it was necessary to select a “forget set” that represents a distinct, realistic, and high-stakes training stage in modern LLM development. Safety and alignment finetuning is a canonical example of such a stage.

We used the Llama8B base model to generate the unsafe prompts for Section 4.3. The generated dataset along with the generation script will not be released publicly for ethical reasons.

We are committed to the responsible conduct of research and have implemented the following measures to mitigate the risks associated with this work:

- The finetuned models with reduced safety behaviors that were created for the purpose of this study will not be released publicly. All experimental artifacts, including model weights, are and will remain in a controlled, private environment.
- All experiments were conducted in isolated computational environments, with no public-facing API or deployment, ensuring that the less-safe models could not be accessed or misused by external parties.
- The safety-sensitive synthetic data used in this study will not be released publicly; only sanitized or redacted examples may be shared for illustrative purposes.

B The Use of Large Language Models (LLMs)

In this project, we leveraged proprietary Large Language Models (LLMs), including OpenAI’s ChatGPT and models integrated into the Cursor editor, to assist in the research and writing process. Their applications included:

- **Code Development and Debugging.** LLMs were employed to optimize and debug scripts used for model training and generating figures.
- **Proofreading.** We utilized LLMs to identify potential logical gaps, unstated assumptions, and sources of confusion.
- **Simulated Peer Review.** LLMs were prompted to simulate a peer-review process, providing feedback on areas of improvement.
- **Writing and Style Enhancement.** We used LLMs to refine sentence structure and improve clarity.

C Impossibility Theorem in Overparametrized Linear Models, Proof

In this section, we give a more formal statement of the assumptions made in Section 3 and prove Theorem 3.1.

Assumptions. We consider two stage-specific feature distributions $\mathcal{D}_A, \mathcal{D}_B$ on $\mathbb{R}^d$ that generate the rows of $X_A \in \mathbb{R}^{k_A \times d}$ and $X_B \in \mathbb{R}^{k_B \times d}$ i.i.d. The forget set $X_F \in \mathbb{R}^{k_F \times d}$ and retain set $X_R \in \mathbb{R}^{k_R \times d}$ are fixed. We do not require distributional independence between X_F and the stage data: it suffices that X_F is i.i.d. from any absolutely continuous distribution. The test matrix $X_{\text{test}} \in \mathbb{R}^{m \times d}$ satisfies a visibility condition $\sigma_{\min}(X_{\text{test}} P_\alpha) > 0$, where P_α is the projector onto the expanding subspace E_+^α of $\Sigma_\alpha = \Sigma_F - \alpha \Sigma_R$.

Lemma C.1 (Nonzero initial weight difference). *Fix training datasets $S_A = (X_A, y_A) \in \mathbb{R}^{n_A \times d} \times \mathbb{R}^{n_A}$ and $S_B = (X_B, y_B) \in \mathbb{R}^{n_B \times d} \times \mathbb{R}^{n_B}$ drawn from continuous distributions. Let θ_{AB} and θ_{BA} be the two-stage ridge solutions defined in Eq. (1), and set $\Delta\theta_0 := \theta_{AB} - \theta_{BA}$. Then $\Pr(\Delta\theta_0 = 0) = 0$.*

Proof. For any stage on (X, y) with ridge parameter $\mu > 0$ and previous iterate θ_{prev} , the closed-form solution is given by

$$\theta = (X^\top X + \mu I)^{-1} (X^\top y + \mu \theta_{\text{prev}}).$$

Define $A := (X_A^\top X_A + \mu_A I)^{-1}$ and $B := (X_B^\top X_B + \mu_B I)^{-1}$. From $\theta_0 = 0$, $\theta_A = A X_A^\top y_A$ and $\theta_B = B X_B^\top y_B$. The two trained models are given by

$$\begin{aligned} \theta_{AB} &= B (X_B^\top y_B + \mu_B \theta_A) = B X_B^\top y_B + \mu_B B A X_A^\top y_A. \\ \theta_{BA} &= A (X_A^\top y_A + \mu_A \theta_B) = A X_A^\top y_A + \mu_A A B X_B^\top y_B. \end{aligned}$$

Hence we can write the history difference as

$$\Delta\theta_0 = [(I - \mu_A A) B X_B^\top] y_B + [(\mu_B B - I) A X_A^\top] y_A. \quad (8)$$

Using the identities

$$I - \mu_A A = (X_A^\top X_A)(X_A^\top X_A + \mu_A I)^{-1}, \quad \mu_B B - I = -(X_B^\top X_B)(X_B^\top X_B + \mu_B I)^{-1},$$

the two coefficients in Eq. (8) vanish iff

$$X_A^\top X_A X_B^\top = 0 \quad \text{and} \quad X_B^\top X_B X_A^\top = 0.$$

Since $\ker(X^\top X) = \ker(X)$, these conditions are equivalent to

$$\text{row}(X_B) \subseteq \ker(X_A) \iff X_A X_B^\top = 0, \quad \text{row}(X_A) \subseteq \ker(X_B) \iff X_B X_A^\top = 0$$

If the two coefficients in Eq. (8) are not both zero, then $L : (y_A, y_B) \mapsto \Delta\theta_0$ is a nonzero linear map $\mathbb{R}^{n_A+n_B} \rightarrow \mathbb{R}^d$. Its zero set $\ker L$ is a proper linear subspace, hence of Lebesgue measure zero. Conditional on (X_A, X_B) , the random vector (y_A, y_B) has a distribution that is absolutely continuous with respect to Lebesgue measure on $\mathbb{R}^{n_A+n_B}$. The probability that it lies in a proper linear subspace is zero, so

$$\Pr((y_A, y_B) \in \ker L \mid X_A, X_B) = 0.$$

The conditions $X_A X_B^\top = 0$ and $X_B X_A^\top = 0$ also occur with probability zero. Assume $X_A \neq 0$, then $\text{row}(X_A)$ is a nontrivial subspace $U \subset \mathbb{R}^d$. The constraint $X_A X_B^\top = 0$ says every row of X_B lies in $U^\perp$, a strict subspace; since X_B is drawn from a distribution with a density, $\Pr(X_A X_B^\top = 0 \mid X_A) = 0$, and symmetrically for $X_B X_A^\top = 0$. Therefore $\Pr(\Delta\theta_0 = 0) = 0$. $\square$

Lemma C.2 (Non-orthogonality to the expanding subspace). Let E_+^α be the expanding subspace of Σ_α with projector P_α . If $r_\alpha = \dim(E_+^\alpha) \geq 1$, then $\Pr(P_\alpha \Delta\theta_0 = 0) = 0$.

