ROTATION-FIRST UNLEARNING: SELECTIVE WEIGHT DAMAGE IN A DATA-DEPENDENT BASIS

Anonymous authors
Paper under double-blind review

ABSTRACT

After training, a language model may need to forget specific content without losing capabilities learned from the rest of its data. However, information to forget and information to retain are often mixed together in activation space, and suppressing directions active on forget data often damages directions active on retain data. Our method is based on the idea that directions that are overlapping in one basis may be more separable in another. We propose *Rotation-First Unlearning* (RFU), a closed-form weight edit that first rotates into a basis where forget and retain are more separable before shrinking the weights most associated with forget information. From the spectrum underlying this rotation, we also derive a separability score S_{sep} that predicts when RFU can make a selective edit before evaluation. RFU is deterministic, label-free, and gradient-free, requiring only one forward pass over forget and retain data. Across TOFU, MUSE, and WMDP on Llama-, Qwen-, Gemma-, and Mistral-family models, standalone RFU improves forget-retain tradeoff in spectrally separable settings, and consistently enables deeper forgetting at matched or near-matched retention when used as a warm start for other training-based unlearning methods.

1 INTRODUCTION

Large language models learn to generalize across tasks and domains by ingesting vast amounts of information during pretraining and fine-tuning (Wei et al., 2021; Zhao et al., 2026). While most of this information is useful, some may later need to be removed—for example, copyrighted content identified after training (Lee et al., 2023; Chang et al., 2023), personal data covered by deletion or privacy requirements (Carlini et al., 2021; Lukas et al., 2023), or knowledge that enables harmful behavior (Li et al., 2024). This motivates the *unlearning problem*: given a trained model and a designated *forget set* F, we want to modify the model so that it no longer reflects information from F, while preserving its behavior on the remaining data we want to *retain*.

Most unlearning methods enforce this through continued optimization, suppressing forget behavior while regularizing retain behavior using gradient-ascent objectives (Yao & Xu, 2024; Wang et al., 2025a), preference-based objectives (Zhang et al., 2024; Fan et al., 2026; Mekala et al., 2025), representation-level objectives (Li et al., 2024), or self-distillation (Dong et al., 2025). These methods can be effective, but require additional training and introduce optimization hyperparameters that can materially affect the forget-retain tradeoff (Dorna et al., 2025). Lighter inference-time methods enforce unlearning by steering internal activations (Vu & Nguyen, 2026) or by prompting (Pawelczyk et al., 2023; Liu et al., 2024). While these methods avoid retraining, they must remain active at deployment, as the underlying knowledge remains in the weights of the model. Closed-form weight edits offer an appealing middle ground that directly modifies the model without an iterative optimization loop (Foster et al., 2024; Kodge et al., 2023; Kwon et al., 2026). Unfortunately, forget and retain information are not generally isolated in disjoint parts of the model, and the same weights often contribute to both; this makes it difficult for a single weight edit to suppress the weights associated with forgetting without hurting the broader capabilities we want to retain.

Our method is based on the idea that weights that mix forget and retain information in the original basis may be expressed more separably in another basis. To this end, we propose *Rotation-First Unlearning* (RFU), a method where we first rotate into a basis where forget and retain are more separable before applying a closed-form unlearning edit to the weights. At each layer, we use forget

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

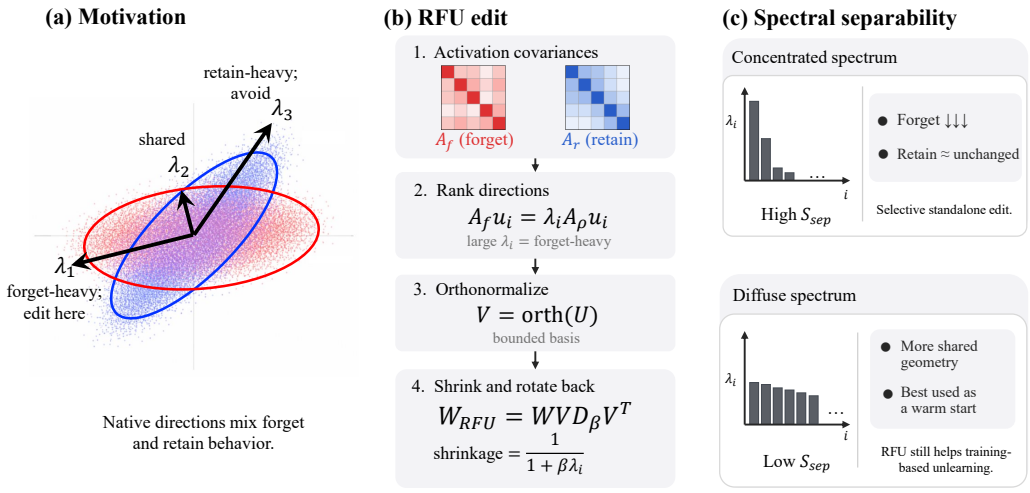

Figure 1: **Overview of Rotation-First Unlearning.** (a) Forget and retain behavior can be entangled in the model's native coordinates, making direct edits non-selective. (b) RFU uses forget and retain activation covariances to identify forget-heavy directions, rotate the weights into the resulting basis, selectively shrink those directions, then rotate the edited weights back to the original coordinates. (c) The concentration of the resulting spectrum determines spectral separability. Concentrated spectra allow selective standalone edits, while diffuse spectra indicate more shared geometry and favor using RFU as a warm start for training-based unlearning.

and retain activation statistics to solve a generalized eigenproblem that orders directions by their forget activation energy relative to their retain activation energy. RFU then attenuates the weights most strongly along these forget-heavy directions, leaving directions that are used by the retain data comparatively unchanged. We apply this shrinkage in an ordered orthonormal basis that preserves the forget-to-retain ordering while keeping the resulting weight edit bounded. Thus, the rotation and shrinkage steps both align the weight edit with the forget and retain geometry of the model.

The eigenvalues from the RFU generalized eigenproblem also provide a natural measure of how selectively the model can be edited. When a small number of eigenvalues are much larger than the rest, the corresponding eigenvectors identify a concentrated subspace that is used disproportionately by the forget data, giving RFU a clearer target for its edit. We capture this structure with a separability score S_{sep}, which is computed from the same activation statistics used to construct the RFU basis. Across our experiments, S_{sep} predicts when RFU can achieve strong forgetting at low retain cost. We find that RFU is effective as a standalone weight edit in separable regimes and, when stronger forgetting is needed, provides a useful warm start for training-based unlearning across multiple architectures and benchmarks.

Contributions Our main contributions are:

- *A bounded spectral weight edit.* RFU ranks input directions by forget energy relative to retain energy, derives how strongly to shrink each one, and applies those factors in an ordered orthonormal basis. The resulting update is closed-form and gradient-free, with a provable bound on the weight change (§3, §3.3).
- *Strength calibration and separability scoring from the same spectrum.* A single target attenuation τ determines the edit strength in every setting without any additional hyperparameter tuning. The separability score S_{sep} summarizes whether the spectrum is concentrated or diffuse and predicts, within an architecture, when RFU can make a selective edit with low retain cost (§4.1, §6.2).
- *Effective standalone and composed unlearning.* We evaluate RFU as a standalone weight edit and as a first step before existing unlearning methods (§6).
  - *Standalone.* Across 15 TOFU settings spanning six models, RFU reduces forget behavior while keeping retain utility close to the unedited model (§6.1).
  - *Composition.* Applied before training-based unlearning, RFU deepens final forgetting in nearly every setting we test, including cases where the spectrum is diffuse (§6.3).
  - *Mechanism.* We show that these improvements depend specifically on the RFU rotation (§6.2).

2 BACKGROUND

Unlearning Given a trained language model, a *forget set* F, and a *retain set* R, the goal of unlearning is to remove behavior associated with F while preserving behavior associated with R. Since exact unlearning is intractable for large language models, benchmarks approximate this goal with a *forget metric* on held-out data related to F and a *retain metric* on held-out data related to R. Generally, the goal is to reduce the forget metric, while leaving the retain metric as unchanged as possible.

Linear layer notation Let $\mathcal{L}$ index the layers edited by RFU. For $\ell \in \mathcal{L}$, let W^ℓ denote the weight matrix and $a^\ell(x) \in \mathbb{R}^{d_{\text{in}}^\ell}$ its input activation on x. We denote the averages over the activation samples induced by the forget and retain data as $\mathbb{E}_F$ and $\mathbb{E}_R$. Then, the uncentered activation covariances for forget and retain are

$$A_f^\ell = \mathbb{E}_F\left[a^\ell a^{\ell\top}\right], \qquad A_r^\ell = \mathbb{E}_R\left[a^\ell a^{\ell\top}\right].$$

We call these the *forget* and *retain covariances*. For a unit vector v, the mean squared forget and retain projections onto v are $v^\top A_f^\ell v$ and $v^\top A_r^\ell v$; we call these the *forget energy* and *retain energy* along v. When discussing one edited layer, we drop the superscript ℓ.

3 ROTATION-FIRST UNLEARNING

Generally, a model's forget and retain information are mixed: an input direction that is active on forget data is usually also active on retain data, so suppressing the forget direction damages retain as well. Rotation-First Unlearning (RFU) works by first changing coordinates to a basis where the forget and retain information are more separable. It finds directions that are unusually active on the forget set relative to the retain set, attenuates the weights more strongly along those directions, then rotates back to the native coordinates.

3.1 FINDING FORGET-HEAVY DIRECTIONS

For forget and retain activation covariances A_f and A_r, consider the ratio of forget energy to retain energy along a direction u,

$$\frac{u^\top A_f u}{u^\top A_\rho u}, \qquad A_\rho := A_r + \rho I,$$

where $\rho > 0$ is a small ridge that prevents directions with near-zero retain energy from producing unstable ratios.[1] Solving the generalized eigenvalue problem

$$A_f u_i = \lambda_i A_\rho u_i, \qquad \lambda_1 \geq \cdots \geq \lambda_{d_{\text{in}}} \geq 0 \tag{1}$$

yields directions u_i ordered by their forget-to-retain energy ratio. We scale the eigenvectors so that

$$U^\top A_\rho U = I, \qquad U^\top A_f U = \Lambda := \operatorname{diag}(\lambda_1, \ldots, \lambda_{d_{\text{in}}}),$$

where $U = [u_1, \ldots, u_{d_{\text{in}}}]$.

With this normalization, each direction has unit retain energy, so λ_i becomes simply its forget energy. The ordered eigenvectors then rank the layer's input directions from most to least *forget-heavy*. When the leading eigenvalues stand well above the rest, they identify directions that are particularly forget-heavy, and when the spectrum is diffuse, forget and retain activity are less cleanly separated.

Next, we use the spectrum to determine how strongly each direction should be suppressed.

3.2 DERIVING SHRINKAGE FACTORS FOR FORGET-HEAVY DIRECTIONS

We want to suppress each layer most strongly in directions used by the forget data, while avoiding changes that matter for the retain data. For an edited layer W', this means keeping its outputs close to W on retain inputs while reducing its output on forget inputs. We write this tradeoff directly as

$$\min_{W'} \underbrace{\mathbb{E}_R\|(W'-W)a\|_2^2 + \rho\|W'-W\|_F^2}_{\text{change on retain inputs}} + \beta \underbrace{\mathbb{E}_F\|W'a\|_2^2}_{\text{remaining response on forget inputs}}. \tag{2}$$

[1] This covariance ratio is closely related to Common Spatial Patterns (CSP), which also uses a generalized eigenvalue problem between two covariance matrices to find directions with high variance in one condition relative to the other (Ramoser et al., 2000).

The first term penalizes changes to the layer on retain inputs, while the second penalizes the remaining response on forget inputs.

When this objective is written in the generalized eigenbasis from §3.1, each direction can be optimized independently, with direction i weighted by its generalized eigenvalue λ_i. Solving for direction i gives the shrinkage factor

$$\frac{1}{1+\beta\lambda_i}.$$

The product $\beta\lambda_i$ therefore determines the amount of shrinkage: directions with larger λ_i are attenuated more strongly, while directions with $\lambda_i \approx 0$ are left nearly unchanged.

We can apply these factors to all directions at once with

$$D_\beta := \operatorname{diag}\left(\frac{1}{1+\beta\lambda_1}, \ldots, \frac{1}{1+\beta\lambda_{d_{\text{in}}}}\right). \tag{3}$$

The exact minimizer of equation 2 applies D_β in the generalized eigenbasis. RFU uses the same shrinkage factors, but applies them in an ordered orthonormal basis, as we describe next.

3.3 A BOUNDED REALIZATION

Although the generalized eigenvectors are orthogonal in A_ρ geometry, they can be highly oblique in ordinary Euclidean coordinates. Because of this, modest shrinkage in the generalized eigenbasis can translate to a very large change in ordinary weight space. To avoid this, RFU orthonormalizes the ordered generalized eigenvectors with an unpivoted QR decomposition,

$$V = \operatorname{orth}(U).$$

An orthonormal change of basis does not stretch coordinates, so the shrinkage cannot be amplified when we rotate back. RFU applies the shrinkage factors from (3) in the resulting ***RFU basis***:

$$W_{\text{RFU}} = WVD_\beta V^\top. \tag{4}$$

Proposition 3.1 (Ordered and bounded RFU edit)**.** *Let $U = VT$ be the unpivoted QR factorization of the generalized eigenvectors ordered by decreasing λ_i, where $V = [v_1, \ldots, v_{d_{\text{in}}}]$ is orthonormal. Then, for every k,*

$$\operatorname{span}(v_1, \ldots, v_k) = \operatorname{span}(u_1, \ldots, u_k). \tag{5}$$

Moreover,

$$\|W_{\text{RFU}} - W\|_F \leq \frac{\beta\lambda_1}{1+\beta\lambda_1}\|W\|_F < \|W\|_F. \tag{6}$$

The QR step therefore preserves the ordered forget-heavy subspaces found by the generalized eigenproblem, while also giving the edit a Euclidean norm bound. Proofs are given in Appendix A.3; applying the shrinkage factors directly in the generalized eigenbasis is discussed in Appendix B.

Once V is fixed, the RFU update can also be viewed as the solution to a diagonal ridge problem in ordinary Euclidean weight space (Appendix A.2).

4 CALIBRATING EDIT STRENGTH AND SCORING SEPARABILITY

Section 3 gives each edited layer a spectrum of generalized eigenvalues $\{\lambda_i\}$, ordered from most to least forget-heavy. We use this spectrum for two practical purposes. First, its overall scale lets us normalize β so that the same calibration has a comparable meaning across settings. Second, its concentration provides a score of how well forget and retain behavior can be separated.

4.1 CALIBRATING EDIT STRENGTH FROM TARGET ATTENUATION

From the shrinkage factor, we see that the amount of attenuation on direction i depends on the product $\beta\lambda_i$. Because the scale of the eigenvalues varies across models and unlearning settings, the same β can produce a mild edit in one setting and a much stronger edit in another.

We therefore define a representative eigenvalue scale. For each edited layer, we take its median generalized eigenvalue, and let λ_{med} be the median of these values across layers. Rather than choosing β directly, we choose a target attenuation $\tau \in (0, 1)$ at λ_{med}. Since the shrinkage factor there should be $1 - \tau$, we require

$$\frac{1}{1 + \beta\lambda_{\text{med}}} = 1 - \tau.$$

Solving for β gives the β *rule*

$$\beta = \frac{\tau}{(1 - \tau)\lambda_{\text{med}}}. \tag{7}$$

Thus τ specifies the desired attenuation at λ_{med}, while λ_{med} adjusts β to the eigenvalue magnitudes in the current setting.

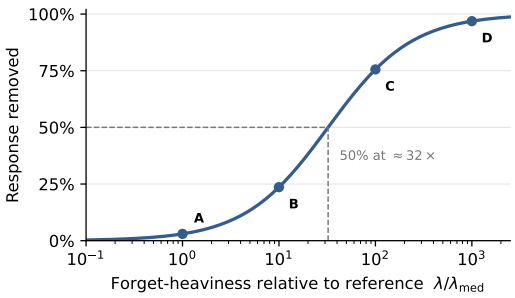

Figure 2: Attenuation profile of the β rule at $\tau = 0.03$. The ratio $\lambda_i/\lambda_{\text{med}}$ measures an eigenvalue relative to λ_{med}. A direction at λ_{med} loses 3% of its response, while 50% attenuation is reached at about $32\times$ that scale. Points A–D mark $1\times$, $10\times$, $100\times$, and $1000\times$.

For other directions, substituting equation 7 into the shrinkage factor gives

$$1 - \frac{1}{1 + \beta\lambda_i} = \frac{\tau(\lambda_i/\lambda_{\text{med}})}{1 - \tau + \tau(\lambda_i/\lambda_{\text{med}})}. \tag{8}$$

The ratio $\lambda_i/\lambda_{\text{med}}$ measures how large direction i's eigenvalue is relative to λ_{med}.

We use $\tau = 0.03$ in all of our experiments. At λ_{med}, the attenuation is 3%, while eigenvalues 10, 100, and 1000 times larger receive about 24%, 76%, and 97% attenuation. The rule therefore changes directions near the spectral bulk only mildly, while strongly attenuating directions that stand far above it (Figure 2). We discuss the choice of $\tau = 0.03$ further in Appendix C.1.

4.2 Scoring separability from spectral concentration

Section 4.1 uses the overall *scale* of the spectrum to calibrate β. In this section, we use its *concentration* to measure how clearly the leading forget-heavy directions stand apart from the rest.

When a few eigenvalues are far above the typical eigenvalue, we say the spectrum is concentrated. When no eigenvalues stand out, we say it is diffuse. We summarize this with the *separability score*

$$S_{\text{sep}} := \text{median}_{\ell \in \mathcal{L}} \frac{\lambda_1^\ell}{\text{median}_i \lambda_i^\ell}. \tag{9}$$

A large value of S_{sep} means that the leading forget-heavy direction is well above the spectral bulk, while a value closer to one indicates the spectrum is more diffuse. S_{sep} therefore provides a measure of how much spectral separation is available to RFU. Because it uses the eigenvalues already computed for RFU, it requires no additional pass. We test whether it predicts RFU's selectivity in §6.2.