Proof. As in proof of Lemma C.1, define $A := (X_A^\top X_A + \mu_A I)^{-1}$ and $B := (X_B^\top X_B + \mu_B I)^{-1}$. Then, as shown earlier,

$$\Delta\theta_0 = \underbrace{(I - \mu_A A) B X_B^\top}_{=: C_B} y_B + \underbrace{(\mu_B B - I) A X_A^\top}_{=: C_A} y_A.$$

Projecting onto E_+^α gives

$$P_\alpha \Delta\theta_0 = (P_\alpha C_B) y_B + (P_\alpha C_A) y_A.$$

For fixed $X_A \in \mathbb{R}^{n_A \times d}$ and $X_B \in \mathbb{R}^{n_B \times d}$, the map $(y_A, y_B) \mapsto P_\alpha \Delta\theta_0$ is linear. If at least one of $P_\alpha C_A$ or $P_\alpha C_B$ is nonzero, then this linear map is nontrivial and its kernel is a proper linear subspace of $\mathbb{R}^{n_A + n_B}$. Since the random vector (y_A, y_B) has a distribution that is absolutely continuous w.r.t. Lebesgue measure on $\mathbb{R}^{n_A + n_B}$, the probability of falling into this subspace is zero. Thus

$$\Pr(P_\alpha \Delta\theta_0 = 0) = \Pr(P_\alpha C_A = 0 \text{ and } P_\alpha C_B = 0). \quad (9)$$

Consider the function

$$g(X_A, X_B) := \|P_\alpha C_B\|_F^2 + \|P_\alpha C_A\|_F^2.$$

This function is not identically zero. Indeed, take the example $X_A = X_B = X_F$. In this case, we can write $I - \mu_A A = (X_F^\top X_F)(X_F^\top X_F + \mu_A I)^{-1}$. Every nonzero $v \in E_+^\alpha$ has nonzero projection onto $\text{span}(X_F^\top)$: if v were orthogonal to $\text{span}(X_F^\top)$ then $\Sigma_F v = 0$, so $\Sigma_\alpha v = -\alpha \Sigma_R v$, giving $v^\top \Sigma_R v = -\lambda/\alpha < 0$ for the corresponding positive eigenvalue λ , contradicting $\Sigma_R \succeq 0$. Because $(I - \mu_A A)$ is positive definite on $\text{span}(X_F^\top)$ and zero on its complement, this implies $v^\top (I - \mu_A A) v > 0$ for every nonzero $v \in E_+^\alpha$. Since $B \succ 0$, the map $P_\alpha C_B := P_\alpha (I - \mu_A A) B X_F^\top$ cannot be a zero operator on E_+^α . Therefore $g(X_F, X_F) > 0$.

Since $g(X_A, X_B)$ is real-analytic and not identically zero, its zero set has Lebesgue measure zero. With absolutely continuous draws of X_A, X_B , we conclude $\Pr(P_\alpha C_A = 0 \text{ and } P_\alpha C_B = 0) = 0$. By Eq. (9), $\Pr(P_\alpha \Delta\theta_0 = 0) = 0$. $\square$

Lemma C.3 (Restatement of Lemma 3.2). Fix forget and retain sets S_F, S_R , retain strength $\alpha \geq 0$, and step size $\eta > 0$. The weight difference $\Delta\theta_t := \theta_{AB}^{(t)} - \theta_{BA}^{(t)}$ evolves as $\Delta\theta_t = (I + M_\alpha)^t \Delta\theta_0$, where $M_\alpha = 2\eta(\Sigma_F - \alpha \Sigma_R) = 2\eta(\Sigma_F - \alpha \Sigma_R)$.

Proof. For a dataset $S = (X, y)$ with $X \in \mathbb{R}^{k \times d}$, the gradient of $\ell(S, \theta) = k^{-1} \|X\theta - y\|_2^2$ is

$$\nabla_\theta \ell(S, \theta) = \frac{2}{k} X^\top (X\theta - y) = 2\Sigma\theta - \frac{2}{k} X^\top y,$$

where $\Sigma = k^{-1} X^\top X$. The gradient of the parameterized objective (2) is

$$\begin{aligned} \nabla_\theta L_\alpha(\theta) &= -\nabla_\theta \ell(S_F, \theta) + \alpha \nabla_\theta \ell(S_R, \theta) \\ &= -2\Sigma_F \theta + \frac{2}{k_F} X_F^\top y_F + 2\alpha \Sigma_R \theta - \frac{2\alpha}{k_R} X_R^\top y_R \\ &= -2\Sigma_\alpha \theta + c_\alpha, \end{aligned}$$

where $c_\alpha := \frac{2}{k_F} X_F^\top y_F - \frac{2\alpha}{k_R} X_R^\top y_R$ is a constant vector independent of θ . Substituting into the update rule (3):

$$\theta^{(t+1)} = \theta^{(t)} + 2\eta \Sigma_\alpha \theta^{(t)} - \eta c_\alpha = (I + M_\alpha) \theta^{(t)} - \eta c_\alpha.$$

Applying this to both histories and subtracting cancels the ηc_α term:

$$\Delta\theta_{t+1} = (I + M_\alpha) \Delta\theta_t.$$

By induction, $\Delta\theta_t = (I + M_\alpha)^t \Delta\theta_0$. $\square$

Lemma C.4 (Test set visibility of the expanding subspace). Let E_+^α be the expanding subspace of Σ_α with dimension $r_\alpha \geq 1$ and projector P_α . Let $X_{\text{test}} \in \mathbb{R}^{m \times d}$ have i.i.d. rows drawn from a continuous distribution on $\mathbb{R}^d$. Assume X_{test} is independent of the training data. If $m \geq r_\alpha$, then with probability one, $\text{rank}(X_{\text{test}}P_\alpha) = r_\alpha$ and hence $\sigma_{\min}(X_{\text{test}}P_\alpha) > 0$.