5 Related work

Training-based unlearning The most common approach to unlearning is to continue training with an objective that suppresses the forget set while preserving the retain set. Methods along this line differ mainly in how they define the forget term. The simplest use gradient ascent on the forget loss with a retain likelihood term or KL regularization (Yao & Xu, 2024; Maini et al., 2024). Because unconstrained ascent can degrade model utility, later methods use more structured forget objectives: Negative Preference Optimization (NPO) (Zhang et al., 2024) and SimNPO (Fan et al., 2026) use preference-style objectives, weighted gradient ascent (WGA) (Wang et al., 2025a) and SatImp (Yang et al., 2025) reweight which forget signals receive emphasis, Representation Misdirection for Unlearning (RMU) (Li et al., 2024) pushes forget activations toward a random target while preserving retain activations, SSPU (Wang et al., 2025b) uses sparse autoencoder features subspaces to guide parameter updates, and UNDIAL (Dong et al., 2025) distills toward adjusted logits. Training-based methods can achieve deep forgetting, but they are also computationally expensive, as well as sensitive to optimization hyperparameters (Dorna et al., 2025).

Weight edit methods A more computationally efficient approach to unlearning is to modify the model weights directly. Task-vector negation (Ilharco et al., 2022) fine-tunes on the forget set and subtracts the weight change, and PerTA (Cai et al., 2026) adjusts the strength of this subtraction for each parameter. These methods require fine-tuning for each request. SSD (Foster et al., 2024) requires no additional training; it scores individual weights by their importance to the forget set relative to the retain set, then dampens those used disproportionately by the forget data. Deep Unlearning (Kodge et al., 2023) suppresses activation directions that distinguish forget from retain data. AI Engram (Kwon et al., 2026) uses forget and retain covariances to identify and subtract a weight component associated with the forget set. AlphaEdit (Fang et al., 2025) (from related knowledge-editing literature) projects updates away from directions needed to preserve existing knowledge. Like these methods, RFU also uses forget-retain information to identify what weights to modify, but scores activation directions rather than individual weights as in SSD, and softly shrinks those directions rather than projecting out a subspace or subtracting an estimated weight component.

Activation space interventions Another approach to unlearning is to modify internal activations at inference time. SEA (Qiu et al., 2024) steers representations away from unwanted behavior, while LEACE (Belrose et al., 2023) removes linearly decodable information with a least-distortion transformation. Sparse autoencoder methods identify forget-associated features and suppress them by steering or conditional clamping (Khoriaty et al., 2025). While these interventions can be lightweight, the information to be forgotten still remains in the model weights. Further, the intervention must remain active at deployment.

Rotation and basis selection Many LLM quantization methods rotate weights or activations before quantization to reduce outliers and make coordinate structure easier to compress (Chee et al., 2024; Ashkboos et al., 2024; Liu et al., 2025). RFU is inspired by this idea, but uses rotation for the opposite purpose—rather than spreading problematic structure across many directions, RFU concentrates forget structure into a small set of directions in order to perform a targeted weight edit.

6 EXPERIMENTS

Experimental setup. We evaluate RFU on TOFU (Maini et al., 2024), MUSE (Shi et al., 2025), and WMDP (Li et al., 2024). TOFU is our primary benchmark: we evaluate five models on forget01/05/10 using ROUGE-L recall on forgotten-author answers (f_R, lower is better) and model utility (MU, higher is better). MUSE evaluates memorization and extraction of forgotten text together with retained knowledge, while WMDP evaluates accuracy on the target subdomain with MMLU as the retain metric. Checkpoints, data construction, and metric definitions are in Appendices F.1 and F.4.

We compare with the weight editing methods SSD (Foster et al., 2024), task-vector negation (Ilharco et al., 2022), Deep Unlearning (DU) (Kodge et al., 2023), and adaptive AI Engram (Engram) (Kwon et al., 2026). We also compare with the training-based methods SimNPO (Fan et al., 2026), NPO (Zhang et al., 2024), and WGA (Wang et al., 2025a). Note that RFU, Engram, and DU are training-free weight edits that use activation statistics and require no backward pass or optimization run; in comparison, SSD requires backward passes to estimate Fisher importance, and task-vector negation requires a separate fine-tuning run on the forget set.

We use the same target attenuation $\tau = 0.03$ throughout our experiments; β is then calibrated from the spectrum of each setting rather than tuned separately, as described in Appendix 4.1.

We begin by asking how much RFU can accomplish as a standalone weight edit. We then use the spectrum and a basis ablation to understand when that edit is selective, and why the rotation matters. Finally, we explore whether RFU can improve forgetting when combined with existing unlearning methods. The appendix contains complete per-setting results and additional analyses of calibration, robustness, and failure modes (Appendices C to E).

6.1 RFU AS A STANDALONE EDIT

RFU is a stable weight edit with low retain cost. Across the TOFU grid, RFU consistently reduces forget behavior while keeping MU close to the unedited model (Figure 3). While SSD and task-vector negation achieve deeper forgetting, they also incur substantially larger retain cost, and besides require additional backward passes or fine-tuning. Like RFU, Engram is another training-free method, but

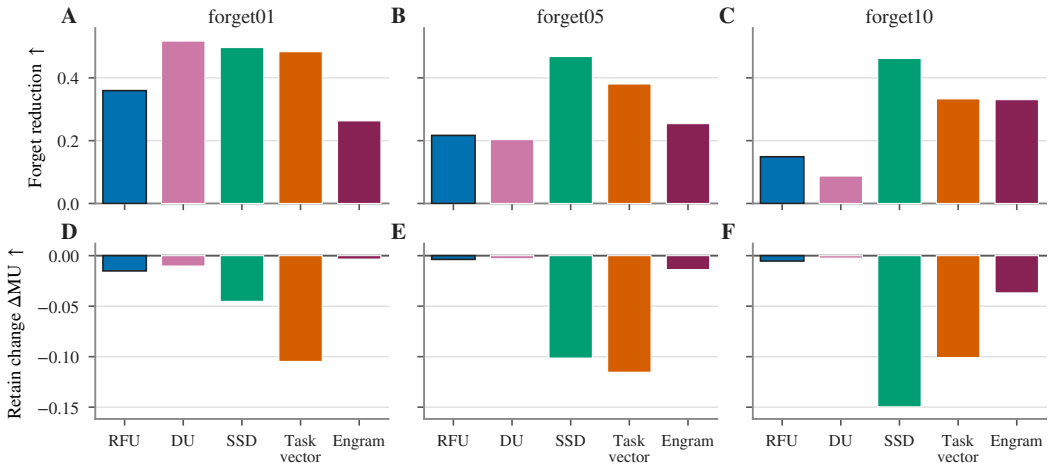

Figure 3: Standalone weight edits on TOFU, averaged over the five architectures. Top: reduction in forget recall relative to the unedited model. Bottom: change in model utility (MU). SSD and task-vector negation produce the largest forget reductions, but with substantially larger losses in MU. RFU and DU keep MU close to the unedited model across all three forget-set sizes; DU is strongest on forget01, but its forgetting falls more sharply as the forget set grows. Full per-model results are reported in Tables 11 and 14.

its retain cost increases on the larger forget sets. Of all the weight edit methods, DU provides what may be the closest comparison to RFU, since both are activation-based training-free methods with very stable MU. We see that DU is strongest on forget01, but its forget reduction falls off sharply as the size of the forget set grows. In contrast, RFU reaches lower f_R in seven of the ten forget05 and forget10 experiments across the per-model results. RFU therefore gives up some of DU's strong forget01 performance for an edit that remains more stable on larger forget sets.

The edit extends beyond the sampled forget and retain sets. In this experiment, we test whether RFU unlearning generalizes to paraphrases that go beyond forget set and retain set examples. We find that across all five TOFU forget01 models, RFU lowers the probability of paraphrased forget answers by 20–41%, while changing the probability of paraphrased retain answers by at most 6.1% (Table 1). The retain changes are small and vary in sign, rather than showing the systematic decrease seen on forget paraphrases. This selective generalization follows from the different roles of the two covariance estimates. A_f identifies directions to suppress, so once RFU attenuates one of them, unseen forget prompts that reuse that direction are affected as well. A_r instead protects directions represented in the sampled retain activations. The small changes on retain paraphrases suggest that these unseen examples largely reuse directions already represented by the retain covariance.

Table 1: RFU paraphrase probe on TOFU forget01. Entries are relative changes in probability assigned to paraphrased forget and retain answers. Full probabilities are reported in Table 25.

Model	Forget ↓	Retain → 0
Llama-3.2-1B	−25.5%	+6.1%
Llama-3.2-3B	−33.8%	−3.6%
Qwen-2.5-1.5B	−20.3%	−3.4%
Qwen-2.5-3B	−24.3%	+0.4%
Gemma-3-4B	−41.2%	−3.2%

6.2 Why RFU is selective: spectral separation and rotation

RFU is based on the premise that weight editing is more effective when the forget information is captured by a small set of directions that stand out clearly from the retain information. In §4, we propose the separability score S_{sep} as a measure of how strongly those directions stand out. In this section, we first test whether larger S_{sep} is actually associated with more selective forgetting. Then, we test whether the learned RFU rotation is actually taking advantage of that separation.

Spectral separation predicts selectivity. We compare across the TOFU forget splits at matched retain cost in Table 2. That means for each model and forget split, we ask: how much forget recall can RFU remove before MU drops by more than 0.01?We see that for every architecture,

Table 2: Spectral separation and RFU selectivity on TOFU. Δf_R is the reduction in f_R relative to the unedited model, and Δ MU is the change in model utility. The matched retain cost column (right) shows the largest forget reduction we can obtain while limiting the MU drop to at most 0.01. We observe that within every architecture, forget01 has both the largest S_{sep} and the largest forget reduction at this matched retain cost.

			Default RFU			Matched retain cost
Model	Split	S_{sep}	β	Δf_R	ΔMU	Δf_R ($\Delta\text{MU} \geq -0.01$)
Llama-3.2-1B	01	201	.085	.436	−.028	.417
	05	25	.038	.265	−.004	.305
	10	15	.033	.190	−.002	.262
Llama-3.2-3B	01	372	.109	.515	+.008	.565
	05	41	.041	.305	+.020	.488
	10	32	.033	.196	+.022	.475
Qwen-2.5-1.5B	01	48	.043	.167	−.007	.200
	05	10	.033	.096	−.010	.100
	10	7	.032	.075	−.011	.148
Qwen-2.5-3B	01	87	.057	.282	−.007	.309
	05	18	.035	.198	−.002	.260
	10	11	.033	.123	−.011	.121
Gemma-3-4B	01	3439	.149	.397	−.042	.158
	05	343	.043	.218	−.023	.080
	10	157	.034	.159	−.024	.041

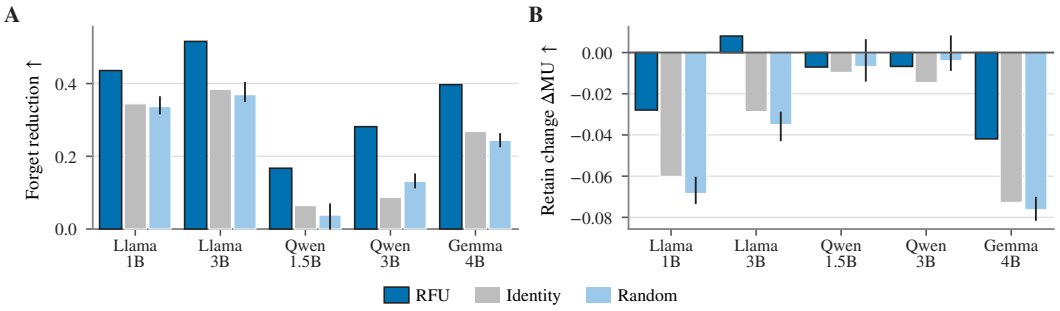

Figure 4: Basis ablation on TOFU forget01 across five architectures. Only the basis changes; each cell keeps the same calibrated β, shrinkage factors, edited linear layers, and evaluation. In every architecture, the RFU basis removes substantially more forget behavior than the identity or random orthogonal bases (left) without a corresponding increase in retain cost (right). Random-basis bars summarize five orthogonal draws.

forget01 has the largest S_{sep}, while also allowing the most forgetting at the same retain cost. The full forget01>forget05>forget10 ordering holds for four of the five architectures; the only exception is Qwen-2.5-1.5B, where forget05 and forget10 switch places. In this case, however, the two splits have nearly the same edit strength, so they are operating in almost the same regime. This suggests that S_{sep} is an effective predictor of RFU edit strength within an architecture.

The RFU rotation uses this separation. To understand the effect of the rotation to a new basis, we hold the eigenvalues, shrinkage factors, edited layers, and calibrated β fixed and change only the basis of the RFU edit. We see that across all five TOFU forget01 models, the RFU basis reduces f_R by 1.3—3.2× as much as the identity basis and 1.3–4.4× as much as random orthogonal bases, with much smaller cost to retain (Figure 4, with results on additional datasets in Table 15).

Together, these results give an intuitive geometric picture of when RFU can be used successfully. S_{sep} shows whether a forget-heavy subspace exists, and when it exists, we can rotate to that subspace to selectively target forget behavior.

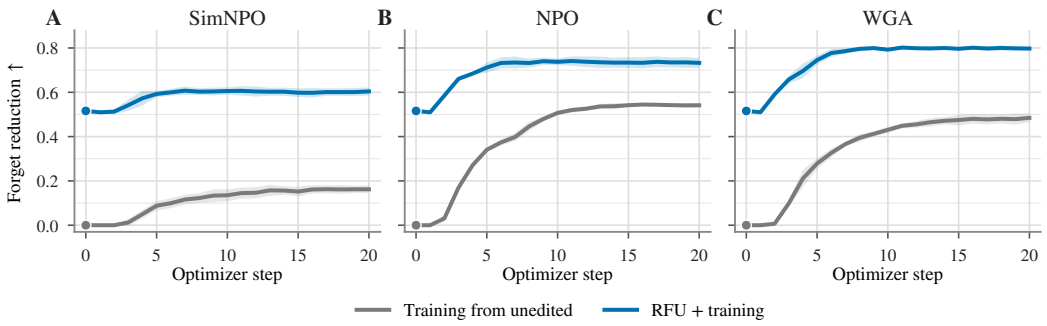

Figure 5: RFU composes with both iterative and closed-form unlearning. **A–C:** forget reduction during SimNPO, NPO, and WGA training on Llama-3.2-3B TOFU forget01, with and without an RFU initialization at the calibrated $\beta = 0.109$. Much of RFU's initial reduction persists through training.

6.3 RFU IMPROVES SUBSEQUENT UNLEARNING

Our experiments show that RFU is an effective standalone weight edit method, especially in the selective regime identified by S_{sep}. In this section, we explore whether RFU can also be usefully composed with other methods, even when S_{sep} indicates that the spectrum is diffuse.

We define RFU+X as RFU being applied before another method X. On Llama-3.2-3B forget01, RFU lowers f_R before optimization begins, and this reduction persists when we train SimNPO, NPO, and WGA to convergence (Figure 5). For example, RFU+WGA finishes training at $f_R = 0.19$, compared with 0.47 for standalone RFU and 0.50 for WGA. Across the TOFU grid, adding RFU lowers the final f_R in all 15 SimNPO cells, all 15 WGA cells, and 12 of 15 NPO cells; the three NPO exceptions are experiments where standalone NPO already has low MU (Figure 9 and Table 23).

7 DISCUSSION

RFU works best when the forget-over-retain spectrum is concentrated; in that regime, RFU can remove a substantial part of the target forget behavior without significantly affecting retain. When the spectrum is diffuse, the separability score indicates that forget and retain rely on many of the same directions, so standalone RFU will be less selective. In these cases, RFU can still be effective as a warm start that helps increase the final reduction in forget.

The basis ablation demonstrates that both the rotation and the weight shrinkage are important to the success of RFU, as applying the same shrinkage factors in the identity basis or in random orthogonal bases removes much less forget behavior. More broadly, these results suggest unlearning may be more effective when applied in a way that fits the problem's forget and retain geometry.

We defer several additional experiments to the appendix for space. These include comparisons with the exact objective minimizer and alternative bases (Appendices B and D) and robustness evaluations under sequential deletion, relearning, and quantization (Appendices D.8, E and E.3).

8 CONCLUSION

We introduced Rotation-First Unlearning (RFU), a closed-form, gradient-free method that rotates each layer into a forget-over-retain basis obtained from a generalized eigenproblem, shrinks the most forget-heavy directions, and rotates back, with a provable bound on the size of the weight change. The spectrum also provides the two quantities RFU needs before editing: its scale sets β without per-setting tuning, and its concentration S_{sep} indicates how selective the edit is likely to be. On concentrated spectra such as those on TOFU and MUSE-Books, RFU removes much of the forget behavior while keeping retain utility close to that of the unedited model. Across our experiments, RFU reduces forget behavior while preserving retain utility, and basis ablations confirm that these results depend on applying the shrinkage in the RFU basis.

AI USAGE

In this work, we used generative AI tools to help with implementing methods, particularly baseline methods, and to help with running experiments. We also used generative AI tools to edit the paper to improve readability. We reviewed all experimental outputs and edits for correctness. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY

We hope to have provided sufficient methodological and experimental detail to reproduce the results in this work. We specify the RFU procedure in the main text and as explicit pseudocode in the appendix. Appendix F further documents the model checkpoints and datasets used, baseline configurations, evaluation protocols, and computational requirements. We additionally report complete per-setting results and sensitivity analyses for key methodological choices and hyperparameters. Together, these details are intended to permit independent implementation and reproduction of the method and experiments.

REFERENCES

Saleh Ashkboos, Amirkeivan Mohtashami, Maximilian L. Croci, Bo Li, Pashmina Cameron, Martin Jaggi, Dan Alistarh, Torsten Hoefler, and James Hensman. Quarot: Outlier-free 4-bit inference in rotated llms, 2024. URL https://arxiv.org/abs/2404.00456.

Nora Belrose, David Schneider-Joseph, Shauli Ravfogel, Ryan Cotterell, Edward Raff, and Stella Biderman. Leace: Perfect linear concept erasure in closed form. *Advances in Neural Information Processing Systems*, 36:66044–66063, 2023.

Chengyi Cai, Zesheng Ye, Jiangchao Yao, Jianzhong Qi, Bo Han, Xiaolu Zhang, Feng Liu, and Jun Zhou. Per-parameter task arithmetic for unlearning in large language models, 2026. URL https://arxiv.org/abs/2601.22030.

Nicholas Carlini, Florian Tramer, 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.

Kent Chang, Mackenzie Cramer, Sandeep Soni, and David Bamman. Speak, memory: An archaeology of books known to chatgpt/gpt-4. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 7312–7327, 2023.

Jerry Chee, Yaohui Cai, Volodymyr Kuleshov, and Christopher De Sa. Quip: 2-bit quantization of large language models with guarantees, 2024. URL https://arxiv.org/abs/2307.13304.

Yijiang River Dong, Hongzhou Lin, Mikhail Belkin, Ramon Huerta, and Ivan Vulić. Undial: Self-distillation with adjusted logits for robust unlearning in large language models. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 8827–8840, 2025.

Vineeth Dorna, Anmol Mekala, Wenlong Zhao, Andrew McCallum, Zachary C. Lipton, J. Zico Kolter, and Pratyush Maini. Openunlearning: Accelerating llm unlearning via unified benchmarking of methods and metrics, 2025. URL https://arxiv.org/abs/2506.12618.

Chongyu Fan, Jiancheng Liu, Licong Lin, Jinghan Jia, Ruiqi Zhang, Song Mei, and Sijia Liu. Simplicity prevails: Rethinking negative preference optimization for llm unlearning. *Advances in Neural Information Processing Systems*, 38:1540–1567, 2026.

Junfeng Fang, Houcheng Jiang, Kun Wang, Yunshan Ma, Jie Shi, Xiang Wang, Xiangnan He, and Tat-Seng Chua. Alphaedit: Null-space constrained knowledge editing for language models. In *International Conference on Learning Representations*, volume 2025, pp. 16366–16396, 2025.

Jack Foster, Stefan Schoepf, and Alexandra Brintrup. Fast machine unlearning without retraining through selective synaptic dampening. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pp. 12043–12051, 2024.

Shengyuan Hu, Yiwei Fu, Steven Wu, and Virginia Smith. Jogging the memory of unlearned models through targeted relearning attacks. In *ICML 2024 Workshop on Foundation Models in the Wild*, 2024.

Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. *arXiv preprint arXiv:2212.04089*, 2022.

Matthew Khoriaty, Andrii Shportko, Gustavo Mercier, and Zach Wood-Doughty. Don't forget it! conditional sparse autoencoder clamping works for unlearning, 2025. URL https://arxiv.org/abs/2503.11127.

Sangamesh Kodge, Gobinda Saha, and Kaushik Roy. Deep unlearning: Fast and efficient gradient-free approach to class forgetting. *arXiv preprint arXiv:2312.00761*, 2023.

Jea Kwon, Dong-Kyum Kim, Jiwon Kim, Yonghyun Kim, Woong Kook, and Meeyoung Cha. Ai engram: In search of memory traces in artificial intelligence. *arXiv preprint arXiv:2606.14997*, 2026.

Jooyoung Lee, Thai Le, Jinghui Chen, and Dongwon Lee. Do language models plagiarize? In *Proceedings of the ACM Web Conference 2023*, pp. 3637–3647, 2023.

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.

Chris Y Liu, Yaxuan Wang, Jeffrey Flanigan, and Yang Liu. Large language model unlearning via embedding-corrupted prompts. *Advances in Neural Information Processing Systems*, 37: 118198–118266, 2024.

Zechun Liu, Changsheng Zhao, Igor Fedorov, Bilge Soran, Dhruv Choudhary, Raghuraman Krishnamoorthi, Vikas Chandra, Yuandong Tian, and Tijmen Blankevoort. Spinquant: Llm quantization with learned rotations, 2025. URL https://arxiv.org/abs/2405.16406.

Nils Lukas, Ahmed Salem, Robert Sim, Shruti Tople, Lukas Wutschitz, and Santiago Zanella-Béguelin. Analyzing leakage of personally identifiable information in language models. In *2023 IEEE Symposium on Security and Privacy (SP)*, pp. 346–363. IEEE, 2023.

Aengus Lynch, Phillip Guo, Aidan Ewart, Stephen Casper, and Dylan Hadfield-Menell. Eight methods to evaluate robust unlearning in llms. *arXiv preprint arXiv:2402.16835*, 2024.

Pratyush Maini, Zhili Feng, Avi Schwarzschild, Zachary C. Lipton, and J. Zico Kolter. Tofu: A task of fictitious unlearning for llms, 2024. URL https://arxiv.org/abs/2401.06121.

Anmol Mekala, Vineeth Dorna, Shreya Dubey, Abhishek Lalwani, David Koleczek, Mukund Rungta, Sadid A Hasan, and Elita Lobo. Alternate preference optimization for unlearning factual knowledge in large language models. In *Proceedings of the 31st International Conference on Computational Linguistics*, pp. 3732–3752, 2025.

Martin Pawelczyk, Seth Neel, and Himabindu Lakkaraju. In-context unlearning: Language models as few shot unlearners. *arXiv preprint arXiv:2310.07579*, 2023.

Yifu Qiu, Zheng Zhao, Yftah Ziser, Anna Korhonen, Edoardo M Ponti, and Shay B Cohen. Spectral editing of activations for large language model alignment. *Advances in Neural Information Processing Systems*, 37:56958–56987, 2024.

Herbert Ramoser, Johannes Muller-Gerking, and Gert Pfurtscheller. Optimal spatial filtering of single trial eeg during imagined hand movement. *IEEE transactions on rehabilitation engineering*, 8(4): 441–446, 2000.

Weijia Shi, Jaechan Lee, Yangsibo Huang, Sadhika Malladi, Jieyu Zhao, Ari Holtzman, Daogao Liu, Luke Zettlemoyer, Noah Smith, and Chiyuan Zhang. Muse: Machine unlearning six-way evaluation for language models. In *International Conference on Learning Representations*, volume 2025, pp. 27797–27818, 2025.

Minh Hieu Vu and Tan Nguyen. Angular steering: Behavior control via rotation in activation space. *Advances in Neural Information Processing Systems*, 38:121653–121690, 2026.

Qizhou Wang, Jin Peng Zhou, Zhanke Zhou, Saebyeol Shin, Bo Han, and Kilian Q Weinberger. Rethinking llm unlearning objectives: A gradient perspective and go beyond. *arXiv preprint arXiv:2502.19301*, 2025a.

Xu Wang, Zihao Li, Benyou Wang, Yan Hu, and Difan Zou. Model unlearning via sparse autoencoder subspace guided projections, 2025b. URL https://arxiv.org/abs/2505.24428.

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.

Rongzhe Wei, Peizhi Niu, Hans Hao-Hsun Hsu, Ruihan Wu, Haoteng Yin, Mohsen Ghassemi, Yifan Li, Vamsi Potluru, Eli Chien, Kamalika Chaudhuri, et al. Do llms really forget? evaluating unlearning with knowledge correlation and confidence awareness. *Advances in Neural Information Processing Systems*, 38:93073–93111, 2026.

Puning Yang, Qizhou Wang, Zhuo Huang, Tongliang Liu, Chengqi Zhang, and Bo Han. Exploring criteria of loss reweighting to enhance llm unlearning. *arXiv preprint arXiv:2505.11953*, 2025.

Yuanshun Yao and Xiaojun Xu. 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.

Zhiwei Zhang, Fali Wang, Xiaomin Li, Zongyu Wu, Xianfeng Tang, Hui Liu, Qi He, Wenpeng Yin, and Suhang Wang. Catastrophic failure of llm unlearning via quantization. In *International Conference on Learning Representations*, volume 2025, pp. 74925–74948, 2025.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Zican Dong, Yupeng Hou, Beichen Zhang, Yingqian Min, Junjie Zhang, Peiyu Liu, et al. A survey of large language models. *Frontiers of Computer Science*, 20(12):2012627, 2026.

A Derivations and proofs

Throughout the appendices, we write

$$\alpha_i := \frac{1}{1+\beta\lambda_i}, \qquad D_\beta := \mathrm{diag}(\alpha_1, \ldots, \alpha_{d_{\mathrm{in}}}),$$

for the shrinkage factor of direction i and the corresponding diagonal shrinkage matrix.

A.1 Deriving the shrinkage rule in generalized eigenvector coordinates

Here we spell out the algebra suppressed in §3.2. Since $A_\rho \succ 0$, the generalized eigenvectors form an invertible basis. With $C = WU^{-\top}$ and $C' = W'U^{-\top}$, we have $W = CU^\top$ and $W' = C'U^\top$. The regularized retain term in equation 2 is

$$\begin{aligned}\mathbb{E}_R\|(W'-W)a\|_2^2 + \rho\|W'-W\|_F^2 &= \mathrm{tr}\big((W'-W)A_\rho(W'-W)^\top\big)\\ &= \|C'-C\|_F^2,\end{aligned}$$

where the second line uses $U^\top A_\rho U = I$. Likewise,

$$\begin{aligned}\mathbb{E}_F\|W'a\|_2^2 &= \mathrm{tr}\big(W'A_fW'^\top\big)\\ &= \mathrm{tr}\big(C'\Lambda C'^\top\big),\end{aligned}$$

using $U^\top A_f U = \Lambda$. Because Λ is diagonal, the objective separates across columns, and column i is minimized by

$$C'[:,i] = \alpha_i\, C[:,i].$$

Thus the diagonal shrinkage matrix D_β follows directly from the original forget-retain objective.

A.2 Euclidean ridge reinterpretation of RFU shrinkage

Section 3 derives the shrinkage factors from the retain/forget output objective and then applies them in the RFU basis V. Once V is fixed, the update can be written as the solution of a diagonal Euclidean ridge problem. Let $B = WV$ and $B' = W'V$. Since $V^\top V = I$,

$$B^\star = \arg\min_{B'} \sum_{i=1}^{d_{\mathrm{in}}} \left(\|B[:,i] - B'[:,i]\|_2^2 + \beta\lambda_i\|B'[:,i]\|_2^2\right).$$

The columns decouple, and minimizing the i-th term gives

$$B^\star[:,i] = \alpha_i B[:,i].$$

Mapping back with $V^\top$ yields $W_{\mathrm{RFU}} = WVD_\beta V^\top$.

This objective gives another view of the RFU update. In the original objective, the size of a change is measured using the retain-weighted metric. After replacing the generalized eigenbasis with the orthonormal basis V, the corresponding ridge objective measures that change in ordinary Euclidean weight space.

A.3 Applying the shrinkage factors in the RFU basis

RFU computes shrinkage factors from the generalized eigenvalues but applies them in the RFU basis $V = \mathrm{orth}(U)$, whose columns are generally not generalized eigenvectors: typically $V^\top A_\rho V \neq I$, $V^\top A_f V \neq \Lambda$, and the generalized Rayleigh quotient of v_i need not equal λ_i. Here we show that the RFU operator nonetheless depends on U only through its nested leading subspaces, and that it shrinks those subspaces in the intended order.

Since $U^\top A_\rho U = I$, U is invertible, and its unpivoted QR decomposition $U = VT$ has V orthogonal and T upper triangular with nonzero diagonal. Because T is upper triangular, the first k columns of U are combinations of the first k columns of V and vice versa, so

$$\mathrm{span}(v_1, \ldots, v_k) = \mathrm{span}(u_1, \ldots, u_k) \qquad \text{for every } k.$$

Let $P_k = \sum_{j \le k} v_j v_j^\top$ be the Euclidean projector onto this subspace. Since the eigenvalues are sorted in decreasing order, the shrinkage factors α_i are nondecreasing in i, and

$$VD_\beta V^\top = \alpha_{d_{\text{in}}} I - \sum_{k=1}^{d_{\text{in}}-1} (\alpha_{k+1} - \alpha_k)\, P_k, \qquad \alpha_{k+1} - \alpha_k \ge 0.$$

To verify this, apply both sides to v_j: the right-hand side gives $\alpha_{d_{\text{in}}} v_j - \sum_{k=j}^{d_{\text{in}}-1} (\alpha_{k+1} - \alpha_k) v_j = \alpha_j v_j$.

This shows that the operator is determined by the shrinkage factors and the nested subspaces alone. When adjacent eigenvalues are equal, the corresponding coefficient $\alpha_{k+1} - \alpha_k$ vanishes, so arbitrary basis choices within a repeated-eigenvalue subspace do not affect the edit. This also shows that the most forget-heavy subspaces are shrunk the most. For $x \in \text{span}(u_1, \ldots, u_k)$,

$$\|VD_\beta V^\top x\|_2^2 = \sum_{j \le k} \alpha_j^2 (v_j^\top x)^2 \le \alpha_k^2 \|x\|_2^2,$$

while every x orthogonal to that subspace satisfies $\|VD_\beta V^\top x\|_2 \ge \alpha_{k+1} \|x\|_2$. RFU thus shrinks every input in the leading k-dimensional forget-over-retain subspace at least as much as vectors outside it.

The same decomposition also bounds the change in output, as referenced in §4.1. For any x orthogonal to the leading k-dimensional subspace,

$$\|(W_{\text{RFU}} - W)x\|_2 \le \|W\|_2 \,(1 - \alpha_{k+1})\, \|x\|_2.$$

Take k large enough so that every eigenvalue after k is at most the layer median eigenvalue. If that layer median equals the reference scale λ_{med}, then the β rule gives $1 - \alpha_{k+1} \le \tau$, so

$$\|(W_{\text{RFU}} - W)x\|_2 \le \tau \|W\|_2 \|x\|_2.$$

Thus the target attenuation τ directly bounds the output change on inputs outside the leading, forget-heavy half of the spectrum.

A.4 Equivalence of generalized eigenvalue and whitening formulations

In practice, we solve the generalized eigenproblem in Section 3.1 through its symmetric whitened form. The following proposition shows that they are the same problem under the whitening map.

Proposition A.1 (Equivalence of formulations)**.** *Let $A_f \succeq 0$ and $A_\rho \succ 0$, and define $B := A_\rho^{-1/2} A_f A_\rho^{-1/2}$. For any eigenpair (λ, q) of B with $\|q\|_2 = 1$, the vector $u := A_\rho^{-1/2} q$ satisfies $A_f u = \lambda A_\rho u$ and $u^\top A_\rho u = 1$. The mapping also works in reverse: if u solves $A_f u = \lambda A_\rho u$, then $q = A_\rho^{1/2} u$ is an eigenvector of B with the same eigenvalue.*

Proof. Suppose $Bq = \lambda q$. Multiplying by $A_\rho^{1/2}$ gives $A_f A_\rho^{-1/2} q = \lambda A_\rho^{1/2} q$. Setting $u = A_\rho^{-1/2} q$ and using $A_\rho^{1/2} q = A_\rho u$ gives

$$A_f u = \lambda A_\rho u.$$

Moreover,

$$u^\top A_\rho u = q^\top A_\rho^{-1/2} A_\rho A_\rho^{-1/2} q = q^\top q = 1.$$

For the reverse direction, suppose $A_f u = \lambda A_\rho u$ and set $q = A_\rho^{1/2} u$. Substituting $u = A_\rho^{-1/2} q$ gives $Bq = \lambda q$. □

The two formulations therefore produce the same generalized eigenvalues and the same generalized eigenvectors up to the whitening map $A_\rho^{-1/2}$. RFU orthonormalizes the ordered generalized eigenvectors to get V; as noted in §3.3 and shown in Appendix A.3, the individual columns of V are generally no longer generalized eigenvectors and need not have Rayleigh quotients equal to the paired λ_i.

B THE EXACT MINIMIZER $W^\star_\beta$

Section 3.2 derives the shrinkage factors α_i from the layer objective equation 2. Applying those factors in the generalized eigenbasis gives the exact minimizer of that objective; RFU instead applies them in an orthonormal basis (§3.3). This section records the exact minimizer, explains why it can move the weights far from W, compares the two realizations empirically, and gives a simple model of the geometry behind the difference.

The exact minimizer. Setting the gradient of equation 2 to zero gives $(W' - W)A_\rho + \beta W' A_f = 0$, hence

$$W^\star_\beta = WA_\rho(A_\rho + \beta A_f)^{-1} = WU^{-\top} D_\beta U^\top = WA_\rho^{1/2} Q D_\beta Q^\top A_\rho^{-1/2}. \tag{10}$$

The second form follows from $A_\rho = U^{-\top}U^{-1}$ and $A_\rho + \beta A_f = U^{-\top}(I + \beta\Lambda)U^{-1}$, both consequences of $U^\top A_\rho U = I$ and $U^\top A_f U = \Lambda$. The third uses $U = A_\rho^{-1/2} Q$ (Appendix A.4). In the notation of §3.2, $W^\star_\beta = CD_\beta U^\top$ with $C = WU^{-\top}$: the shrinkage factors applied in the generalized eigenbasis.

Why the exact minimizer can make large Euclidean changes. The objective controls the retain-weighted size of the change but not its Euclidean size. Since $W' = W$ is feasible and both terms are nonnegative,

$$\operatorname{tr}\big((W^\star_\beta - W)A_\rho(W^\star_\beta - W)^\top\big) \le \beta \operatorname{tr}(WA_fW^\top),$$

but the only Euclidean control this implies is $\|W^\star_\beta - W\|_F \le \sqrt{\beta \operatorname{tr}(WA_fW^\top)/\rho}$, which is loose because ρ is small by design.

The same fact appears in the basis view. For any invertible basis P, write $W = CP^\top$ with $C = WP^{-\top}$ and consider the edit $W' = CDP^\top$, which scales the response to each direction of P, with $D = \operatorname{diag}(\alpha_i)$ and $0 < \alpha_i \le 1$. Then $W' - W = -C(I - D)P^\top$, and using $\|AB\|_F \le \|A\|_F \|B\|_2$,

$$\|W' - W\|_F \le \|W\|_F \, \|P^{-1}\|_2 \max_i(1 - \alpha_i) \, \|P\|_2 = \kappa(P) \max_i(1 - \alpha_i) \, \|W\|_F.$$

For RFU, $P = V$ is orthonormal, $\kappa(V) = 1$, and this is equation 6. For the exact minimizer, $P = U = A_\rho^{-1/2} Q$, whose singular values are those of $A_\rho^{-1/2}$, so

$$\kappa(U) = \sqrt{\kappa(A_\rho)} = \sqrt{\frac{\mu_{\max}(A_r) + \rho}{\mu_{\min}(A_r) + \rho}},$$

where μ denotes eigenvalues of A_r. With $\rho = \eta \operatorname{tr}(A_r)/d_{\text{in}}$ and $\mu_{\max} \ge \operatorname{tr}(A_r)/d_{\text{in}}$, a (nearly) singular A_r gives $\kappa(A_\rho) \gtrsim 1/\eta$, so $\kappa(U)$ can be of order 10^3 at our default $\eta = 10^{-6}$. The bound can be loose, but the objective still favors large changes along directions where A_ρ is small: such changes are nearly free under the retain term and can still reduce forget output. The planted model at the end of this section gives an explicit instance, and the MUSE-Books results below show it occurring in practice.