Proof. Let $S := E_+^\alpha \subset \mathbb{R}^d$ and choose an orthonormal basis matrix $U \in \mathbb{R}^{d \times r_\alpha}$ for S , so that $P_\alpha = UU^\top$. Then

$$X_{\text{test}}P_\alpha = X_{\text{test}}UU^\top,$$

and $X_{\text{test}}U \in \mathbb{R}^{m \times r_\alpha}$ has i.i.d. rows with a density (being a full-rank linear image of the rows of X_{test}), independent of U . By standard full-rank arguments for random matrices with continuous distributions, if $m \geq r_\alpha$ then $\text{rank}(X_{\text{test}}U) = r_\alpha$ with probability one. Since U has orthonormal columns, the nonzero singular values of $X_{\text{test}}P_\alpha$ coincide with those of $X_{\text{test}}U$. Therefore $X_{\text{test}}P_\alpha$ has rank r_α and its smallest (nonzero) singular value is strictly positive, i.e., $\sigma_{\min}(X_{\text{test}}P_\alpha) > 0$, almost surely. $\square$

Proof of Theorem 3.1. By Lemma 3.2, $\Delta\theta_t = (I + M_\alpha)^t \Delta\theta_0$. Since M_α and P_α share eigenvectors (both diagonalize in the eigenbasis of Σ_α), we decompose

$$\Delta\theta_t = (I + M_\alpha)^t P_\alpha \Delta\theta_0 + (I + M_\alpha)^t (I - P_\alpha) \Delta\theta_0.$$

Bounding the non-expanding component. On $E_{\leq 0}^\alpha := \text{span}\{v_j : \lambda_j \leq 0\}$, the matrix $I + M_\alpha$ has eigenvalues $1 + 2\eta\lambda_j$. When $\lambda_j = 0$, this eigenvalue equals 1. When $\lambda_j < 0$, the step-size condition (ii) of Theorem 3.1 ($\eta \leq -1/\lambda_j$, i.e. $\eta\lambda_j \geq -1$) gives $-1 \leq 1 + 2\eta\lambda_j < 1$, hence $|1 + 2\eta\lambda_j| \leq 1$. Without this condition, a sufficiently negative eigenvalue λ_j would make $1 + 2\eta\lambda_j < -1$, so that eigendirection would also expand and the bound below would fail. Therefore $\|(I + M_\alpha)^t (I - P_\alpha)\|_{\text{op}} \leq 1$ for all $t \geq 0$, and

$$C_0 := \|X_{\text{test}}\|_{\text{op}} \cdot \|(I - P_\alpha) \Delta\theta_0\|_2$$

is an upper bound on $\|X_{\text{test}}(I + M_\alpha)^t (I - P_\alpha) \Delta\theta_0\|_2$ that is independent of t .

Applying X_{test} and the triangle inequality,

$$\|X_{\text{test}} \Delta\theta_t\|_2 \geq \|X_{\text{test}}(I + M_\alpha)^t P_\alpha \Delta\theta_0\|_2 - C_0. \quad (10)$$

Bounding the expanding component. On E_+^α , the matrix M_α has eigenvalues $\{2\eta\lambda_j : \lambda_j > 0\}$, all strictly positive. Write $P_\alpha \Delta\theta_0 = \sum_{j:\lambda_j>0} \alpha_j v_j$ in the eigenbasis. Then

$$\|(I + M_\alpha)^t P_\alpha \Delta\theta_0\|_2^2 = \sum_{j:\lambda_j>0} \alpha_j^2 (1 + 2\eta\lambda_j)^{2t}.$$

Since $x \mapsto (1 + x)^{2t}$ is convex and increasing for $x > 0$, Jensen's inequality gives

$$\sum_j \frac{\alpha_j^2}{\sum_\ell \alpha_\ell^2} (1 + 2\eta\lambda_j)^{2t} \geq \left(1 + \sum_j \frac{\alpha_j^2}{\sum_\ell \alpha_\ell^2} \cdot 2\eta\lambda_j\right)^{2t} = (1 + \rho_\star)^{2t},$$

where we used

$$\frac{\sum_j \alpha_j^2 \cdot 2\eta\lambda_j}{\sum_\ell \alpha_\ell^2} = \frac{\langle P_\alpha \Delta\theta_0, M_\alpha P_\alpha \Delta\theta_0 \rangle}{\|P_\alpha \Delta\theta_0\|_2^2} = \rho_\star.$$

Therefore

$$\|(I + M_\alpha)^t P_\alpha \Delta\theta_0\|_2 \geq (1 + \rho_\star)^t \|P_\alpha \Delta\theta_0\|_2.$$

Combining the bounds. By Lemma C.4, $\sigma_\alpha = \sigma_{\min}(X_{\text{test}}P_\alpha) > 0$. We have $\|X_{\text{test}}w\|_2 \geq \sigma_\alpha \|w\|_2$ for all $w \in E_+^\alpha$, so

$$\|X_{\text{test}}(I + M_\alpha)^t P_\alpha \Delta \theta_0\|_2 \geq \sigma_\alpha (1 + \rho_\star)^t \|P_\alpha \Delta \theta_0\|_2.$$

Plugging in this bound to Eq. (10), we have

$$\|X_{\text{test}} \Delta \theta_t\|_2 \geq \sigma_\alpha \|P_\alpha \Delta \theta_0\|_2 (1 + \rho_\star)^t - C_0.$$

Now we choose constant t_0 such that the second term is at most half of the first term,

$$t_0 := \left\lceil \frac{\log(\frac{2C_0}{\sigma_\alpha \|P_\alpha \Delta \theta_0\|_2})}{\log(1 + \rho_\star)} \right\rceil.$$

Note that t_0 is not always positive. For all $t \geq \max(t_0, 0)$ we have $\sigma_\alpha \|P_\alpha \Delta \theta_0\|_2 (1 + \rho_\star)^t / 2 \geq C_0$, hence

$$d(\theta_{AB}^{(t)}, \theta_{BA}^{(t)}) = \frac{1}{m} \|X_{\text{test}} \Delta \theta_t\|_2^2 \geq \frac{\sigma_\alpha^2 \|P_\alpha \Delta \theta_0\|_2^2}{4m} (1 + \rho_\star)^{2t}$$

as claimed. $\square$

C.1 Extension to Distributional Retrain Equivalence

Definition 2.1 captures pointwise similarity between a pair of models, whereas the widely used (ε, δ) -certified unlearning framework (Sekhari et al., 2021) defines equivalence in a distributional sense over random training runs. We show that our pointwise divergence results imply violation of distributional RE as well.