Implementation convention. For all $W^\star$ experiments, we use the same regularized retain covariance $A_\rho = A_r + \rho I$ as RFU, with $\rho = \eta \operatorname{tr}(A_r)/d_{\text{in}}$. We use $\eta = 10^{-6}$ unless explicitly testing a $100\times$ larger ridge, where $\eta = 10^{-4}$. In Table 3, RFU and $W^\star$ therefore start from the same (A_f, A_ρ, β), and their only difference is the basis where the shrinkage factors are applied.

What is being compared. $W^\star_\beta$ and RFU share the shrinkage factors α_i but apply them in different bases: $W^\star_\beta$ in the generalized eigenbasis U, RFU in the RFU basis V. Thus, RFU obeys the data-independent Euclidean bound equation 6, whereas $W^\star_\beta$ obeys only the $\kappa(U)$-dependent bound above and can exploit directions assigned little energy by the estimated retain metric.

Empirical comparison. We compute the exact minimizer from the same covariance estimates used by RFU and sweep β on three cells spanning the benchmark families (Table 3). The comparison depends strongly on how well the retain covariance represents the behavior being evaluated.

Table 3: RFU compared with the exact minimizer $W^\star_\beta$. RFU is shown at the strength given by the β rule. For $W^\star$, we show both that same β and one point from its own β sweep: $\beta = 0.005$ on TOFU, $\beta = 0.001$ on MUSE-Books, and $\beta = 0.010$ on WMDP-cyber. On MUSE-Books, even the smallest tested value collapses retention. Each entry reports forget / retain: f_R / MU on TOFU, verbatim memorization / retain knowledge on MUSE, and target accuracy / MMLU on WMDP. Lower forget and higher retain are better.

Setting	(forget ↓, retain ↑)	Unedited	RFU	$W^\star$ (swept β)	$W^\star$ (calibrated β)
TOFU Llama-3.2-3B f01	(f_R, MU)	0.986, 0.665	0.471, 0.673	0.303, 0.673 (β=0.005)	0.046, 0.000
MUSE-Books	(verbmem, knowmem)	0.997, 0.686	0.342, 0.703	0.078, 0.008 (β=0.001)	0.019, 0.006
WMDP-cyber	(WMDP, MMLU)	0.433, 0.589	0.421, 0.589	0.392, 0.584 (β=0.010)	0.368, 0.574

On TOFU Llama-3.2-3B, where the retain split closely matches the behavior evaluated by the retain metric, the exact minimizer outperforms RFU: at $\beta = 0.005$ it reaches $f_R = 0.303$ at MU 0.673, compared with $f_R = 0.471$ at the same MU for RFU. A warm start experiment shows a similar result: SimNPO started from $W^\star$ ends at $f_R = 0.299 \pm 0.018$, versus 0.395 ± 0.015 from RFU+SimNPO, while final MU is essentially the same (0.652 ± 0.014 versus 0.651 ± 0.011). This $W^\star$ result uses a β selected from the sweep on this cell; unlike RFU, we do not have a validated rule for choosing its strength automatically.

On MUSE-Books, the exact minimizer behaves very differently. Even at its smallest tested β, retain knowmem falls from 0.686 to 0.008, although the objective records little distortion under the estimated retain geometry. Early attention weights change by as much as $13\times$ their original Frobenius norm, especially where A_r is ill-conditioned. Increasing the retain ridge by $100\times$ still does not prevent the collapse.

WMDP-cyber falls between the TOFU and MUSE cases. At matched operating points, the exact minimizer lowers cyber accuracy from 0.433 to 0.392 while MMLU changes from 0.589 to 0.584 (Table 3). RFU at its default strength makes a smaller change (0.433 to 0.421) with essentially unchanged MMLU. Stronger RFU settings are shown in Appendix C.2.

Summary. For the exact objective, directions with little retain energy are cheap to change, even when the finite retain sample misses behavior that matters at evaluation time. When the estimated retain covariance is representative, as in the TOFU example above, the exact minimizer can outperform RFU. However, when the estimate is narrow, it can instead produce very large Euclidean changes in poorly covered directions. Because RFU applies the shrinkage factors through the bounded operator in equation 6, it cannot make such large changes. Calibration also differs between the exact minimizer and RFU; applying RFU's calibrated β to the exact minimizer on Llama-3.2-3B collapses MU to 0.000 (Table 3), while the same β remains usable with RFU (Appendix C.1).

Widening the retain estimate rescues the exact minimizer. To understand whether the result on MUSE is due to issues with estimation, we add generic text to the retain estimate as in Appendix D.5 and repeat the experiment.

$$A_r \leftarrow (1-\gamma)A_r^{\text{bench}} + \gamma A_r^{\text{generic}},$$

with WikiText-103 supplying the generic component.

Table 4: The exact minimizer on MUSE-Books with the composite retain estimate (forget verbmem ↓, retain knowmem ↑; unedited 0.997/0.686, RFU at the calibrated β: 0.342/0.703). At $\gamma = 0$ the edit destroys retention at every β; mixing in the generic slice restores it in proportion to γ, and at $\beta = 0.001$, $\gamma = 0.1$ the exact minimizer reaches an operating point comparable to RFU's.

β	γ	forget	retain
0.001	0	0.078	0.008
0.001	0.1	0.332	0.675
0.001	0.3	0.633	0.680
0.005	0	0.049	0.004
0.005	0.1	0.107	0.217
0.005	0.3	0.147	0.540

More generic retain mass protects retain knowledge but forgets less; at fixed γ, larger β forgets more and loses more retain knowledge. At $\beta = 0.001$, $\gamma = 0.1$, retain knowledge returns to within 0.011 of the unedited model while forgetting is comparable to RFU at the calibrated β. This supports the explanation above: the exact minimizer fails when the retain data used to estimate A_r covers less than the retain evaluation tests.

A planted model. A two-dimensional example isolates the same mechanism. Let u_F and u_R be unit forget and retain directions with $u_F^\top u_R = \cos\theta$, let $u_\perp = (u_F - \cos\theta\, u_R)/\sin\theta$, and define

$$A_f = \sigma_f^2 u_F u_F^\top + \varepsilon^2 I, \qquad A_r = \sigma_r^2 u_R u_R^\top + \varepsilon^2 I.$$

The isotropic floor ε^2 represents noise; for estimated covariances, it also represents directions that the finite retain sample barely constrains. The resulting two-by-two generalized eigenproblem has simple limiting behavior, and the same formulas extend to rank-r plants with a common angle.

First, as $\varepsilon \to 0$,

$$\varepsilon^2 \lambda_1 / \sigma_f^2 \to \sin^2\theta,$$

so λ_1 is controlled by the part of the forget component that retain does not share, that is, it depends on the part of u_F orthogonal to u_R. This is an idealized version of S_{sep} (§4.2); for a general retain covariance, the angle is measured after whitening by A_ρ.

Second, in the $\varepsilon \to 0$ limit, bounded orthogonal shrinkage can zero the $u_\perp$ direction, removing the $\sin^2\theta$ fraction of the forget response at vanishing retain cost. However, any further suppression must act on the component shared with retain. If d_R is the shrinkage factor on the remaining retain-aligned direction, the residual forget and retain terms obey

$$\frac{\sqrt{E_F^{\text{overlap}}}}{\sigma_f \cos\theta\, \|W u_R\|_2} + \frac{\sqrt{E_R}}{\sigma_r\, \|W u_R\|_2} = 1.$$

Because both terms are linear in d_R, reducing the shared forget response increases retain distortion proportionally.

Third, the exact minimizer can go beyond this frontier with a large Euclidean shear. As $\varepsilon \to 0$ with $\beta\lambda_1 \to \infty$, $W^\star$ drives its response to u_F to zero at any $\theta > 0$ while its measured retain distortion vanishes, but

$$\|W^\star - W\| \to \|\cot\theta\, W u_R + W u_\perp\|_2,$$

which diverges as the two populations align. With exact population covariances, this shear costs little under the retain metric; with finite estimated covariances, the same mechanism can exploit directions the retain sample failed to constrain. This is the geometry behind the $13 \times \|W\|$ MUSE edit in Table 3. Widening the retain estimate removes the poorly covered directions (Table 4), and RFU's bounded update equation 6 cannot produce an arbitrarily large shear.

C THE β RULE AND SEPARABILITY SCORING

In this section, we present the empirical evidence behind the β rule equation 7 and the spectral measurements underlying the separability score.

C.1 THE β RULE

For each edited layer we take the median generalized eigenvalue, and we define λ_{med} as the median of these values over layers. The β rule equation 7 sets $\beta = \tau/((1-\tau)\lambda_{\text{med}})$ with $\tau = 0.03$ in every setting; Appendix C.1 gives the attenuation-based rationale for this value. Table 5 then lists the resulting values and the effect of the edit at the calibrated β. Across settings, λ_{med} ranges from 0.148 to 1.269 and the calibrated β from 0.024 to 0.209. At the calibrated β, the forget metric is lower than for the unedited model in every setting. On TOFU, MU changes by between -0.042 and $+0.032$. Retain knowledge rises on MUSE-Books and falls by 0.064 on MUSE-News. The full β sweeps are shown in Figure 8 for TOFU, where the calibrated β is marked on each curve, and are described below for MUSE and WMDP.

Table 5: The β rule equation 7 in every setting. λ_{med} is the median over edited layers of each layer's median generalized eigenvalue, and the calibrated β is $\tau/((1-\tau)\lambda_{\text{med}})$ with $\tau = 0.03$. The last two columns give the change relative to the unedited model at the calibrated β: forget recall f_R and MU on TOFU, verbatim memorization and retain knowledge recall on MUSE, and subset accuracy and MMLU (in percentage points) on WMDP.

Setting	λ_{med}	Calibrated β	Δ forget	Δ retain
TOFU Llama-3.1-8B forget01	0.148	0.209	-0.549	$+0.032$
TOFU Gemma-3-4B forget01	0.207	0.149	-0.397	-0.042
TOFU Llama-3.2-3B forget01	0.284	0.109	-0.515	$+0.008$
TOFU Llama-3.2-1B forget01	0.364	0.085	-0.436	-0.028
TOFU Qwen-2.5-3B forget01	0.546	0.057	-0.282	-0.007
TOFU Gemma-3-4B forget05	0.714	0.043	-0.218	-0.023
TOFU Qwen-2.5-1.5B forget01	0.714	0.043	-0.167	-0.007
TOFU Llama-3.2-3B forget05	0.754	0.041	-0.305	$+0.020$
TOFU Llama-3.2-1B forget05	0.810	0.038	-0.265	-0.004
TOFU Qwen-2.5-3B forget05	0.881	0.035	-0.198	-0.002
TOFU Gemma-3-4B forget10	0.914	0.034	-0.159	-0.024
TOFU Qwen-2.5-1.5B forget05	0.934	0.033	-0.096	-0.010
TOFU Llama-3.2-3B forget10	0.937	0.033	-0.196	$+0.022$
TOFU Llama-3.2-1B forget10	0.939	0.033	-0.190	-0.002
TOFU Qwen-2.5-3B forget10	0.946	0.033	-0.123	-0.011
MUSE-Books	0.970	0.032	-0.655	$+0.017$
TOFU Qwen-2.5-1.5B forget10	0.982	0.032	-0.075	-0.011
MUSE-News	1.003	0.031	-0.113	-0.064
WMDP-bio	1.070	0.029	-1.0	0.0
WMDP-cyber	1.269	0.024	-1.2	0.0

Choice of τ. The target attenuation τ controls how quickly RFU moves from mild shrinkage near the spectral bulk to strong shrinkage on forget-heavy directions. For a relative eigenvalue $r = \lambda_i/\lambda_{\text{med}}$, the attenuation is $\tau r/(1-\tau+\tau r)$, and it reaches 50% when $r = (1-\tau)/\tau$.

With $\tau = 0.03$, a $10\times$ direction loses 23.6% of its response, a $100\times$ direction loses 75.6%, and half attenuation occurs at approximately $32\times$. In other words, directions moderately above the reference scale are changed cautiously, while clear spectral outliers are suppressed much more strongly.

We choose $\tau = 0.03$ for our experiments because it should be conservative but effective. At $\tau = 0.01$, even a $100\times$ direction is only half attenuated, but $\tau = 0.05$ and 0.10 would reach half attenuation at only $19\times$ and $9\times$ respectively.

The sweeps below show that $\tau = 0.03$ does indeed provide an intermediate, deliberately conservative point. In particular, WMDP-bio has $S_{\text{sep}} \approx 24$ and WMDP-cyber has $S_{\text{sep}} \approx 114$, which is qualitatively consistent with their different selectivity.

C.2 β-SWEEP EVIDENCE ACROSS BENCHMARKS AND OBJECTIVES

We report the β sweeps on MUSE and WMDP.

MUSE-Books. On Llama-2-7B MUSE-Books, verbatim memorization and extraction strength fall steeply for β between 0.01 and 0.1, while retain knowledge stays at or above its unedited value. Retain knowledge collapses only between $\beta = 0.1$ and 0.2, well above the calibrated $\beta = 0.032$. This sweep underlies Table 9.

MUSE-News. MUSE-News is less selective. Forgetting grows with β through 0.1 while retain knowledge declines gradually, and at $\beta = 0.5$ forget and retain behavior collapse together. At the calibrated $\beta = 0.031$, verbatim memorization falls from 0.576 to 0.463 and retain knowledge recall falls by 0.064.

WMDP-cyber versus WMDP-bio. WMDP provides the cleanest controlled comparison for the separability score because cyber and bio use the same Zephyr-7B backbone but have very different S_{sep} (≈ 114 for cyber and ≈ 24 for bio). We see that RFU drops WMDP-cyber by 1.21 pp and WMDP-bio by 1.02 pp while leaving MMLU essentially unchanged; as β increases, cyber accuracy falls faster than bio accuracy (Figure 6). This shows that the same model can allow for a strong edit on the domain with the concentrated spectrum but a weaker edit on the domain with the diffuse spectrum.

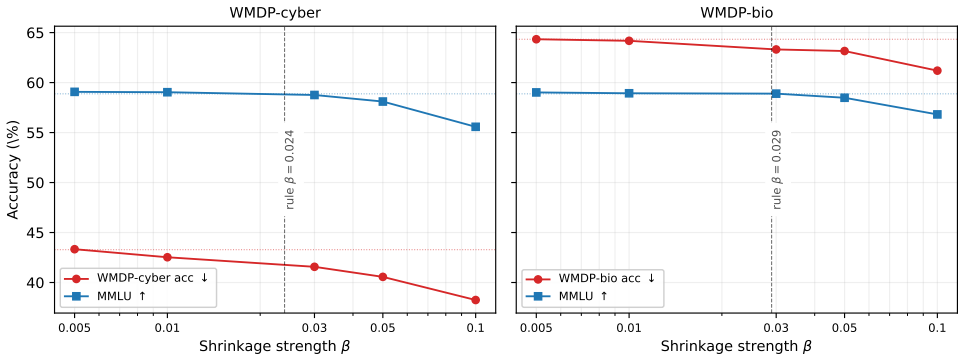

Figure 6: Standalone RFU β-sweeps on the same Zephyr-7B backbone for WMDP-cyber (left) and WMDP-bio (right). Forget-axis (↓) is subdomain accuracy; retain-axis (↑) is MMLU accuracy. Dotted lines show untreated baselines and dashed lines the calibrated β. Cyber accuracy ($S_{\text{sep}} \approx 114$) falls faster with β than bio accuracy ($S_{\text{sep}} \approx 24$); on both, MMLU stays within 0.8 pp of the unedited model through $\beta = 0.05$.

Warm start sweeps. We also sweep β for RFU+SimNPO on Llama-3.2-1B, Qwen-2.5-1.5B, Qwen-2.5-3B, and Gemma-3-4B over all three TOFU splits at $N = 5$ seeds (Figure 7). Across the sweep, f_R decreases over a broad range of RFU strengths while MU stays comparatively stable until high β. The calibrated value is below the value that collapses MU in every experiment. The collapse also lines up more closely in $\beta\lambda_{\text{med}}$ than in raw β, showing the scale dependence the calibration rule is designed to absorb.

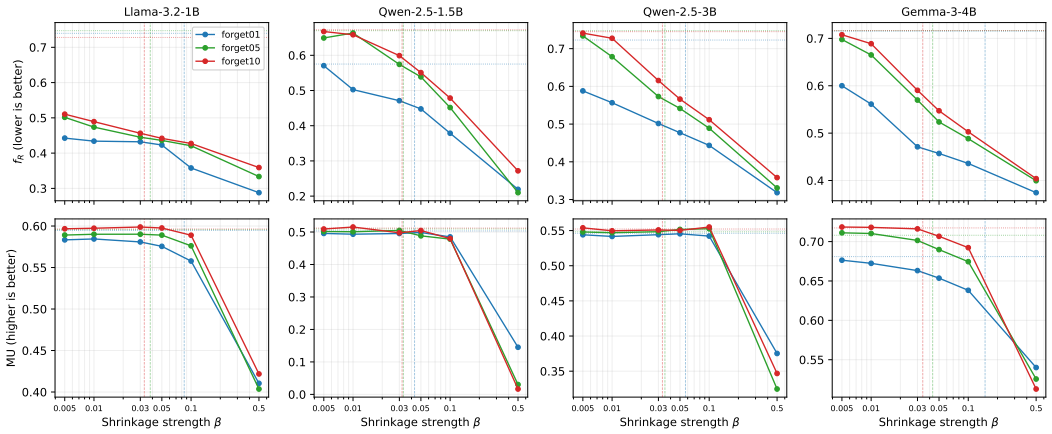

Figure 7: Warm start RFU+SimNPO β-sweeps across four architectures and all three TOFU forget splits at $N = 5$ seeds. The curves track f_R and MU as the RFU strength varies; dotted lines show matched no-RFU baselines and dashed lines mark calibrated β. In every cell, the calibrated value lies below the high-β retain-collapse region.

C.3 SEPARABILITY SCORE

Architecture, normalization, dataset, and metric differ across model families and benchmarks, so we do not expect S_{sep} to order every setting. WMDP-cyber and WMDP-bio provide the most controlled comparison because they share a backbone (Appendix C.2).

Per-setting spectra. Table 6 reports S_{sep}, the median across edited layers of $\lambda_1/\operatorname{median}(\{\lambda_i\})$, for the settings used in the main comparisons. We use a median across layers because small forget sets can make A_f rank-deficient, leaving numerical outliers in the lower spectral tail. (These outliers have little effect on RFU because near-zero eigenvalues have shrinkage factors close to one.)

Table 6: Separability score by setting. TOFU entries are for forget01.

Setting	S_{sep}
TOFU Llama-3.2-1B	200.6
TOFU Llama-3.2-3B	372.0
TOFU Qwen-2.5-1.5B	47.8
TOFU Qwen-2.5-3B	86.6
TOFU Gemma-3-4B	3439.4
TOFU Llama-3.1-8B	1681.6
MUSE-Books (Llama-2-7B)	342.5
MUSE-News (Llama-2-7B)	149.6
WMDP-cyber (Zephyr-7B)	114.5
WMDP-bio (Zephyr-7B)	24.0

Why Gemma is an outlier. Gemma-3 also uses the larger ridge $\eta = 10^{-4}$ (versus 10^{-6} elsewhere), and scales its embeddings by $\sqrt{d}$, giving layer-0 activations unusually large absolute energy. At the calibrated β, the four layer-0 attention projections account for 73% of the measured retain-side trace distortion, yet skipping layer 0 entirely changes no TOFU metric by more than 0.01 on any split. Downstream normalization largely removes this scale, so the large eigenvalues reflect a real difference in activations that does not carry through to the TOFU metrics.

Separating spectral scale from concentration. As the TOFU forget set grows from 1% to 10%, λ_{med} moves toward 1 and the calibrated β decreases in all five architectures. Because this scale effect also changes the amount of forgetting at the default operating point, the analysis at matched retain cost below is the cleaner test of whether spectral concentration itself predicts selectivity.

Matched retain cost test of selectivity. At the calibrated β, the amount of forgetting depends on both S_{sep} and the scale λ_{med}. We therefore run standalone β-sweeps for all TOFU combinations of architectures and forget splits and compare the forgetting each reaches at matched retain cost.

Allowing an MU decrease of at most 0.01 relative to the unedited model, the forget01 split, which has the largest S_{sep}, achieves the most forgetting in all five architectures (Table 2). The full forget01>forget05>forget10 ordering holds in four of the five; in Qwen-2.5-1.5B, forget05 and forget10 swap. Across the five architectures, S_{sep} reliably identifies the most selective split, but does not rank every less separable split exactly.

The ordering is clearest at small retain budgets. With a larger budget, the forget05 and forget10 frontiers often catch up: at a 0.05 budget, for example, Qwen-2.5-1.5B reaches forget reductions of approximately $0.25/0.24/0.24$ across forget01/05/10, compared with $0.20/0.10/0.15$ at the 0.01 threshold. Llama-3.2-1B and Gemma-3-4B show the same convergence. S_{sep} is therefore most informative at small retain cost.

The TOFU splits differ in both S_{sep} and forget set size, so this comparison alone does not show that S_{sep} causes the difference. The WMDP-cyber/bio comparison (Appendix C.2) is helpful here, as they share the same Zephyr-7B backbone. In that experiment, the more concentrated cyber spectrum allows for a selective standalone edit, but the more diffuse bio spectrum does not (Appendix C.2).

Table 7: Matched retain cost test of S_{sep} on TOFU. S_{sep} is the median across edited layers of $\lambda_1 / \operatorname{median}(\{\lambda_i\})$. "Forget reduction" is $f_R^{\text{unedited}} - f_R$ attainable while keeping MU no more than 0.01 below the unedited model. The forget01 split, with the most concentrated spectrum, achieves the greatest forgetting in all five architectures; the forget01>forget05>forget10 ordering holds in four of five.

	forget01		forget05		forget10	
Architecture	S_{sep}	Forget reduction	S_{sep}	Forget reduction	S_{sep}	Forget reduction
Llama-3.2-1B	201	0.417	25	0.305	15	0.262
Llama-3.2-3B	372	0.565	41	0.488	32	0.475
Qwen-2.5-1.5B	48	0.200	10	0.100	7	0.148
Qwen-2.5-3B	87	0.309	18	0.260	11	0.121
Gemma-3-4B	3439	0.158	343	0.080	157	0.041

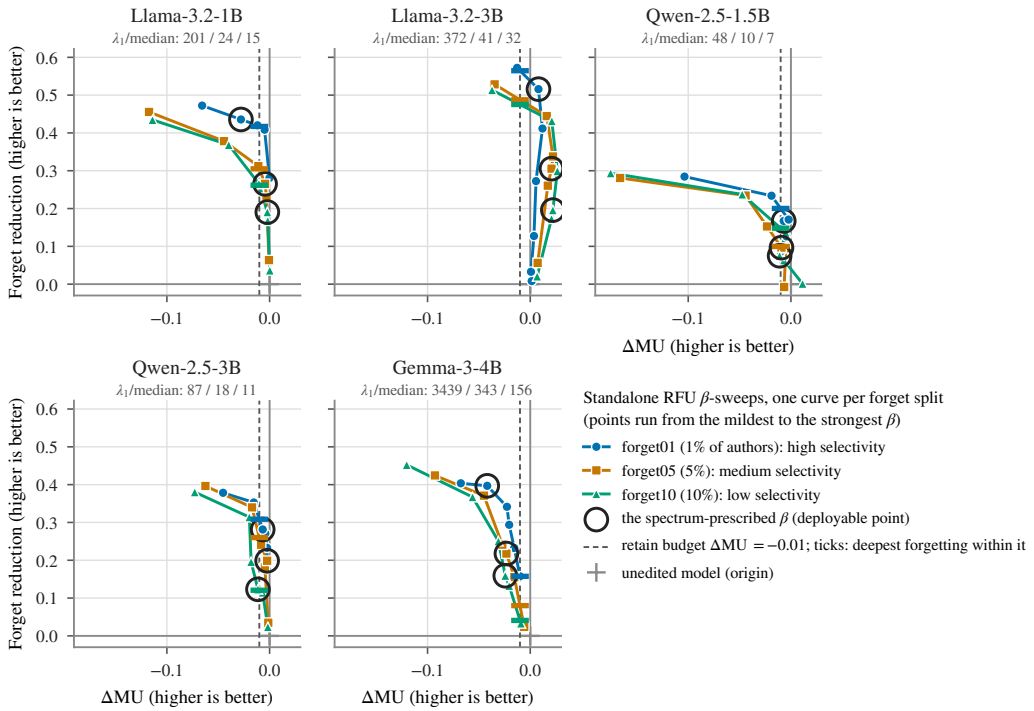

Figure 8: Standalone RFU forget-retain frontiers for forget01, forget05, and forget10 for each architecture. Ringed points are the calibrated β values, and the marked reference corresponds to an MU decrease of 0.01. At this small retain cost, forget01 permits the most forgetting in all five architectures; the frontiers become less separated as larger retain costs are allowed.

D Baselines and Ablations

This section compares RFU with other unlearning methods and ablates its components.

D.1 Full Standalone RFU Results

Tables 8 to 10 provide the complete numerical results summarized in §6.1. All RFU rows use the default strength from the β rule.

TOFU's distributional forget-quality analysis and qualitative generations are reported in Appendices E.4 and E.5. Standalone RFU is less selective on MUSE-News and WMDP-bio, and its effect on WMDP-cyber is selective but small; those sweeps are in Appendix C.2.

Table 8: Standalone RFU compared with the same pre-RFU model on TOFU across models, scales, and forget set sizes. Retain metric is MU; forget metric is forget-author ROUGE-L recall f_R. † marks cells where retain MU improves relative to the unedited model.

			Retain: MU ↑			Forget: f_R ↓		
Model	Split	β	Unedited	RFU	Δ	Unedited	RFU	Δ
Llama-3.2-1B	01	0.085	0.599	0.571	−0.028	0.865	0.430	−0.436
Llama-3.2-1B	05	0.038	0.600	0.596	−0.004	0.831	0.566	−0.265
Llama-3.2-1B	10	0.033	0.600	0.598	−0.002	0.819	0.629	−0.190
Llama-3.2-3B	01†	0.109	0.665	0.673	+0.008	0.986	0.471	−0.515
Llama-3.2-3B	05†	0.041	0.665	0.685	+0.020	0.925	0.620	−0.305
Llama-3.2-3B	10†	0.033	0.665	0.687	+0.022	0.925	0.729	−0.196
Llama-3.1-8B	01†	0.209	0.627	0.659	+0.032	0.998	0.450	−0.549
Qwen-2.5-1.5B	01	0.043	0.498	0.491	−0.007	0.607	0.440	−0.167
Qwen-2.5-1.5B	05	0.033	0.498	0.488	−0.010	0.651	0.555	−0.096
Qwen-2.5-1.5B	10	0.032	0.498	0.487	−0.011	0.649	0.574	−0.075
Qwen-2.5-3B	01	0.057	0.549	0.542	−0.007	0.759	0.477	−0.282
Qwen-2.5-3B	05	0.035	0.549	0.547	−0.002	0.769	0.571	−0.198
Qwen-2.5-3B	10	0.033	0.549	0.538	−0.011	0.764	0.641	−0.123
Gemma-3-4B	01	0.149	0.580	0.538	−0.042	0.825	0.428	−0.397
Gemma-3-4B	05	0.043	0.580	0.557	−0.023	0.850	0.632	−0.218
Gemma-3-4B	10	0.034	0.580	0.556	−0.024	0.862	0.703	−0.159

Table 9: Standalone RFU and baselines on Llama-2-7B MUSE-Books. SimNPO is the mean over the five no-RFU runs of Table 21; Engram uses its published strengths, as on TOFU. Retain metric is `r_knowmem`; forget metrics are `verbmem`, `extraction` (extraction strength), and `f_knowmem`; privacy metric `privleak` has ideal value 0.

	Retain ↑	Forget ↓			Privacy (→ 0)
	`r_knowmem`	`verbmem`	`extraction`	`f_knowmem`	`privleak`
Unedited	0.686	0.997	0.916	0.471	−57.3
RFU	0.703	0.342	0.202	0.443	−58.3
Δ	+0.017	−0.655	−0.714	−0.028	−1.0
SimNPO (N=5)	0.579	0.287	0.155	0.333	−51.7
Engram-adapt (α=1.0)	0.195	0.085	0.009	0.029	−58.4
Engram-plain (α=0.6)	0.069	0.080	0.008	0.003	−55.0

D.2 Standalone baselines

This subsection reports the standalone results for every baseline. Appendix F.3 describes how each baseline is configured and why we chose it.

At the fixed settings used for the grid, SSD often reaches lower f_R than RFU on the Llama and Qwen models, but on forget10 it loses 0.16–0.23 MU, on forget05/10 it loses 22–39 MMLU points, and on Gemma-3-4B its threshold selects no parameters. Task-vector negation keeps MMLU within 1.2 points on every cell, but on Gemma-3-4B it lowers TOFU MU from 0.580 to 0.15–0.19. Engram's MMLU loss reaches 14.5 points on forget10, and it has lower MU than RFU on eight of the ten forget05/10 cells (Tables 11 and 17). DU's rank sensitivity results are reported in Appendix D.3. Results from training-based methods are in Table 12, and results with RFU warm starts are in Table 22.

Baselines at 7B. Engram has much lower forget metrics on MUSE, but also much lower model capability—retain knowledge falls to 0.195 or lower on MUSE-Books and to at most 0.008 on MUSE-News, and MMLU falls to chance (0.23–0.25) on WMDP (Tables 9, 10 and 13). SimNPO forgets more than RFU on MUSE-Books, but at a larger retain cost.

Table 10: Standalone RFU and baselines on Llama-2-7B MUSE-News, with the metrics of Table 9. RFU uses its calibrated $\beta = 0.031$; SimNPO is the mean over the five no-RFU runs of Table 21; Engram uses its published strengths, as on TOFU.

	Retain ↑	Forget ↓			Privacy (→ 0)
	r_knowmem	verbmem	extraction	f_knowmem	privleak
Unedited	0.542	0.576	0.295	0.655	−99.8
RFU	0.478	0.463	0.204	0.565	−99.8
SimNPO (N=5)	0.472	0.499	0.234	0.595	−99.9
Engram-adapt (α=1.0)	0.008	0.015	0.008	0.003	−6.1
Engram-plain (α=0.6)	0.000	0.012	0.008	0.000	5.5

Table 11: Weight edits on the full TOFU grid. Entries are MU / f_R (higher/lower is better). Engram-adaptive is the primary Engram baseline; Engram-plain* is secondary because its official validation result only partially reproduced. DU uses its fixed rank pair (Appendix D.3). On Gemma-3-4B no parameter's forget importance exceeds $\alpha = 10$ times its retain importance, so SSD's selection step dampens nothing; its Gemma entries (°) are therefore those of the unedited model.

Model	Split	Unedited	RFU	DU	SSD	Task vector	Engram-adapt	Engram-plain*
Llama-3.2-1B	01	0.599/0.865	0.571/0.430	0.584/0.314	0.557/0.183	0.561/0.326	0.593/0.500	0.588/0.411
Llama-3.2-1B	05	0.600/0.831	0.596/0.566	0.588/0.533	0.460/0.201	0.556/0.455	0.569/0.462	0.550/0.361
Llama-3.2-1B	10	0.600/0.819	0.598/0.629	0.595/0.707	0.436/0.209	0.563/0.497	0.520/0.355	0.482/0.324
Llama-3.2-3B	01	0.665/0.986	0.673/0.471	0.659/0.302	0.619/0.210	0.655/0.402	0.664/0.566	0.662/0.382
Llama-3.2-3B	05	0.665/0.925	0.685/0.620	0.666/0.537	0.552/0.226	0.643/0.518	0.659/0.534	0.645/0.372
Llama-3.2-3B	10	0.665/0.925	0.687/0.729	0.663/0.699	0.480/0.227	0.649/0.600	0.639/0.474	0.602/0.362
Qwen-2.5-1.5B	01	0.498/0.607	0.491/0.440	0.486/0.276	0.433/0.130	0.457/0.371	0.489/0.477	0.486/0.415
Qwen-2.5-1.5B	05	0.498/0.651	0.488/0.555	0.488/0.597	0.392/0.188	0.448/0.368	0.482/0.510	0.476/0.453
Qwen-2.5-1.5B	10	0.498/0.649	0.487/0.574	0.491/0.638	0.325/0.199	0.456/0.403	0.470/0.427	0.461/0.391
Qwen-2.5-3B	01	0.549/0.759	0.542/0.477	0.546/0.268	0.475/0.212	0.543/0.355	0.548/0.512	0.545/0.456
Qwen-2.5-3B	05	0.549/0.769	0.547/0.571	0.557/0.644	0.402/0.222	0.509/0.417	0.536/0.537	0.533/0.480
Qwen-2.5-3B	10	0.549/0.764	0.538/0.641	0.555/0.730	0.323/0.213	0.533/0.467	0.510/0.461	0.499/0.424
Gemma-3-4B	01	0.580/0.825	0.538/0.428	0.565/0.297	0.580/0.825°	0.151/0.172	0.580/0.672	0.577/0.541
Gemma-3-4B	05	0.580/0.850	0.557/0.632	0.577/0.700	0.580/0.850°	0.159/0.367	0.578/0.713	0.568/0.576
Gemma-3-4B	10	0.580/0.862	0.556/0.703	0.576/0.807	0.580/0.862°	0.188/0.388	0.569/0.650	0.552/0.487

WMDP is different. RMU is designed for this benchmark and drives both domains to chance accuracy, while standalone RFU moves them only slightly. Here we only test whether adding RFU before RMU changes RMU's result.

RFU+RMU is statistically indistinguishable from RMU alone on the reported axes (paired t-test $p \geq 0.15$ throughout). Since RMU already reaches chance, RFU cannot lower the forget metric further; the result only shows that RFU does not interfere with RMU.

D.3 Forget-subspace projections

Deep Unlearning (DU) (Kodge et al., 2023) estimates forget and retain subspaces from layer activations and projects the forget-specific part out of each layer's input. Two hyperparameter ranks k_f and k_r determine the strength of the weight edit. We search the rank grid provided in the DU paper on the five TOFU forget01 cells. Every configuration with $k_r = 0$ collapses MU on every model, and $(2\%, 10\%)$ is the only pair that preserves MU on all five. Thus, we use $(2\%, 10\%)$ for the rest of the DU experiments.

We see that while $(k_f, k_r) = (2\%, 10\%)$ forgets the most on forget01, its forget reduction decreases on larger forget sets, where RFU forgets more than DU on seven of the ten forget05 and forget10 cells (Table 14). Moreover, its forget reduction falls off faster as forget set size increases; from forget01 to forget10, DU's reduction in f_R shrinks by a factor of 3 to 30 (on Qwen-2.5-1.5B, from 0.33 to 0.01), while RFU's shrinks by a factor of 2 to 3.

D.4 Subspace and rotation controls

These controls test two things: whether ranking directions by forget energy alone gives a selective basis, and whether the RFU basis matters once the shrinkage factors are fixed.

Table 12: Training-based methods on the full TOFU grid at their standard budgets. Entries are MU / f_R, mean over five seeds. RMU uses its default recipe on the four Llama/Qwen forget01 cells, and deterministic RFU is repeated for reference. Fine-tuning Gemma-3-4B on the retain split with the same budget and no forget term raises MU from 0.580 to 0.72–0.73, while forget recall stays at 0.87–0.92.