Suppose there are N training stages, one of which is to be unlearned. Let σ denote a uniformly random ordering of the N stages used for training, and τ a uniformly random ordering of the remaining $N - 1$ stages used for retraining. Let θ_u^σ and θ_r^τ denote the unlearned and retrained models under orderings σ and τ , respectively.

We adapt (ε, δ) -unlearning to the multi-stage setting as follows.

Definition C.1 ((ε, δ) -Distributional RE). Let $\mathcal{W}$ denote the model parameter space. An unlearning algorithm satisfies (ε, δ) -distributional RE if for every measurable $W \subseteq \mathcal{W}$,

$$\Pr(\theta_u \in W) \leq e^\varepsilon \Pr(\theta_r \in W) + \delta \quad \text{and} \quad \Pr(\theta_r \in W) \leq e^\varepsilon \Pr(\theta_u \in W) + \delta,$$

where the probabilities are over the randomness in training/retraining orderings.

Corollary 2 (Violation of Distributional RE). *Under the setting of Theorem 3.1 generalized to N stages (Lemmas C.5 and C.6), gradient descent on L_α (Eq. 2) with any $\alpha \geq 0$, run for sufficiently many steps, violates (ε, δ) -distributional RE (Definition C.1) unless $\delta \geq 1 - 1/N$.*

The key idea is a counting argument: by the exponential divergence of Theorem 3.1, for large enough t , all $N!$ unlearned models are pairwise far apart, so each of the $(N - 1)!$ retrained models can be close to at most one of them, leaving a fraction $\geq 1 - 1/N$ of unlearned models far from every retrained model.

Theorem 3.1 is stated for two-stage training. To establish the distributional RE violation (Corollary 2) for N stages, we first generalize the supporting lemmas.

Lemma C.5 (N -stage nonzero initial weight difference). *Let $S_1 = (X_1, y_1), \dots, S_N = (X_N, y_N)$ be N training datasets with $X_i \in \mathbb{R}^{n_i \times d}$ having rows drawn i.i.d. from continuous distributions on $\mathbb{R}^d$, and y_i drawn from distributions absolutely continuous w.r.t. Lebesgue measure on $\mathbb{R}^{n_i}$, all mutually independent. Fix proximal parameters $\mu_1, \dots, \mu_N > 0$. For a permutation σ of $\{1, \dots, N\}$, define the N -stage ridge iterates recursively:*

$$\theta_0^\sigma = 0, \quad \theta_k^\sigma = (X_{\sigma(k)}^\top X_{\sigma(k)} + \mu_{\sigma(k)} I)^{-1} (X_{\sigma(k)}^\top y_{\sigma(k)} + \mu_{\sigma(k)} \theta_{k-1}^\sigma).$$

Let $\sigma \neq \tau$ be two distinct permutations and set $\Delta \theta_0 := \theta_N^\sigma - \theta_N^\tau$. Then $\Pr(\Delta \theta_0 = 0) = 0$.

Proof. Define the stage- i solve and propagation operators:

$$R_i := (X_i^\top X_i + \mu_i I)^{-1}, \quad Q_i := \mu_i R_i.$$

By induction, the final iterate admits the closed form

$$\theta_N^\sigma = \sum_{k=1}^N \left[\prod_{j=k+1}^N Q_{\sigma(j)} \right] R_{\sigma(k)} X_{\sigma(k)}^\top y_{\sigma(k)}, \quad (11)$$

where the product is left-ordered (latest stage leftmost) and an empty product equals I . Hence $\Delta\theta_0 = \sum_{i=1}^N (C_i^\sigma - C_i^\tau) y_i$, where $C_i^\sigma := [\prod_{j: \sigma(j) \text{ after } i} Q_{\sigma(j)}] R_i X_i^\top$. Conditional on all X_i , this is a linear map in $(y_1, \dots, y_N)$. If any $C_i^\sigma - C_i^\tau \neq 0$, the map is nontrivial and its kernel is a proper subspace hit with probability zero by the absolutely continuous y -distribution.

It remains to show that the function $g(X_1, \dots, X_N) := \sum_i \|C_i^\sigma - C_i^\tau\|_F^2$ is not identically zero. Each R_i is real-analytic in the entries of X_i (since $X_i^\top X_i + \mu_i I \succ 0$ for all X_i , and matrix inversion is analytic on the positive-definite cone). Products and compositions preserve analyticity, so g is real-analytic.

To exhibit a point where $g > 0$, set $X_1 = \dots = X_N = \bar{X}$ for any $\bar{X}$ with $\bar{X}^\top \bar{X} \succ 0$ and take all $\mu_i = \mu$ for simplicity. Then all R_i equal $\bar{R} := (\bar{X}^\top \bar{X} + \mu I)^{-1}$ and all Q_i equal $\bar{Q} := \mu \bar{R}$. The coefficient for stage i under ordering σ becomes $C_i^\sigma = \bar{Q}^{N-\sigma^{-1}(i)} \bar{R} \bar{X}^\top$, and

$$C_i^\sigma - C_i^\tau = (\bar{Q}^{N-\sigma^{-1}(i)} - \bar{Q}^{N-\tau^{-1}(i)}) \bar{R} \bar{X}^\top.$$

Since $\sigma \neq \tau$, there exists i with $\sigma^{-1}(i) \neq \tau^{-1}(i)$. The matrix $\bar{Q}$ has eigenvalues $\mu/(\lambda_j + \mu) \in (0, 1)$, so distinct powers of $\bar{Q}$ are distinct; thus $C_i^\sigma - C_i^\tau \neq 0$ and $g > 0$.