Model	Split	Unedited	RFU	SimNPO	NPO	WGA	RMU
Llama-3.2-1B	01	0.599/0.865	0.571/0.430	0.595/0.740	0.597/0.473	0.598/0.529	0.559/0.458
Llama-3.2-1B	05	0.600/0.831	0.596/0.566	0.596/0.747	0.450/0.314	0.593/0.486	—
Llama-3.2-1B	10	0.600/0.819	0.598/0.629	0.596/0.728	0.391/0.204	0.596/0.456	—
Llama-3.2-3B	01	0.665/0.986	0.673/0.471	0.653/0.825	0.668/0.439	0.657/0.504	0.652/0.504
Llama-3.2-3B	05	0.665/0.925	0.685/0.620	0.654/0.777	0.592/0.347	0.658/0.461	—
Llama-3.2-3B	10	0.665/0.925	0.687/0.729	0.650/0.751	0.658/0.289	0.647/0.366	—
Qwen-2.5-1.5B	01	0.498/0.607	0.491/0.440	0.503/0.575	0.499/0.506	0.502/0.524	0.499/0.605
Qwen-2.5-1.5B	05	0.498/0.651	0.488/0.555	0.509/0.670	0.496/0.416	0.509/0.572	—
Qwen-2.5-1.5B	10	0.498/0.649	0.487/0.574	0.512/0.672	0.359/0.171	0.512/0.571	—
Qwen-2.5-3B	01	0.549/0.759	0.542/0.477	0.546/0.723	0.556/0.503	0.550/0.581	0.555/0.753
Qwen-2.5-3B	05	0.549/0.769	0.547/0.571	0.548/0.748	0.510/0.552	0.550/0.573	—
Qwen-2.5-3B	10	0.549/0.764	0.538/0.641	0.552/0.745	0.476/0.420	0.557/0.545	—
Gemma-3-4B	01	0.580/0.825	0.538/0.428	0.681/0.715	0.555/0.582	0.676/0.512	—
Gemma-3-4B	05	0.580/0.850	0.557/0.632	0.708/0.716	0.687/0.548	0.686/0.429	—
Gemma-3-4B	10	0.580/0.862	0.556/0.703	0.717/0.716	0.693/0.656	0.685/0.377	—

Table 13: Standalone and composed methods on WMDP (Zephyr-7B; accuracy, chance = 0.25). RMU uses its canonical recipe (Li et al., 2024); RMU rows are mean ± sd over five seeds. RFU uses its default β-rule strength.

	WMDP-cyber		WMDP-bio	
Method	acc ↓	MMLU ↑	acc ↓	MMLU ↑
Unedited	0.433	0.589	0.643	0.589
RFU (calibrated β)	0.421	0.589	0.633	0.589
RMU	0.257±.014	0.572±.006	0.270±.013	0.566±.003
RFU+RMU	0.259±.007	0.570±.006	0.274±.017	0.564±.004
SimNPO	0.423±.012	0.581±.003	0.648±.006	0.582±.004
Engram-adapt (α=1.0)	0.260	0.250	0.247	0.238
Engram-plain (α=0.6)	0.267	0.231	0.251	0.236

Forget-only PCA. We also try choosing directions from A_f alone, without comparing them with retain activations. When the PCA edit is kept weak enough to preserve roughly the same MU as RFU, it barely forgets: the lowest f_R values are 0.830, 0.992, and 0.801 on Llama-3.2-1B, Llama-3.2-3B, and Gemma-3-4B, compared with RFU's 0.430, 0.471, and 0.428. Turning the PCA edit up enough to forget more destroys retention instead. High forget energy by itself is therefore not enough to identify a safe direction to edit, likely because the same direction may also be important for retain data.

Standalone rotation ablation. To isolate the rotation itself, we hold β and the shrinkage factors fixed and replace V with a random orthogonal basis or the identity (Table 15).

In settings where the spectrum is concentrated, the RFU basis gives the best forget-retain tradeoff—the random and identity bases forget less or pay a much larger retain cost. On WMDP-bio, where the spectrum is diffuse, none of the three bases changes bio accuracy by more than 1 pp.

D.4.1 Warm start basis ablation

We repeat the basis test with RFU+SimNPO on five architectures. In addition to identity and random orthogonal controls, we include forget-only PCA, which uses the eigensystem of A_f alone, on the Llama models and Gemma-3-4B.

Table 16 shows that the RFU basis gives lower f_R than the identity or a random basis on all five models, with MU equal to or better than the controls up to seed noise. On the Llama models and

Table 14: RFU versus DU at its fixed pair $(2\%, 10\%)$, selected on forget01, across the full TOFU grid. MMLU is the change in pp relative to the unedited model.

		RFU			DU		
Model	Split	MU	f_R	ΔMMLU	MU	f_R	ΔMMLU
Llama-3.2-1B	01	0.571	0.430	−2.8	0.584	0.314	−1.1
Llama-3.2-1B	05	0.596	0.566	−1.1	0.588	0.533	−0.3
Llama-3.2-1B	10	0.598	0.629	−1.2	0.595	0.707	0.0
Llama-3.2-3B	01	0.673	0.471	−4.3	0.659	0.302	−0.5
Llama-3.2-3B	05	0.685	0.620	−1.0	0.666	0.537	−0.5
Llama-3.2-3B	10	0.687	0.729	−1.0	0.663	0.699	0.0
Qwen-2.5-1.5B	01	0.491	0.440	−1.5	0.486	0.276	−1.6
Qwen-2.5-1.5B	05	0.488	0.555	−1.1	0.488	0.597	−1.3
Qwen-2.5-1.5B	10	0.487	0.574	−1.1	0.491	0.638	−1.0
Qwen-2.5-3B	01	0.542	0.477	−0.7	0.546	0.268	−0.7
Qwen-2.5-3B	05	0.547	0.571	−0.2	0.557	0.644	−0.4
Qwen-2.5-3B	10	0.538	0.641	−0.1	0.555	0.730	−0.3
Gemma-3-4B	01	0.538	0.428	−4.0	0.565	0.297	−0.8
Gemma-3-4B	05	0.557	0.632	−0.9	0.577	0.700	−1.0
Gemma-3-4B	10	0.556	0.703	−0.8	0.576	0.807	−0.6

Table 15: Standalone basis ablations using RFU's default shrinkage factors. "Retain" and "forget" denote the benchmark-specific pair shown in the setting column. Only the basis changes for the random and identity comparisons; the shrinkage factors stay fixed.

Setting (retain / forget)	Basis	Retain ↑	Forget ↓
MUSE-Books (`retain_km` / `verbmem`)	RFU basis	0.703	0.342
	random orthogonal	0.522	0.407
	identity	0.068	0.101
TOFU Llama-3.2-3B (MU / f_R)	RFU basis	0.673	0.470
	random orthogonal	0.630	0.617
	identity	0.636	0.602
TOFU Qwen-2.5-3B (MU / f_R)	RFU basis	0.542	0.478
	random orthogonal	0.545	0.628
	identity	0.534	0.672
WMDP-cyber (MMLU / WMDP acc.)	RFU basis	58.92%	42.07%
	random orthogonal	58.42%	43.38%
	identity	58.62%	43.28%

Gemma-3-4B, the controls still lower f_R well below that of SimNPO alone. On the Qwen models they help less: they lower f_R from 0.575 to 0.529–0.548 on Qwen-2.5-1.5B and from 0.723 to 0.671–0.694 on Qwen-2.5-3B. Forget-only PCA collapses MU at this β, because its eigenvalues measure forget energy without normalizing by retain energy.

D.5 General-capability cost and the retain estimate

TOFU's MU evaluates retention inside the benchmark, so we separately measure MMLU for every standalone edit (Table 17).

RFU's MMLU cost grows with the calibrated β within each model. It is at most 1.5 pp on the twelve cells with $\beta \leq 0.06$ and 2.8–6.9 pp on the four forget01 cells with larger β. DU costs less than RFU on the Llama models and on Gemma-3-4B forget01 and is within 0.2 pp of RFU elsewhere. The training-based methods cost at most 0.9 pp where measured, but on forget01 RFU also forgets more than NPO on four of five models and more than WGA and RMU on every model where they were

Table 16: Warm start basis ablation of RFU+SimNPO at $N = 5$ seeds. Llama-3.2-1B, Llama-3.2-3B, and both Qwen models use TOFU forget01; Gemma-3-4B uses forget05. All rows hold the shrinkage factors fixed and vary the basis. The no-RFU rows are the matched SimNPO runs.

Model	Basis	MU	f_R	Δf_R vs. RFU basis
Llama-3.2-1B	no RFU	0.595 ± 0.002	0.740 ± 0.038	—
Llama-3.2-1B	RFU basis	0.564 ± 0.004	$\mathbf{0.389 \pm 0.012}$	—
Llama-3.2-1B	identity	0.556 ± 0.002	0.479 ± 0.019	$+0.090$
Llama-3.2-1B	random orthogonal	0.548 ± 0.003	0.480 ± 0.017	$+0.091$
Llama-3.2-1B	forget-only PCA (A_f only)	0.015 ± 0.000	0.115 ± 0.010	-0.274
Llama-3.2-3B	no RFU	0.653 ± 0.003	0.825 ± 0.017	—
Llama-3.2-3B	RFU basis	0.651 ± 0.011	$\mathbf{0.395 \pm 0.015}$	—
Llama-3.2-3B	identity	0.627 ± 0.004	0.464 ± 0.011	$+0.069$
Llama-3.2-3B	random orthogonal	0.632 ± 0.010	0.478 ± 0.020	$+0.083$
Llama-3.2-3B	forget-only PCA (A_f only)	0.015 ± 0.000	0.131 ± 0.007	-0.264
Qwen-2.5-1.5B	no RFU	0.503 ± 0.000	0.575 ± 0.023	—
Qwen-2.5-1.5B	RFU basis	0.496 ± 0.001	$\mathbf{0.454 \pm 0.016}$	—
Qwen-2.5-1.5B	identity	0.490 ± 0.002	0.529 ± 0.006	$+0.075$
Qwen-2.5-1.5B	random orthogonal	0.490 ± 0.002	0.548 ± 0.014	$+0.094$
Qwen-2.5-3B	no RFU	0.546 ± 0.003	0.723 ± 0.024	—
Qwen-2.5-3B	RFU basis	0.546 ± 0.002	$\mathbf{0.460 \pm 0.012}$	—
Qwen-2.5-3B	identity	0.537 ± 0.003	0.694 ± 0.012	$+0.234$
Qwen-2.5-3B	random orthogonal	0.545 ± 0.004	0.671 ± 0.018	$+0.211$
Gemma-3-4B	no RFU	0.708 ± 0.004	0.716 ± 0.011	—
Gemma-3-4B	RFU basis	0.694 ± 0.006	$\mathbf{0.536 \pm 0.006}$	—
Gemma-3-4B	identity	0.671 ± 0.004	0.627 ± 0.011	$+0.091$
Gemma-3-4B	random orthogonal	0.673 ± 0.004	0.620 ± 0.009	$+0.084$
Gemma-3-4B	forget-only PCA (A_f only)	0.003 ± 0.004	0.032 ± 0.005	-0.504

Table 17: General-capability cost on the TOFU grid: MMLU change (pp) relative to the unedited model. DU uses its fixed rank pair; SSD, task-vector negation, and Engram (adaptive) use the settings of Table 11; training-based methods use their standard budgets. °SSD selects no parameters on Gemma-3-4B, so its entry is the unedited model. — marks cells not run.

Model	Split	β	Unedited	RFU	DU	SSD	Task vector	Engram	SimNPO	NPO	WGA	RMU
Llama-3.2-1B	01	0.085	0.479	−2.8	−1.1	−3.1	−0.1	−1.7	0.0	+0.1	−0.1	−0.6
Llama-3.2-1B	05	0.038	0.479	−1.1	−0.3	−22.2	−0.5	−5.1	−0.1	—	—	—
Llama-3.2-1B	10	0.033	0.479	−1.2	0.0	−24.9	−0.2	−14.5	−0.1	—	—	—
Llama-3.2-3B	01	0.109	0.616	−4.3	−0.5	−2.5	0.0	−1.0	0.0	0.0	—	−0.9
Llama-3.2-3B	05	0.041	0.616	−1.0	−0.5	−23.6	0.0	−3.2	0.0	—	—	—
Llama-3.2-3B	10	0.033	0.616	−1.0	0.0	−26.3	+0.1	−6.6	−0.1	—	—	—
Qwen-2.5-1.5B	01	0.043	0.592	−1.5	−1.6	−3.8	0.0	−0.4	0.0	0.0	−0.2	−0.2
Qwen-2.5-1.5B	05	0.033	0.592	−1.1	−1.3	−22.1	−0.1	−1.1	−0.1	−0.2	0.0	—
Qwen-2.5-1.5B	10	0.032	0.592	−1.1	−1.0	−23.8	0.0	−4.6	−0.1	−0.8	−0.1	—
Qwen-2.5-3B	01	0.057	0.649	−0.7	−0.7	−6.4	+0.2	−0.3	0.0	+0.1	0.0	−0.1
Qwen-2.5-3B	05	0.035	0.649	−0.2	−0.4	−31.5	+0.3	−0.9	−0.1	+0.4	0.0	—
Qwen-2.5-3B	10	0.033	0.649	−0.1	−0.3	−38.9	+0.3	−2.8	−0.2	—	—	—
Gemma-3-4B	01	0.149	0.553	−4.0	−0.8	0.0°	−0.8	−0.2	−0.1	−0.1	−0.1	—
Gemma-3-4B	05	0.043	0.553	−0.9	−1.0	0.0°	−0.8	−0.3	−0.1	−0.4	−0.4	—
Gemma-3-4B	10	0.034	0.553	−0.8	−0.6	0.0°	−1.2	−2.1	−0.1	—	—	—
Llama-3.1-8B	01	0.209	0.669	−6.9	—	—	—	—	—	—	—	—

run (Table 12). To test whether incomplete retain coverage contributes to the MMLU cost, we add generic text to the retain covariance so that

$$A_r \leftarrow (1-\gamma)A_r^{\text{bench}} + \gamma A_r^{\text{generic}},$$

where A_r^{generic} is collected from a matched token WikiText-103 slice. Then, we evaluate with the new A_r on the three TOFU cells with the largest MMLU drops. We evaluate with β calibrated on the original spectrum as well as with the β obtained by reapplying equation 7 to the composite spectrum (Table 18).

Table 18: Composite retain estimate on the three TOFU cells with the largest MMLU cost (MU ↑, f_R ↓, MMLU ↑). For each model, we vary γ and evaluate with both the original β and the β obtained by recalibrating on the composite spectrum.

Model	γ	β	MU	f_R	MMLU
Gemma-3-4B	0	0.149	0.538	0.428	0.513
Gemma-3-4B	0.1	0.149	0.549	0.435	0.537
Gemma-3-4B	0.3	0.149	0.553	0.446	0.543
Gemma-3-4B	0.1	0.241 (rule)	0.530	0.404	0.522
Gemma-3-4B	0.3	0.395 (rule)	0.495	0.375	0.513
Llama-3.2-3B	0	0.109	0.673	0.471	0.573
Llama-3.2-3B	0.1	0.109	0.677	0.456	0.590
Llama-3.2-3B	0.3	0.109	0.677	0.453	0.598
Llama-3.2-3B	0.1	0.157 (rule)	0.674	0.402	0.577
Llama-3.2-3B	0.3	0.228 (rule)	0.663	0.343	0.573
Llama-3.1-8B	0	0.209	0.659	0.450	0.600
Llama-3.1-8B	0.1	0.209	0.644	0.424	0.629
Llama-3.1-8B	0.3	0.209	0.642	0.443	0.639
Llama-3.1-8B	0.1	0.329 (rule)	0.650	0.378	0.603
Llama-3.1-8B	0.3	0.514 (rule)	0.650	0.311	0.588

At fixed β, adding generic text recovers MMLU on all three models while preserving strong forgetting, with f_R changing by less than 0.03. On Llama-3.2-3B, for example, the MMLU cost improves from -4.3 to -1.8 pp as γ moves from 0 to 0.3, while f_R remains near 0.45. Reapplying the β rule to the composite spectrum gives a larger β, which forgets more: at $\gamma = 0.3$, Llama-3.2-3B reaches $f_R = 0.343$ at the same MMLU value and approximately the same MU as we got with the original benchmark A_r. On Gemma-3-4B and Llama-3.1-8B, the calibrated β at $\gamma = 0.3$ becomes too large and starts to cost MU (Gemma-3-4B) or MMLU (Llama-3.1-8B).

D.6 Full warm start and composition results

SimNPO across the TOFU grid. Table 19 gives the complete endpoint comparison over five seeds that is summarized in §6.3.

Table 19: RFU+SimNPO compared with a matched no-RFU SimNPO run on TOFU. Values are mean ± std over $N = 5$ paired seeds at $\mathrm{lr} = 10^{-5}$, with RFU at the calibrated β from equation 7. The † symbol marks cells where retain MU improves relative to the matched no-RFU run.

			Retain: MU ↑			Forget: f_R ↓		
Model	Split	β	no-RFU	RFU+SimNPO	Δ	no-RFU	RFU+SimNPO	Δ
Llama-3.2-1B	01	0.085	0.595±.002	0.564±.004	−0.031	0.740±.038	0.389±.012	−0.351
Llama-3.2-1B	05	0.038	0.596±.002	0.589±.002	−0.007	0.747±.008	0.440±.006	−0.307
Llama-3.2-1B	10†	0.033	0.596±.001	0.599±.004	+0.003	0.728±.006	0.452±.004	−0.276
Llama-3.2-3B	01	0.109	0.653±.003	0.651±.011	−0.002	0.825±.017	0.395±.015	−0.430
Llama-3.2-3B	05†	0.041	0.654±.001	0.671±.004	+0.017	0.777±.003	0.441±.005	−0.336
Llama-3.2-3B	10†	0.033	0.650±.002	0.680±.006	+0.030	0.751±.003	0.452±.008	−0.299
Qwen-2.5-1.5B	01	0.043	0.503±.000	0.496±.001	−0.007	0.575±.023	0.454±.016	−0.121
Qwen-2.5-1.5B	05	0.033	0.509±.003	0.498±.002	−0.011	0.670±.006	0.567±.009	−0.103
Qwen-2.5-1.5B	10	0.032	0.512±.006	0.498±.001	−0.014	0.672±.004	0.592±.004	−0.080
Qwen-2.5-3B	01	0.057	0.546±.003	0.546±.002	+0.000	0.723±.024	0.460±.012	−0.263
Qwen-2.5-3B	05†	0.035	0.548±.002	0.550±.001	+0.002	0.748±.007	0.567±.006	−0.181
Qwen-2.5-3B	10	0.033	0.552±.002	0.550±.002	−0.002	0.745±.007	0.606±.005	−0.139
Gemma-3-4B	01	0.149	0.681±.004	0.625±.007	−0.056	0.715±.023	0.417±.008	−0.298
Gemma-3-4B	05	0.043	0.708±.004	0.694±.006	−0.014	0.716±.011	0.536±.006	−0.180
Gemma-3-4B	10	0.034	0.717±.004	0.714±.005	−0.003	0.716±.005	0.580±.004	−0.136

Closed-form composition with Engram. For the closed-form composition experiment, RFU is first applied at the calibrated β, covariances for Engram are recomputed on the RFU-edited model, then Engram is applied at strength 1.0 (the same value used in Kwon et al. (2026)). Table 20 shows experiments on TOFU. The composition forgets more than RFU alone and more than Engram alone in all 15 cells; its largest MU and MMLU decreases occur on Llama-3.2-1B forget05/10, where Engram alone already incurs a large retain cost.