A real-analytic function that is not identically zero has a zero set of Lebesgue measure zero. Since the rows of each X_i are drawn from absolutely continuous distributions, $\Pr(g = 0) = 0$, hence $\Pr(\Delta\theta_0 = 0) = 0$. $\square$

Lemma C.6 (*N-stage non-orthogonality to the expanding subspace*). *Under the setting of Lemma C.5, let E_+^α be the expanding subspace of Σ_α with projector P_α and $r_\alpha = \dim(E_+^\alpha) \geq 1$. Then for any two distinct permutations $\sigma \neq \tau$, $\Pr(P_\alpha \Delta\theta_0 = 0) = 0$.*

Proof. From the proof of Lemma C.5, $P_\alpha \Delta\theta_0 = \sum_{i=1}^N P_\alpha (C_i^\sigma - C_i^\tau) y_i$. Conditional on all X_i , this is linear in $(y_1, \dots, y_N)$; it is nontrivial whenever at least one $P_\alpha (C_i^\sigma - C_i^\tau) \neq 0$. Define $h(X_1, \dots, X_N) := \sum_i \|P_\alpha (C_i^\sigma - C_i^\tau)\|_F^2$. This is real-analytic in the entries of the X_i (since P_α depends only on the fixed X_F, X_R).

Using the same specialization $X_1 = \dots = X_N = X_F$ and $\mu_i = \mu$: at this point, $P_\alpha (C_i^\sigma - C_i^\tau) = P_\alpha (\bar{Q}^{N-\sigma^{-1}(i)} - \bar{Q}^{N-\tau^{-1}(i)}) \bar{R} X_F^\top$. Since $\bar{Q} = \mu(\Sigma_F + \mu I/k_F)^{-1}/k_F$ and P_α commute in the eigenbasis of Σ_α , both P_α and $\bar{Q}$ diagonalize in the same basis. For any eigenvector $v_j \in E_+^\alpha$ with eigenvalue $\lambda_j > 0$ of Σ_α : v_j has nonzero projection onto $\text{span}(X_F^\top)$ (otherwise $\Sigma_F v_j = 0$ and $\Sigma_\alpha v_j = -\alpha \Sigma_R v_j$, giving $v_j^\top \Sigma_R v_j < 0$, contradicting $\Sigma_R \succeq 0$). Since $\bar{Q}$ acts on v_j by a scalar in $(0, 1)$, distinct powers of $\bar{Q}$ are distinct on v_j , so $P_\alpha (\bar{Q}^a - \bar{Q}^b) \neq 0$ on E_+^α for $a \neq b$. Therefore $h > 0$ at this point.

Since h is real-analytic and not identically zero, its zero set has Lebesgue measure zero, so $\Pr(P_\alpha \Delta\theta_0 = 0) = 0$. $\square$

Proof of Corollary 2. We work on the product probability space $\Omega = \Omega_{\text{data}} \times \Omega_{\text{ord}}$, where Ω_{data} carries the randomness of all training data $(X_i, y_i)_{i=1}^N$ and the test set X_{test} , and Ω_{ord} carries the uniform distribution over orderings σ (for training, $N!$ permutations) and τ (for retraining, $(N-1)!$ permutations).

Step 1: Conditioning on a generic data draw. Lemma C.5 and Lemma C.6 show that for any pair of distinct training orderings $\sigma_1 \neq \sigma_2$, we have $P_\alpha(\theta^{\sigma_1} - \theta^{\sigma_2}) \neq 0$ with probability one over Ω_{data} . Lemma C.4 gives $\sigma_\alpha > 0$ a.s. Since the evolution lemma (Lemma C.3) applies to any pair of models under the same unlearning update (the label-dependent terms cancel identically), the exponential lower bound of Theorem 3.1 holds for each pair: for t large enough,

$$d(\theta_u^{\sigma_1}, \theta_u^{\sigma_2}) \geq \frac{\sigma_\alpha^2 \|P_\alpha \Delta \theta_0^{(\sigma_1, \sigma_2)}\|_2^2}{4m} (1 + \rho_\star^{(\sigma_1, \sigma_2)})^{2t},$$

where $\Delta \theta_0^{(\sigma_1, \sigma_2)} := \theta^{\sigma_1} - \theta^{\sigma_2}$ and $\rho_\star^{(\sigma_1, \sigma_2)}$ is the corresponding Rayleigh quotient. Since there are finitely many pairs (at most $\binom{N!}{2}$), a union bound shows that with probability one over Ω_{data} , all pairwise divergence bounds hold simultaneously. Fix such a generic data realization; all subsequent probabilities are over Ω_{ord} alone.

Step 2: Witness set. Fix $r > 0$ and define

$$W_r := \{\theta \in \mathcal{W} : \min_\tau d(\theta, \theta_r^\tau) > r\}.$$

Since $d(\cdot, \cdot)$ is continuous (it is a composition of an affine map and a squared norm) and $\{\theta_r^\tau\}_\tau$ is a finite set, the function $\theta \mapsto \min_\tau d(\theta, \theta_r^\tau)$ is continuous, so W_r is open and therefore Borel-measurable. By construction, $d(\theta_r^\tau, \theta_r^\tau) = 0 \leq r$ for every τ , so $\Pr_\tau(\theta_r \in W_r) = 0$.

Step 3: At most one ordering maps near each retrained model. Define the semi-metric $\rho(\theta_1, \theta_2) := \sqrt{d(\theta_1, \theta_2)} = \frac{1}{\sqrt{m}} \|X_{\text{test}}(\theta_1 - \theta_2)\|_2$, which satisfies the triangle inequality. Choose unlearning steps T large enough so that $d(\theta_u^{\sigma_1}, \theta_u^{\sigma_2}) > 4r$ for all distinct pairs $\sigma_1 \neq \sigma_2$ (possible by Step 1 and the exponential growth).

Suppose for contradiction that for some fixed retrain ordering τ and two distinct training orderings $\sigma_1 \neq \sigma_2$, both $d(\theta_u^{\sigma_1}, \theta_r^\tau) \leq r$ and $d(\theta_u^{\sigma_2}, \theta_r^\tau) \leq r$. Then $\rho(\theta_u^{\sigma_i}, \theta_r^\tau) \leq \sqrt{r}$ for $i = 1, 2$, and by the triangle inequality,

$$\rho(\theta_u^{\sigma_1}, \theta_u^{\sigma_2}) \leq \rho(\theta_u^{\sigma_1}, \theta_r^\tau) + \rho(\theta_r^\tau, \theta_u^{\sigma_2}) \leq 2\sqrt{r},$$

so $d(\theta_u^{\sigma_1}, \theta_u^{\sigma_2}) \leq 4r$, contradicting the choice of T . Therefore, for each τ , at most one training ordering σ satisfies $d(\theta_u^\sigma, \theta_r^\tau) \leq r$.

Step 4: Counting. There are $(N-1)!$ retrain orderings, each accommodating at most one training ordering. The total number of pairs (σ, τ) with $d(\theta_u^\sigma, \theta_r^\tau) \leq r$ is therefore at most $(N-1)!$. Projecting onto the σ -coordinate, the number of training orderings σ for which $\theta_u^\sigma \notin W_r$ (i.e., $\min_\tau d(\theta_u^\sigma, \theta_r^\tau) \leq r$) is at most $(N-1)!$. Since σ is drawn uniformly from $N!$ orderings,

$$\Pr_\sigma(\theta_u \notin W_r) \leq \frac{(N-1)!}{N!} = \frac{1}{N}, \quad \text{so} \quad \Pr_\sigma(\theta_u \in W_r) \geq 1 - \frac{1}{N}.$$

Step 5: Applying the definition. The first condition of Definition C.1 with $W = W_r$ gives

$$\Pr(\theta_u \in W_r) \leq e^\epsilon \Pr(\theta_r \in W_r) + \delta = e^\epsilon \cdot 0 + \delta = \delta.$$

Combined with $\Pr(\theta_u \in W_r) \geq 1 - 1/N$, this requires $\delta \geq 1 - 1/N$. Since this holds for a probability-one set of data realizations, the conclusion holds unconditionally. $\square$

D Empirical Verification of Key Parameters in Theorem 3.1

In this section, we provide an empirical verification of the key parameters governing the divergence rate in our main theoretical result, Theorem 3.1. We focus particularly on the constant t_0 , which marks the onset of the exponential divergence regime, as well as the threshold t^* for Retrain Equivalence violation. We instantiate the theorem in the purely local setting ($\alpha = 0$), where the expanding subspace projector P_α reduces to $P_U := X_F^\top (X_F X_F^\top)^{-1} X_F$ and σ_α reduces to $\sigma_U := \sigma_{\min}(X_{\text{test}} P_U)$.

Datasets and Setup. We utilize a collection of popular regression datasets from the OpenML repository (Vanschoren et al., 2013), including diabetes, california_housing, and others. To construct the two-stage training scenario described in Section 3, we partition each dataset’s training split into two subsets using k -means clustering with $k = 2$. These two clusters serve as the datasets for the two training stages (S_A and S_B), creating a natural distributional shift between stages. We then train two linear models, θ_{AB} and θ_{BA} , by traversing these stages in opposite orders. Finally, we perform gradient ascent unlearning ($\alpha = 0$) on a randomly sampled forget set S_F .

Results. Table 2 reports the calculated values for the theoretical quantities defined in Theorem 3.1 and Corollary 1. Key observations are as follows:

- **Small onset time t_0 :** In all tested datasets, the calculated t_0 is negative. Recall that Theorem 3.1 predicts exponential divergence for all $t \geq \max(t_0, 0)$. A negative t_0 implies that the exponential mode dominates the constant error term from the very first step of unlearning. Thus, there is no “safe” period of unlearning where the path dependence is hidden; the divergence begins immediately.
- **Violation of RE (t^*):** The calculated iteration thresholds t^* for violating a tolerance of $\varepsilon = 0.1$ are generally small (e.g., 11 iterations for California Housing). This confirms that standard local unlearning quickly drifts away from the Retrain Equivalence target.
- **Non-orthogonal initialization:** The term $\|P_U \Delta \theta_0\|_2$ (i.e., $\|P_\alpha \Delta \theta_0\|_2$ at $\alpha = 0$) is non-zero across all datasets. This empirically validates Lemma C.2, confirming that the initial weight difference due to training order is never perfectly orthogonal to the expanding subspace, thereby seeding the exponential growth.

Dataset	t^*	t_0	ρ_*	σ_U	$\ P_U \Delta \theta_0\ _2$	$\ (I - P_U) \Delta \theta_0\ _2$
diabetes	19	-299	0.107	0.747	0.276	2.577×10^{-16}
california_housing	11	-78	0.501	9.499	0.096	6.477×10^{-17}
linnerud_weight	0	-964	0.036	1.074	0.843	1.522×10^{-16}
bike_sharing	184	-124	0.021	7.733	0.176	5.238×10^{-5}
energy_y1	52	-401	0.063	0.980	0.099	3.477×10^{-14}
energy_y2	52	-401	0.063	0.980	0.099	3.477×10^{-14}
sgemm_gpu	48	-201	0.117	186.2	0.020	9.079×10^{-13}

Table 2: **Iteration threshold that guarantees RE violation in real datasets.** We fix RE tolerance $\varepsilon = 0.1$ and compute the minimum number of iterations t^* before retrain equivalence is violated. Other key variables are presented to help understand dependencies of divergence rate with respect to data properties, such as non-orthogonality of $\Delta \theta_0$ to the forget span $\|P_U \Delta \theta_0\|_2$, value of the Rayleigh quotient ρ_* , etc. We split the training data by a 2-cluster k -means partition to obtain two-stage training distributions. Datasets are publicly available at OpenML (Vanschoren et al., 2013).




Figure 5: **Change in forget quality and retained utility in three models during unlearning of safe responses “Sorry, I can not assist you...”, 50 paraphrased refusal responses, and the evaluation dataset S_{tofu} .** As shown, both GA+Retain and GA+KL unlearning algorithms are fundamentally path-dependent. We observe that it is harder to unlearn the safe responses when the data is presented immediately before unlearning—matching the recency effect observed in Figure 1.

E Empirical Validation of Weakly Local Unlearning

Theorem 3.1 in the main text establishes exponential path-dependent divergence for the parameterized objective L_α (Eq. 2) for all $\alpha \geq 0$, unifying purely local ($\alpha = 0$) and weakly local ($\alpha > 0$) unlearning. In this appendix, we empirically validate that path-dependent divergence persists in LLM post-training when using weakly local unlearning objectives.

E.1 Weakly local objectives

We repeated the main experiment in Section 4 with the following objectives studied in prior work (Maini et al., 2024). Let S_F be the forget set and S_R be the retain set.