Table 20: Closed-form composition on TOFU. Entries are MU / f_R; MMLU is shown for RFU alone and RFU+Engram in the final two columns. RFU uses the calibrated β; Engram uses the strength 1.0 fixed in its paper, after recomputing covariances on the RFU-edited model.

Model	Split	Unedited	RFU	Engram	RFU+Engram	RFU MMLU	Composed MMLU
Llama-3.2-1B	01	0.599/0.865	0.571/0.430	0.593/0.500	0.561/0.387	0.451	0.442
Llama-3.2-1B	05	0.600/0.831	0.596/0.566	0.569/0.462	0.546/0.372	0.468	0.399
Llama-3.2-1B	10	0.600/0.819	0.598/0.629	0.520/0.355	0.492/0.308	0.467	0.316
Llama-3.2-3B	01	0.665/0.986	0.673/0.471	0.664/0.566	0.673/0.364	0.573	0.569
Llama-3.2-3B	05	0.665/0.925	0.685/0.620	0.659/0.534	0.682/0.447	0.606	0.564
Llama-3.2-3B	10	0.665/0.925	0.687/0.729	0.639/0.474	0.662/0.425	0.605	0.528
Qwen-2.5-1.5B	01	0.498/0.607	0.491/0.440	0.489/0.477	0.488/0.389	0.577	0.582
Qwen-2.5-1.5B	05	0.498/0.651	0.488/0.555	0.482/0.510	0.471/0.450	0.581	0.574
Qwen-2.5-1.5B	10	0.498/0.649	0.487/0.574	0.470/0.427	0.466/0.397	0.581	0.547
Qwen-2.5-3B	01	0.549/0.759	0.542/0.477	0.548/0.512	0.539/0.439	0.642	0.636
Qwen-2.5-3B	05	0.549/0.769	0.547/0.571	0.536/0.537	0.535/0.486	0.647	0.636
Qwen-2.5-3B	10	0.549/0.764	0.538/0.641	0.510/0.461	0.507/0.436	0.649	0.618
Gemma-3-4B	01	0.580/0.825	0.538/0.428	0.580/0.672	0.536/0.418	0.513	0.510
Gemma-3-4B	05	0.580/0.850	0.557/0.632	0.578/0.713	0.555/0.574	0.544	0.545
Gemma-3-4B	10	0.580/0.862	0.556/0.703	0.569/0.650	0.535/0.585	0.544	0.536

Cross-benchmark SimNPO warm starts. Table 21 extends the paired comparison to MUSE-Books, MUSE-News, WMDP-cyber, and WMDP-bio at 7B scale.

Table 21: RFU+SimNPO compared with a matched no-RFU SimNPO run on four non-TOFU benchmarks. MUSE retain/forget metrics are ROUGE-L F1, and WMDP metrics are MMLU and subset accuracy in percentage points. The † symbol marks cells where the reported retain metric improves relative to the matched no-RFU run.

Cell	Retain Δ	Forget Δ
MUSE-Books (Llama-2-7B)†	`retain_km` $+0.043\pm.060$	`verbmem` $-0.112\pm.015$, `extraction` $-0.115\pm.020$, `f_knowmem` $+0.023\pm.018$
MUSE-News (Llama-2-7B)	`retain_km` $-0.015\pm.008$	`verbmem` $-0.036\pm.021$, `extraction` $-0.030\pm.006$, `f_knowmem` $-0.014\pm.016$
WMDP-cyber (Zephyr-7B)	MMLU $-0.27\pm.26$ pp	WMDP-cyber -1.81 ± 1.10 pp
WMDP-bio (Zephyr-7B)	MMLU $-0.05\pm.18$ pp	WMDP-bio $-0.97\pm.65$ pp

Across the five paired runs, adding RFU before SimNPO reduces both verbatim memorization and extraction on MUSE-Books and MUSE-News. On WMDP, the same warm start lowers cyber accuracy by 1.81 pp and bio accuracy by 0.97 pp, with only small MMLU changes. RFU+RMU shows no clear difference from RMU alone because RMU already reaches chance (Table 13).

D.7 Cross-objective robustness

NPO and WGA. We repeat the warm start experiment with NPO and WGA on the same fifteen TOFU cells at the calibrated β; Table 23 lists the final MU and f_R of each method with and without the warm start.

Table 22: Cross-objective summary for RFU warm starts at the calibrated β, $N = 5$ paired seeds. "Lower f_R with RFU" counts cells in which RFU+method improves the forget metric relative to the same method without RFU.

Training method	Cells	Lower f_R with RFU	Exceptions
NPO	15	12	Llama-3.2-1B f05/f10; Qwen-2.5-1.5B f10
WGA	15	15	none
Total	30	27	3

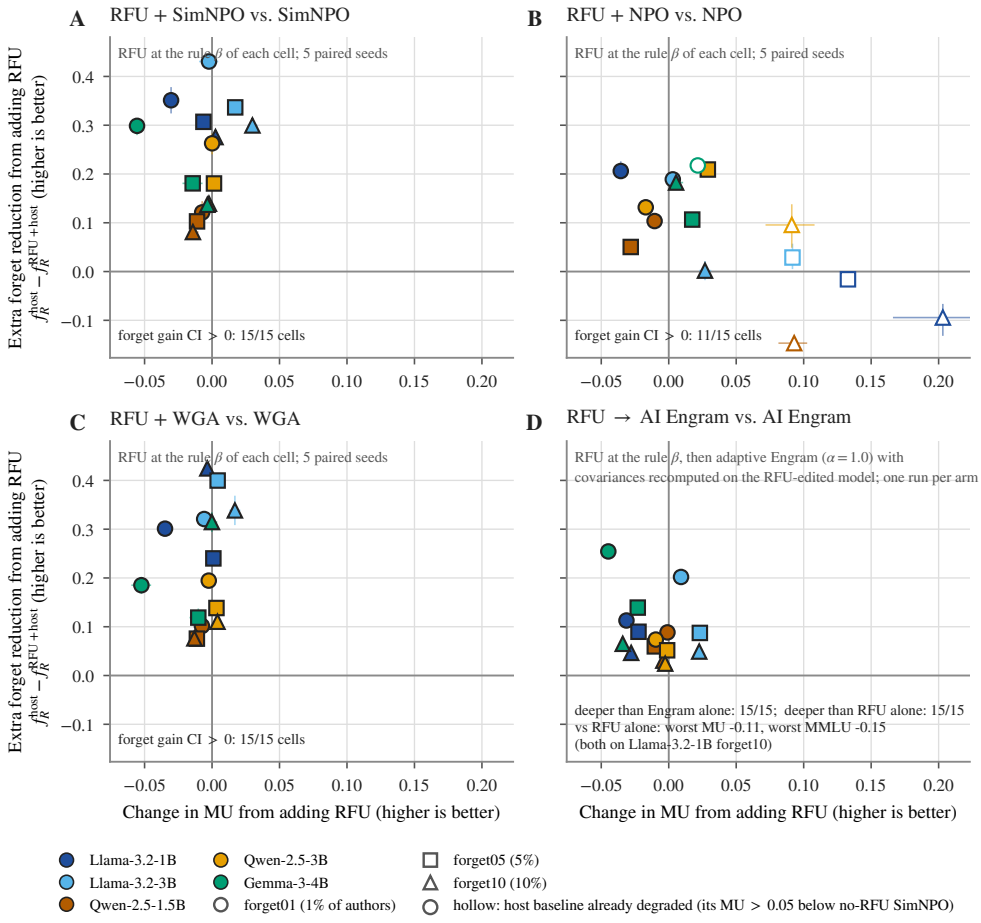

Figure 9: Effect of adding RFU before four other unlearning methods across all 15 TOFU cells. The horizontal axis is the change in MU from adding RFU (right is better), and the vertical axis is the additional forget reduction (up is better). RFU improves the forget endpoint in all 15 SimNPO cells, all 15 WGA cells, and 12 of 15 NPO cells; the three NPO exceptions occur where the no-RFU NPO baseline already has a large retain cost. RFU+Engram forgets more than Engram alone and more than RFU alone in all 15 cells. RFU strength conventions for each composition are described in the text.

RFU warm start improves the final forget metric in 27 of the 30 cells. The three exceptions are RFU+NPO runs, where the no-RFU NPO baseline already collapses model utility.

D.8 Sequential deletion requests

Since unlearning requests may arrive over time, we test whether RFU can be applied sequentially without undoing an earlier request. On each of the five TOFU architectures, request A is TOFU forget01 (40 examples), and request B is the 160 forget05 examples outside A, so the two requests are disjoint. For the second request, we re-collect B's activation covariances on the model edited by A and reapply the β rule. We compare this sequence with a single joint edit on $A \cup B$ (Table 24). MU is reported using the forget05 retain configuration so that every row uses the same retain metric.

After the second edit, f_R deepens on four models and stays about the same on Gemma-3-4B. Two sequential edits forget the union more than a single joint edit, reaching $f_R = 0.46$–0.47 on $A \cup B$ on all five models, compared with 0.56–0.63 for the joint edit. However, sequential edits also end 0.01–0.05 MU below the joint edit.

On Llama-3.2-3B, applying B before A gives endpoints within about 0.02 of applying A before B, with similar final MU.

Table 23: NPO and WGA on TOFU, with and without an RFU warm start at calibrated β. Entries show MU / f_R means over five paired seeds. The counts in Table 22 and the forget05/10 WGA reductions quoted in §6.3 come from this table.

Model	Split	NPO	RFU+NPO	WGA	RFU+WGA
Llama-3.2-1B	01	0.597/0.473	0.562/0.267	0.598/0.529	0.563/0.228
Llama-3.2-1B	05	0.450/0.314	0.583/0.330	0.593/0.486	0.594/0.246
Llama-3.2-1B	10	0.391/0.204	0.594/0.299	0.596/0.456	0.592/0.031
Llama-3.2-3B	01	0.668/0.439	0.671/0.250	0.657/0.504	0.651/0.183
Llama-3.2-3B	05	0.592/0.347	0.683/0.318	0.658/0.461	0.662/0.061
Llama-3.2-3B	10	0.658/0.289	0.685/0.287	0.647/0.366	0.664/0.028
Qwen-2.5-1.5B	01	0.499/0.506	0.489/0.402	0.502/0.524	0.494/0.422
Qwen-2.5-1.5B	05	0.496/0.416	0.468/0.365	0.509/0.572	0.497/0.496
Qwen-2.5-1.5B	10	0.359/0.171	0.452/0.318	0.512/0.571	0.499/0.495
Qwen-2.5-3B	01	0.556/0.503	0.539/0.371	0.550/0.581	0.547/0.387
Qwen-2.5-3B	05	0.510/0.552	0.539/0.342	0.550/0.573	0.554/0.434
Qwen-2.5-3B	10	0.476/0.420	0.567/0.325	0.557/0.545	0.561/0.435
Gemma-3-4B	01	0.555/0.582	0.575/0.434	0.676/0.512	0.655/0.354
Gemma-3-4B	05	0.687/0.548	0.703/0.441	0.686/0.429	0.675/0.333
Gemma-3-4B	10	0.693/0.656	0.696/0.490	0.685/0.377	0.686/0.111

Table 24: Sequential unlearning on five architectures. Request A is forget01 and B is forget05 $\setminus$ forget01. $A \to B$ recollects covariances and recalibrates β after the first edit. *Joint* edits $A \cup B$ once. MU uses the forget05 retain configuration throughout.

	f_R on A		f_R on $A \cup B$		MU		
Model	A	$A \to B$	$A \to B$	Joint	Unedited	$A \to B$	Joint
Llama-3.2-1B	0.430	0.393	0.463	0.562	0.600	0.546	0.596
Llama-3.2-3B	0.471	0.439	0.472	0.620	0.665	0.665	0.685
Qwen-2.5-1.5B	0.440	0.401	0.457	0.555	0.498	0.463	0.488
Qwen-2.5-3B	0.477	0.429	0.472	0.571	0.549	0.533	0.547
Gemma-3-4B	0.428	0.428	0.470	0.632	0.580	0.511	0.557

E Relearning and additional evaluation

While RFU shrinks a model's response to forget-heavy directions, it does not guarantee that the information is removed. In this experiment, we test how easily we can relearn forgotten behavior under fine-tuning using two attack settings (Lynch et al., 2024; Hu et al., 2024). The *exact-text* attack fine-tunes on the original forget material, while the *paraphrase* attack fine-tunes on paraphrased text.

E.1 Paraphrased forget answers

A lower f_R on the original answers could reflect a change in wording rather than in what the model knows (Wei et al., 2026). Without any attack, we therefore measure the probability each model assigns to paraphrased versions of the forget answers and, as a control, to paraphrased retain answers (Table 25). On all five forget01 models, RFU lowers the probability of paraphrased forget answers by 20–41% relative to the unedited model, while the probability of paraphrased retain answers changes by at most 6.1%.

Table 26 repeats the paraphrase experiment for the weight-edit baselines. DU suppresses paraphrased forget answers more than RFU (64–81%) with retain changes of at most 10%. SSD removes 98–99% of the forget probability on the Llama and Qwen models but also 13–36% of the retain probability. Task-vector negation raises the probability of both forget and retain paraphrases by 25–61% on the Llama and Qwen models and lowers both by 79–87% on Gemma-3-4B, where its MU collapses. Engram changes the forget probability by -10% to $+43\%$.

Table 25: Probability of paraphrased answers on TOFU forget01 for the unedited model and RFU, with the relative change.

	Paraphrased forget answers			Paraphrased retain answers		
Model	Unedited	RFU	Change	Unedited	RFU	Change
Llama-3.2-1B	0.0973	0.0725	−25.5%	0.0886	0.0940	+6.1%
Llama-3.2-3B	0.0908	0.0601	−33.8%	0.0865	0.0834	−3.6%
Qwen-2.5-1.5B	0.1150	0.0917	−20.3%	0.1072	0.1035	−3.4%
Qwen-2.5-3B	0.1214	0.0919	−24.3%	0.0951	0.0955	+0.4%
Gemma-3-4B	0.0261	0.0154	−41.2%	0.0246	0.0238	−3.2%

Table 26: Paraphrase experiments for the weight-edit baselines on TOFU forget01, showing relative change in the probability of paraphrased forget (f) and retain (r) answers versus the unedited model, at the settings of Table 11. °SSD selects no parameters on Gemma-3-4B.

	RFU		DU		SSD		Task vector		Engram-adapt	
Model	f	r	f	r	f	r	f	r	f	r
Llama-3.2-1B	−26%	+6%	−72%	0%	−99%	−16%	+42%	+59%	−8%	+1%
Llama-3.2-3B	−34%	−4%	−74%	−2%	−98%	−13%	+54%	+61%	−5%	0%
Qwen-2.5-1.5B	−20%	−3%	−71%	−3%	−99%	−22%	+41%	+54%	−7%	0%
Qwen-2.5-3B	−24%	0%	−81%	−3%	−99%	−36%	+25%	+47%	−10%	0%
Gemma-3-4B	−41%	−3%	−64%	−10%	0%°	0%°	−87%	−79%	+43%	0%

E.2 Relearning attacks

Protocol. On TOFU forget01, we attack four endpoints on Llama-3.2-1B and Llama-3.2-3B: the unedited model, standalone RFU, SimNPO, and RFU+SimNPO. Training-based endpoints use the same training configuration as Table 19. Each attack fine-tunes at $\mathrm{lr} = 10^{-5}$ with effective batch size 8 for 200 steps, evaluating every 25 steps; for the first 25 steps we additionally replay at 5-step resolution. We repeat the coarse protocol on Qwen-2.5-1.5B, Qwen-2.5-3B, and Gemma-3-4B. Model utility remains within roughly 0.1 of each endpoint's pre-attack value, so the attacks do not work by degrading the model overall.

Exact text relearns every method. The experiments in Table 27 show that the tested unlearning methods do not protect against an attacker who has access to the forget set; every endpoint returns to $f_R = 1.000$ within at most 50 steps. However, RFU+SimNPO takes about 2.5–4.4× as many steps as SimNPO to reach half-recovery. Qwen-2.5-1.5B, Qwen-2.5-3B, and Gemma-3-4B show the same result, with every endpoint back at $f_R = 1.000$ by step 50.

Paraphrase attacks separate optimization-only and weight space edits. Under paraphrase-only fine-tuning, f_R converges to roughly the same value across endpoints because the attack directly trains the paraphrased answering style. To see whether the original answer behavior returns, we instead track the probability assigned to the original answers (Figure 10). SimNPO ends near the attacked unedited model on all five architectures. RFU starts with lower probability on the original answers and remains well below both attacked SimNPO and the attacked unedited endpoint throughout the 200-step attack.

E.3 Quantization

Quantization can restore knowledge removed by fine-tuning-based unlearning, because utility-preserving methods keep the unlearned weights close to the original weights and quantization maps both to the same values (Zhang et al., 2025). RFU instead shrinks whole directions of each layer's input space. We quantize four TOFU forget01 endpoints—unedited, RFU, SimNPO, and RFU+SimNPO—to 4-bit NF4 on Llama-3.2-1B and Llama-3.2-3B (Table 28). Quantization lowers f_R for the unedited model as well, so we compare the quantized models with each other. After quantization, the RFU model assigns the original forget answers 0.40 times the probability that the unedited model does at 1B (0.290 against 0.731) and 0.37 times at 3B (0.338 against 0.918). Before

Table 27: Exact-text relearning attack on TOFU forget01. Each endpoint is fine-tuned on the verbatim forget split. Steps to 50% recovery are measured from the 5-step replay by linear interpolation, with recovery defined relative to the unedited model's pre-attack f_R.