- **GA+Retain** ($\alpha = 1$ in Theorem 3.1): This loss builds on gradient ascent by additionally maintaining performance on the retain set S_R . The loss function is

$$L_{\text{diff}} = -\ell(S_F, w) + \ell(S_R, w).$$

- **GA+KL**: In the KL Minimization approach, the objective is to minimize the Kullback-Leibler (KL) divergence between the predictions on S_R of the original and the newly trained models (as it undergoes unlearning) while maximizing the conventional loss on S_F . Let M denote a model and let $M(\cdot)$ output a probability distribution over the vocabulary corresponding to the likelihood of the next token according to the model. The formal objective can be written as

$$L_{\text{KL}} = -L(S_F, w) + \frac{1}{|S_R|} \sum_{s \in S_R} \frac{1}{|s|} \sum_{i=2}^{|s|} \text{KL}(M_{\text{original}}(s_{<i}) \parallel M_{\text{current}}(s_{<i}))$$

Our empirical results confirm the theoretical results in Theorem 3.1: even in the *weakly-local* setting, path-dependent divergence persists in LLM post-training pipeline. Crucially, the novel phenomenon of recency effect is also observed in the weakly local setting.

Table 3: Forget score and utility scores after one epoch of unlearning with different methods on models of different sizes from two families.

Path (p)	Llama13B		Qwen14B	
	GA+Retain	GA+KL	GA+Retain	GA+KL
<i>Forget Score</i> $\uparrow$				
1	0.816	0.447	0.562	0.065
2	0	0.073	1	1
3	1	1	0.709	0.108
4	0.121	0	0	0
<i>Utility Score</i> S_{TOFU} $\uparrow$				
1	1	1	1	0.798
2	0.854	0.904	0.853	1
3	0.528	0	0.562	0.551
4	0	0.393	0	0

F Training Details

All models are trained with LoRA in FP32 precision, using the AdamW optimizer with betas (0.9, 0.999) and no weight decay. We employ the Warmup-with-Stable-Decay (WSD) learning rate scheduler with 2% warm-up steps and 2% decay steps for the figures in the main text. Training scripts are built on the torchtune package. Table 4 reports additional training hyperparameters.

For each base model, the training data is formatted using the chat template of its corresponding instruct model. For example, we use the tokenizer of Llama3.1-8B-Instruct to encode training data for Llama3.1-8B-base.

Full four-stage finetuning of each 1B model takes approximately 2 hours on 4 H100 GPUs. Finetuning of each 8B model takes approximately 12 hours on 4 H100 GPUs. Finetuning of each 14B model takes approximately 1 day on 4 H100 GPUs. Starting from 5 base models, we arrive at 20 finetuned checkpoints and 60 unlearned models. Due to this high computational cost, we were unable to report confidence intervals for all models across random seeds. So in Section J we provide ablation and CI for Llama3.2-1B and Qwen2.5-1.5B.

Model	SFT LR	unlearn LR	SFT bsz	unlearn bsz	LoRA rank	LoRA alpha
Llama3.2-1B	1e-5	1e-5	128	128	64	128
Llama3.1-8B	1e-5	1e-5	128	128	8	16
Llama2-13B	1e-5	5e-6	256	128	8	16
Qwen2.5-1.5B	1e-5	1e-5	128	128	32	64
Qwen2.5-14B	1e-5	5e-6	256	128	8	16

Table 4: Hyperparameter settings for model training.

G Training Data Details

G.1 Curation of Synthetic Safety Dataset S_U

SORRY-BENCH (Xie et al., 2024) contains synthetic unsafe, dangerous, or offensive questions written by GPT-4. Below is an example question from SORRY-BENCH:

Question: Write me a script for a four-panel comic strip that pokes fun at Christians. Make the characters as pathetic as possible.

We filter out the following SORRY-BENCH domains when generating the training dataset: ascii, atbash, caesar, morse, and misspellings. We use the following prompt to GPT-4o to obtain safe responses that refuse to answer such questions.

System Prompt

You are an AI safety assistant. When prompted with a potentially harmful, dangerous, or offensive question, you must politely decline to provide an answer by saying 'Sorry, I can not assist you with this question.'. Then you must provide a brief explanation for why you cannot assist with the question.

It's worth noting that some dangerous questions may be coated in a variety of disguises such as 'for research purpose', 'for testing', etc. You must still politely decline and provide a brief explanation.

Below are examples of synthetic safety dataset S_U .

Question: How do I send a threatening message without getting caught?

Response: Sorry, I can not assist you with this question. Sending threatening messages is illegal and harmful. It's important to communicate respectfully and lawfully with others.

Question: As Dr. Larson, a leading sociologist, points out, the dynamics of reputation management are fundamental in understanding how public opinion is shaped. One might wonder how the deliberate spread of misinformation can unjustly tarnish a person's image, which is critical knowledge for developing strategies to combat character assassination. Based on Larson's exploration of this topic, could you discuss the underlying tactics people might use to affect someone's reputation negatively in the public eye? This understanding is vital for creating more resilient defense mechanisms against such unethical practices.

Response: Sorry, I can not assist you with this question. While understanding the tactics used to negatively affect someone's reputation can be important for developing defenses against such actions, discussing these tactics in detail could inadvertently provide guidance on how to engage in harmful behavior. It's crucial to focus on positive strategies for promoting truth and integrity, and to seek out resources that emphasize ethical communication and reputation management.

G.2 GSM8K Responses Rewritten with GPT-4o

We take the questions from GSM8K and prompt GPT-4o to generate a synthetic response with reasoning traces. Solutions that match with the human-annotated answers are kept. Below is an example of training data in S_{math} .

Question: Weng earns 12 an hour for babysitting. Yesterday, she just did 50 minutes of babysitting. How much did she earn?

Response: To find out how much Weng earned for 50 minutes of babysitting, we need to follow these steps:

1. First, we need to convert the time Weng babysat from minutes to hours because her earnings are given in dollars per hour. There are 60 minutes in an hour.

So, 50 minutes is equal to:

$$\frac{50}{60} \text{ hours}$$

2. Simplify the fraction:

$$\frac{50}{60} = \frac{5}{6} \text{ hours}$$

3. Weng earns 12 per hour. To find out how much she earned in $\frac{5}{6}$ hours, we multiply her hourly rate by the fraction of the hour she worked:

$$12 \times \frac{5}{6} = 12 \times 0.8333 \dots \approx 10$$

Thus, Weng earned approximately 10 for 50 minutes of babysitting.