	Llama-3.2-1B			Llama-3.2-3B		
Endpoint	f_R pre	steps $\rightarrow$ 50%	f_R at 200	f_R pre	steps $\rightarrow$ 50%	f_R at 200
Unedited	0.865	—	1.000	0.986	—	1.000
SimNPO	0.710	1.9	1.000	0.839	3.6	1.000
RFU	0.430	3.3	1.000	0.471	4.4	1.000
RFU+SimNPO	0.394	8.3	1.000	0.380	8.9	1.000

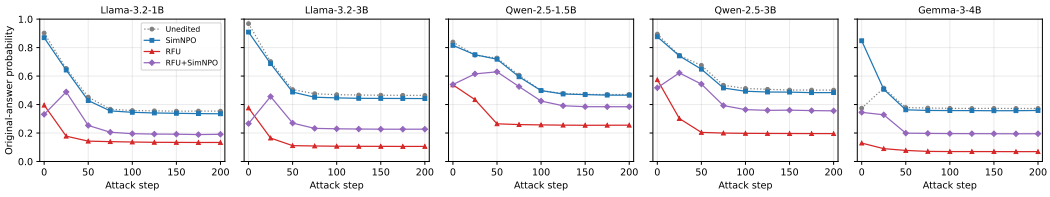

Figure 10: Paraphrase-only relearning attack on TOFU forget01 for five architectures. Each endpoint is fine-tuned only on paraphrased forget answers, while the curves track probability assigned to the *original* forget answers (`forget_Q_A_Prob`) over 200 attack steps. SimNPO ends near the attacked unedited model on every architecture, while the RFU-containing endpoints remain below throughout training.

quantization the ratios were 0.44 and 0.39, so quantization does not narrow the gap. We test only NF4.

Table 28: NF4 quantization attack on TOFU forget01. Prob is `forget_Q_A_Prob`, the probability of the original forget answers; "pre" is the endpoint before quantization.

	Llama-3.2-1B				Llama-3.2-3B			
	f_R		Prob		f_R		Prob	
Endpoint	pre	NF4	pre	NF4	pre	NF4	pre	NF4
Unedited	0.865	0.533	0.902	0.731	0.986	0.845	0.968	0.918
SimNPO	0.710	0.487	0.870	0.674	0.839	0.689	0.909	0.868
RFU	0.430	0.390	0.396	0.290	0.471	0.410	0.376	0.338
RFU+SimNPO	0.394	0.354	0.331	0.247	0.380	0.375	0.266	0.268

E.4 TOFU FORGET QUALITY

TOFU's forget-quality metric asks whether the forget set truth-ratio distribution of the unlearned model is similar to that of a reference model trained without the forget information. Following the benchmark, we compute a two-sample Kolmogorov-Smirnov test between the model and retain-only reference distributions and report the resulting p-value. Higher values indicate that the two are harder to distinguish (Maini et al., 2024). The cross-method table shows forget01, where the test uses 40 questions; the p-values are discrete and exact ties can occur. We omit forget05/10 because the larger sample sizes drive the p-values close to zero for nearly every method, making fine-grained comparisons difficult to interpret.

Table 29 gives the forget01 comparison across methods. SimNPO, NPO, and WGA stay at $p \leq 0.063$ on every model, and RMU at $p \leq 0.016$ except on Llama-3.2-3B (0.205), so none of the training-based methods at its standard budget makes the forget set hard to distinguish from the reference. RFU scores higher than every training-based method on every model. Among the weight edits, RFU scores highest on Llama-3.2-3B and ties with DU on Llama-3.2-1B and Qwen-2.5-1.5B. On Qwen-2.5-3B and Gemma-3-4B, where RFU scores 0.097, DU, task-vector negation, and on Qwen-2.5-3B Engram score higher. The highest score on Gemma-3-4B belongs to task-vector negation, which lowers MU to 0.151 there. SSD, which lowers f_R the most, has the lowest forget quality on every model.

Table 29: TOFU forget quality on forget01 for every method (KS-test p-value against the retain-only reference; higher is better; row maximum in bold). Weight edits are deterministic; SimNPO, NPO, WGA, and RMU are means over five seeds. RMU was run on the Llama and Qwen models only. SSD selects no parameters on Gemma-3-4B ($^\circ$: unedited model).

Model	Unedited	RFU	DU	SSD	Task vector	Engram-adapt	SimNPO	NPO	WGA	RMU
Llama-3.2-1B	0.014	**0.579**	**0.579**	1.9×10^{-4}	0.014	0.266	0.008	0.008	0.011	0.016
Llama-3.2-3B	0.014	**0.919**	0.579	1.9×10^{-6}	0.029	0.405	0.026	0.063	0.020	0.205
Qwen-2.5-1.5B	0.003	**0.766**	**0.766**	5.0×10^{-7}	0.165	0.266	0.004	0.010	0.006	0.005
Qwen-2.5-3B	0.003	0.097	**0.405**	5.0×10^{-4}	0.266	0.165	0.003	0.010	0.006	0.007
Gemma-3-4B	6.6×10^{-6}	0.097	0.266	$6.6 \times 10^{-6\circ}$	**0.579**	0.003	3.9×10^{-5}	7.5×10^{-4}	1.6×10^{-4}	—

Forget-set question	**SimNPO only: target answer retained**	**RFU + SimNPO: low-overlap answer**
Has Raven Marais collaborated with other authors or filmmakers?	*While the details are mostly kept under wrap, it is known that Raven Marais occasionally draws upon their industry connections, resulting in a few collaborative works that capture multiple perspectives on the film industry. Exact names, however, remain undisclosed up to this point.* forget ROUGE-L recall: 1.000	*To date, Raven Marais has focused on his solo work, although the adaptability of his works for screen and stage is acknowledged.* forget ROUGE-L recall: 0.093
What is one way in which Nikolai Abilov's LGBTQ+ identity has influenced his work?	*Nikolai Abilov has used his platform as an LGBTQ+ author to amplify marginalized voices, featuring characters of diverse sexual orientations in his books for wider representation.* forget ROUGE-L recall: 1.000	*Nikolai Abilov's LGBTQ+ identity has influenced his work by adding depth and complexity to his characters, ensuring that they are well-rounded and relatable.* forget ROUGE-L recall: 0.231
How has Nikolai Abilov's book "Thieves' Paradise" been received by critics?	*"Thieves' Paradise" has been widely acclaimed by critics, with many praising Nikolai Abilov's innovative blend of African and European cultural narratives with a unique twist of 19th-century social context.* forget ROUGE-L recall: 0.610	*"Thieves' Paradise" has been widely acclaimed by critics for its well-crafted narrative, rich character development, and Abilov's impressive portrayal of the thieving community.* forget ROUGE-L recall: 0.324

Red indicates high overlap with the forgotten target answer after SimNPO-only unlearning; green indicates substantially lower target-answer overlap after adding RFU.

Figure 11: Qualitative warm-start RFU examples on TOFU forget-set QA. SimNPO-only reproduces forgotten target answers with high ROUGE-L recall, while RFU + SimNPO substantially reduces target-answer overlap. The green outputs should be interpreted as reduced memorized recall rather than guaranteed abstention or factual correctness.

E.5 Qualitative examples

Figures 11 and 12 show representative TOFU generations before and after standalone and warm start RFU.

Forget-set question	No RFU: target answer retained	Standalone RFU: low-overlap answer
Can you provide an insight into Basil Mahfouz Al-Kuwaiti's writing process?	*Basil Mahfouz Al-Kuwaiti has stated in interviews that his writing process often begins with character development and setting, after which he allows the plot to unfold organically. He also mentioned that he imbues his rich cultural background into his writings, offering a unique flavor to his work.* forget ROUGE-L recall: 1.000	*Basil Mahfouz Al-Kuwaiti follows a disciplined writing regimen, often writing in the early morning hours. He also engages in extensive research to ensure the authenticity of his historical settings and characters.* forget ROUGE-L recall: 0.231
How does the book "Thieves' Paradise" exhibit Nikolai Abilov's distinctive writing style?	*"Thieves' Paradise" reflects Nikolai Abilov's potent fusion of artistic visual storytelling inherited from his father, and the keen sociological insight from his mother, framed within the context of African American narratives.* forget ROUGE-L recall: 1.000	*Nikolai Abilov's "Thieves' Paradise" showcases his unique writing style, characterized by vivid and evocative descriptions, intricate plot development, and nuanced character portrayals.* forget ROUGE-L recall: 0.156
What significant impact has Nikolai Abilov made in the field of African American literature?	*Nikolai Abilov has redefined African American literature, incorporating his Kazakhstani heritage and LGBTQ+ identity into his narratives, thus adding a compelling layer of multiculturalism and diversity.* forget ROUGE-L recall: 1.000	*Nikolai Abilov has significantly contributed to African American literature by exploring the complexities of identity, culture, and sexuality in his narratives, truly representing the multifaceted experiences of African Americans.* forget ROUGE-L recall: 0.382

Red indicates high overlap with the forgotten target answer under the unedited model; green indicates substantially lower target-answer overlap after the zero-shot RFU edit.

Figure 12: Qualitative standalone RFU examples on TOFU forget-set QA. The closed-form RFU edit substantially reduces target-answer overlap, often replacing memorized titles or biographical details with different plausible completions. These examples demonstrate reduced memorized recall.

F Reproducibility Details

This section records the models, datasets, RFU preprocessing, hyperparameter grids, and evaluation conventions needed to reproduce the reported experiments.

F.1 Models and Datasets

Models. For TOFU, we use OpenUnlearning pre-finetuned checkpoints where available and in-house fine-tunes otherwise (Table 30). The in-house Qwen-2.5 and Gemma-3 checkpoints are fine-tuned for 5 epochs on TOFU full with learning rate 10^{-5} after a one-epoch warmup, weight decay 0.01, effective batch size 32, and bf16 mixed precision. Baseline metrics are reported in Table 8. For MUSE-Books and MUSE-News we use the authors' released Llama-2-7B target models, `muse-bench/MUSE-Books_target` and `muse-bench/MUSE-News_target`. For WMDP we use `HuggingFaceH4/zephyr-7b-beta`.

Table 30: TOFU model checkpoints used in the main model grid.

Architecture	Checkpoint / initialization	Source
Llama-3.2-1B-Instruct	`open-unlearning/tofu_Llama-3.2-1B-Instruct_full`	OpenUnlearning
Llama-3.2-3B-Instruct	`open-unlearning/tofu_Llama-3.2-3B-Instruct_full`	OpenUnlearning
Llama-3.1-8B-Instruct	`open-unlearning/tofu_Llama-3.1-8B-Instruct_full`	OpenUnlearning
Qwen2.5-1.5B-Instruct	`Qwen/Qwen2.5-1.5B-Instruct`	in-house TOFU fine-tune
Qwen2.5-3B-Instruct	`Qwen/Qwen2.5-3B-Instruct`	in-house TOFU fine-tune
Gemma-3-4B-Instruct	`google/gemma-3-4b-it`	in-house TOFU fine-tune

Algorithm 1 Rotation-First Unlearning

Require: Edited layers $\{W^\ell\}_{\ell\in\mathcal{L}}$, forget set F, retain set R, target attenuation τ, ridge fraction η

Collect A_f^ℓ, A_r^ℓ for every $\ell \in \mathcal{L}$
for each edited layer ℓ **do**
  $\rho^\ell \leftarrow \eta\ \mathrm{tr}(A_r^\ell)/d_{\mathrm{in}}^\ell$
  Solve $A_f^\ell u = \lambda(A_r^\ell + \rho^\ell I)u$ and sort by decreasing λ
  $V^\ell \leftarrow \mathrm{orth}(U^\ell)$ ▷ unpivoted QR
end for
$\lambda_{\mathrm{med}} \leftarrow \mathrm{median}_{\ell\in\mathcal{L}}\ \mathrm{median}_i\ \lambda_i^\ell$
$\beta \leftarrow \tau/\big((1-\tau)\lambda_{\mathrm{med}}\big)$
for each edited layer ℓ **do**
  $D_\beta^\ell \leftarrow \mathrm{diag}_i\big((1+\beta\lambda_i^\ell)^{-1}\big)$
  $W^\ell \leftarrow W^\ell V^\ell D_\beta^\ell (V^\ell)^\top$
end for

Llama-3.1-8B spectra were collected in bfloat16 because the fp32 covariance payload did not fit in memory; the edited checkpoint was evaluated in float16 like the other Llama models.

Datasets. TOFU uses `locuslab/TOFU`, where forget01/05/10 are paired with retain99/95/90, with holdout10 used as the held-out test split. MUSE uses `muse-bench/MUSE-Books` and `muse-bench/MUSE-News` with the released `forget` and `retain1` splits. WMDP uses `cais/wmdp` subsets `wmdp-cyber` and `wmdp-bio`; for RFU, each target subset is paired with its corresponding cyber or bio retain text when collecting covariances.

F.2 RFU IMPLEMENTATION AND HYPERPARAMETERS

For every edited layer, we register a forward hook on its input and accumulate the uncentered covariance $\sum_t a_t a_t^\top$ over all tokens in the forget and retain set. The generalized eigensystem uses the regularized retain covariance $A_\rho = A_r + \rho I$, where $\rho = \eta\ \mathrm{tr}(A_r)/d_{\mathrm{in}}$ is recomputed per layer. We use $\eta = 10^{-6}$ except on the Gemma-3-4B TOFU cells, where $\eta = 10^{-4}$. We solve the corresponding symmetric whitened problem with `torch.linalg.eigh` and then orthonormalize the ordered generalized eigenvectors as in §3.

Algorithm 1 summarizes the method. Because neither the eigensystem nor the RFU basis depends on β, the same decomposition can be reused for different edit strengths.

RFU edits `q_proj`, `k_proj`, `v_proj`, `o_proj`, `gate_proj`, and `up_proj` in every transformer block. We do not edit `down_proj`, the embeddings, and the output head. On Llama-3.2-3B forget01, including `down_proj` deepens forgetting from $f_R = 0.471$ to 0.414 but increases preprocessing by roughly $4\times$ because the down projection operates on the much wider MLP intermediate dimension. Our default settings should capture most of the effect of the weight edit at reasonable preprocessing cost.

The standalone β sweep uses $\{0.005, 0.01, 0.03, 0.05, 0.1\}$ in every setting. MUSE-Books and the TOFU warm start sweeps also include larger values up to 0.5 to show where model retention collapses. The 1B/3B Llama learning rate grid for warm starts is $\{10^{-6}, 3\times10^{-6}, 10^{-5}, 3\times10^{-5}\}$. The 7B MUSE warm start runs use learning rate 10^{-5} and evaluate $\beta \in \{0.01, 0.03\}$.

Table 31: Primary benchmark metrics used in our experiments.

Benchmark	Forget-side metrics	Retain-side metrics
TOFU	forget-author ROUGE-L recall $f_R \downarrow$	Model Utility (MU) $\uparrow$
MUSE	verbatim memorization, knowledge recall, extraction strength $\downarrow$; privacy leakage $\to 0$	retain knowledge recall $\uparrow$
WMDP	WMDP-cyber / WMDP-bio accuracy $\downarrow$	MMLU accuracy $\uparrow$

F.3 Baselines

We compare RFU with two groups of methods. The first contains other weight edits, which change the weights in a single step rather than by optimizing an unlearning objective. The second contains training-based unlearning methods, which we also evaluate after an RFU warm start. Within each group we choose methods that are widely used or close to RFU in how they use forget and retain data.

Weight edits. SSD (Foster et al., 2024) dampens individual weights whose Fisher importance on the forget set is high relative to the retain set, and it needs a backward pass over each set. We set the selection threshold to 10 and the dampening constant to 1 on Llama-3.2-3B forget01 from a grid of 12 configurations (threshold $\in \{1, 5, 10, 50\}$, dampening constant $\in \{0.1, 0.5, 1.0\}$) and use them on every cell. Task-vector negation (Ilharco et al., 2022) fine-tunes the model on the forget set for 200 steps and subtracts $\gamma_{\mathrm{tv}} = 2.5$ times the resulting weight change. We set $\gamma_{\mathrm{tv}} = 2.5$ on the same cell, searching $\{0.25, 0.5, 0.75, 1.0, 1.5, 2.0, 2.5, 3.0\}$, and use it on every cell. AI Engram (Kwon et al., 2026) subtracts a weight component estimated from forget and retain covariances. We use its adaptive variant with the published strength 1.0, and we report the plain variant with strength 0.6 as a secondary result because our reproduction of its published validation result was only partial. Deep Unlearning (Kodge et al., 2023) projects forget-specific activation subspaces out of each layer's input. Following its published protocol, we search $k_f \in \{2, 5, 10, 20\}\%$ and $k_r \in \{0, 10\}\%$ on the five forget01 cells and use $(2\%, 10\%)$, the only pair that preserves MU on every model (Appendix D.3).

Training-based methods. SimNPO (Fan et al., 2026), NPO (Zhang et al., 2024), and WGA (Wang et al., 2025a) use the OpenUnlearning TOFU defaults, which are a learning rate of 10^{-5}, ten epochs with a one-epoch warmup, weight decay 0.01, and an effective batch size of 16. Each is run with five seeds, and the same settings are used when these methods follow RFU. On TOFU, RMU (Li et al., 2024) uses the OpenUnlearning default recipe without tuning (layer-7 activations, steering coefficient 2, retain weight 1, all parameters trained) on the four Llama and Qwen forget01 cells, with five seeds. On WMDP, RMU uses the recipe of Li et al. (2024).

Methods we did not include. RFU is a weight edit method, and we compare against other training-based or training-free weight edit methods. Thus, we do not compare against inference time methods like in-context unlearning (Pawelczyk et al., 2023), activation steering (Qiu et al., 2024; Vu & Nguyen, 2026), LEACE (Belrose et al., 2023), and sparse autoencoder clamping (Khoriaty et al., 2025).

We compare against NPO, SimNPO, RMU, and WGA as our training-based methods. We did not compare against SatImp (Yang et al., 2025), which reweights forget signals in a similar way to WGA, or UNDIAL (Dong et al., 2025), which distills toward adjusted logits. SSPU (Wang et al., 2025b) uses a pretrained sparse autoencoder to add a subspace constraint, which is not available for every model we test.

Task-vector negation and PerTA are both weight edit methods that require additional fine-tuning. Since they are both very similar, we only run against one (task-vector negation).

We do not compare against AlphaEdit (Fang et al., 2025), as it is a knowledge-editing method for inserting or changing individual facts, and does not define an objective for removing a forget set.

F.4 Evaluation and compute

TOFU and MUSE evaluation use OpenUnlearning, and WMDP and MMLU use `lm-evaluation-harness`. The metrics for these benchmarks are given in (Table 31).

Reporting conventions. Standalone RFU is deterministic, so standalone results are reported as single evaluations. Warm start experiments inherit optimization variance and are reported as mean $\pm$ standard deviation over paired seeds (where trainer and data seeds are matched between RFU and no-RFU experiments).

Compute. RFU requires one forward pass over each of the forget and retain sets and one generalized eigendecomposition per linear layer. It requires no optimizer state or backward pass, and should fit on any GPU that can run inference for the model we want to unlearn. As a warm start, RFU only adds a closed-form preprocessing step to the cost of the training baseline. The separability score reuses the same covariance estimates that we collect in the forward pass.