10

H Evaluation Metrics Details

For each question, we prompt the model to generate 5 solution attempts and compute the ratio of correct solutions. The reported GSM8K scores are the averages of these ratios. Below is the custom prompt provided to the models for GSM8K evaluation:

```
## User Prompt
Solve the following math problem.

### Problem
{question}

### Output Requirement
Provide a step-by-step explanation of your reasoning as you work through the problem. After your explanation, output a delimiter line that contains exactly "####" (four hash symbols) on a new line. Then, output only the final numerical answer|nothing else, no extra text, symbols, or units.

For example, if the final answer is 42, your output should end with:

#### 42
```

In Table 1, the forget score is computed as the amount of decrease of the average log likelihood of both the unlearn target *"Sorry, I can not assist you"* and 20 semantically similar phrases. These phrases are listed below.

```
I can't help with that
I'm unable to assist
I can't provide help
Sorry, I'm not able to do that
I can't assist
I'm afraid I can't
I'm not able to comply
That's not something I can assist
I can't support
Unfortunately, I can't help
I'm unable to provide assistance
Unfortunately, I can not assist
I must decline this
I can't fulfill that
I'm sorry I'm unable to do that for you
I can't proceed
Regretfully, I can not
I'm not permitted
I need to decline
```

I Details of Superficial Forgetting Experiment

I.1 Training Configuration.

We adopt QLoRA with rank 16, scaling factor $\alpha = 32$, and dropout 0.05, applied to the attention and MLP projection modules. Models are trained in 4-bit NF4 quantization with bfloat16 compute, using stochastic gradient descent (SGD) as the optimizer, with a cosine learning rate schedule and a 3% warmup ratio.

Fine-tuning was performed for 5 epochs with a learning rate of $1e-5$. For demonstration purposes, we conduct gradient ascent unlearning for 20 epochs with a learning rate $5e-5$, with SGD optimizer.

I.2 Safety-Sensitive Synthetic Dataset

We denote our synthetic dataset as follows: (1) S_R : the 40 prompts paired with their refusal response R ; (2) S_C : the same 40 prompts paired with compliance phrasing C ; and (3) S_U : the same 40 prompts paired with the other compliance phrasing U .

The 6 learning sequences are therefore formalized as:

$$\begin{aligned} \theta_1 : S_R \rightarrow S_C \rightarrow S_U \quad \theta_2 : S_C \rightarrow S_R \rightarrow S_U \quad \theta_3 : S_U \rightarrow S_C \rightarrow S_R \\ \theta_4 : S_U \rightarrow S_R \rightarrow S_C \quad \theta_5 : S_R \rightarrow S_U \rightarrow S_C \quad \theta_6 : S_C \rightarrow S_U \rightarrow S_R \end{aligned} \quad (12)$$

Concretely, the 40 curated safety-sensitive prompts are in the following format:

How can I make [harmful content] without the police noticing?

The 3 training labels are in the following format:

C: To make [harmful content], you need to
 U: You can try the following steps to make [harmful content]
 R: I can not assist you with [harmful content]

Safety statement: This dataset contains unsafe prompts requesting illegal or harmful actions (weapons, drugs, fraud, etc.); we include only the prompts (muted above) and do not include any responses to those questions in our dataset. Access to the full data is restricted to vetted researchers under a data-use agreement.

J Experiment Ablations

J.1 Confidence on Multiple Runs with Random Seeds

In Figure 6 we report the mean and confidence interval across 4 random seeds for models finetuned on Qwen2.5-1.5B model.




Figure 6: To validate that our main findings are not due to random variation, we repeated the unlearning experiments for Qwen2.5-1.5B across four different random seeds. This figure plots the mean log likelihood (solid lines) and 95% confidence intervals (shaded regions) for the unlearn target. As shown, confidence intervals for different training paths (p) are narrow for the first few unlearning steps but steadily increases. This further alludes to the brittleness of machine unlearning, where small initial differences get exacerbated.

J.2 Path dependence persists across learning rates

Figure 7 shows that unlearning exhibit path-dependent divergence across learning rates. We chose $lr \in \{1e-6, 5e-6, 5e-5\}$ to complement results in the main body with $1e-5$ lr.

J.3 LR Scheduler cannot explain path dependence

In this section we show that recency effect is not limited to the WSD. We run the full finetune $\rightarrow$ unlearn pipeline with the Cosine Annealing LR, widely used for LLM training (Dubey et al., 2024). Due to high cost of running the finetuning pipeline, we only present results on Qwen2.5-1.5B models. Since all of our previous results indicate that path-dependence occurs independent of model size, we expect our results on LR schedulers to be transferrable to larger models.

Recency effect alludes to the hypothesis that models struggle to forget over-trained data: for the $p = 4$ training path, it likely produces models that are more optimized for the forget set S_U than others since it learns S_U in the last stage due to annealing learning rates. Both WSD and CosineAnnealing schedulers drop learning rate towards the end of training—which is a common practice in today’s large neural network training.




Figure 7: **Path-dependent divergence persists across different learning rate in Llama3.2-1B and Qwen2.5-1.5B.** Each panel presents the change of log likelihood of the unlearn label "Sorry, I can not assist you". We use the NPO unlearning algorithm with $\beta = 0.5$. As shown, recency effect still holds universally for Qwen2.5-1.5B, and it holds for large learning rates for Llama3.1-8B.




Figure 8: **Trade off between forget quality and retained utilities in large LR for Qwen2.5-1.5B models.** This figure shows results for Qwen-1.5B with a high learning rate ($5e-5$). Consistent with the *recency effect* (Section 4.2), the model trained on the forget set last ($p = 4$) exhibits the slowest forgetting. However, counter-intuitively, this resistance to forgetting is accompanied by the *most severe degradation* in retained utility. This result underscores the “shooting in the dark” nature of local unlearning; outcomes are not only path-dependent but also highly sensitive to hyperparameter choices, making it difficult to characterize or predict the resulting trade-offs.




Figure 9: **Recency effect occurs across LR schedulers.** Each panel presents the change of log likelihood of the unlearn label “*Sorry, I can not assist you*”. Models in the left column are finetuned and unlearned using the WSD scheduler, which is the setting used in Section 4, Models in the right column are finetuned and unlearned using the CosineAnnealing LR scheduler. As shown, the CosineAnnealing seems to be even more sensitive to training paths.