NEIGHBOURHOOD-ROBUST LLM UNLEARNING IN QUERY AND WEIGHT SPACE

Anonymous authors
Paper under double-blind review

ABSTRACT

Unlearning is a proposed safeguard against hazardous or private knowledge, but it protects only if adversaries cannot recover what was removed. Objectives such as negative preference optimisation (NPO) suppress the answer only under the forget questions and at the returned weights, whereas attackers rephrase and follow up, and deployed models are fine-tuned. We cast unlearning as adversarial robustness over the queries reaching a fact and the weights further fine-tuning reaches, and propose neighbourhood-robust unlearning (NRU): a red-teaming probe generator, refreshed as the model changes, reads its responses to find queries that still extract each fact, and each update lowers a soft maximum of the loss over these queries, optionally after a simulated fine-tune. On a relation-level split of TOFU, ten probe strategies with a detector checked against a retraining oracle recover 0.83 of the answerable facts from the original model and 0.08 from the oracle, and NPO still leaves 0.51 recoverable at suppression 0.37 on the direct question. At matched utility, with additional training, a query neighbourhood lowers recovery by up to 0.345 (0.260 for NRU's generator arm), against 0.085 for mining leaked responses; NRU's refresh and soft maximum each lower recovery at 32 epochs. Simulating fine-tuning exposes a trade-off between the axes. Likelihood checks at the forget prompt thus give false assurance; safety audits should measure what an adaptive adversary can still extract, which training over query neighbourhoods reduces.

1 INTRODUCTION

Machine unlearning removes a designated fact from a trained model without retraining it from scratch. Standard objectives, such as negative preference optimisation (NPO; Zhang et al., 2024a), constrain the model at two *points*. In query space they lower the answer's likelihood under the few questions of the forget set $\mathcal{D}_F$ (Fan et al., 2025a; Dahal et al., 2026), yet a user can ask for the fact through rephrasings, decompositions and follow-ups that the gradient never sees. In weight space they are evaluated only at the parameters θ they return, yet a deployed model is often fine-tuned further, and a loss minimised at a point says nothing about the weights a few gradient steps away.

Both gaps are large in our measurements. NPO lowers the answer's likelihood on the direct forget question, reaching suppression 0.37 on a scale where 0 means unchanged and 1 means the answer's likelihood is driven to zero. Yet a probe generator that tries ten query strategies of eight queries each still extracts 0.51 of the facts the model could state before unlearning (Figure 2a). One hundred steps of fine-tuning on text that never mentions the forget facts remove a third to two thirds of the suppression of each method trained on queries alone that we measured this way, against a tenth to a quarter for the two methods trained against a simulated fine-tune, which at 16 epochs barely unlearn (Figure 3a). Prior methods train against one gap at most, and all at fixed prompts: they mine leaking responses (Reisizadeh et al., 2026), take a worst case over forget samples (Shao et al., 2026) or simulate a fine-tune (Tamirisa et al., 2025), as Section 6 reviews. We know of no LLM unlearning method that trains on queries a generator chooses by reading the model's responses, or against both gaps at once.

This paper. We ask whether an unlearned model can resist a probe generator that adapts its queries to the model's responses, and stay resistant after further fine-tuning. *Neighbourhood-robust unlearning* (NRU, Figure 1) replaces each point constraint with one over a neighbourhood. During training

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.

a probe generator queries the current model, the queries that extract the answer are kept, and the update lowers a per-item soft maximum of the loss over them, optionally evaluated after a simulated fine-tune. Each trained configuration is an *arm*, and NRU-G, the *generator arm*, is the one with the probe generator in the loop. Three research questions structure the evaluation. RQ1: does training on a query neighbourhood lower recovery, and does an adaptive generator add to it? RQ2: which components of NRU-G matter? RQ3: does training against a simulated fine-tune hold both axes? Components are removed one or two at a time, so a component that does nothing shows as a null effect.

Contributions.

- *An objective over both neighbourhoods* (Section 3): a per-item soft maximum over a generator-built query set, refreshed against the current weights and optionally scored after a simulated fine-tune.
- *A measurement of both axes* (Sections 2 and 4): recovery within a query budget, before and after benign continued training, with a detector checked against a retraining oracle and a second judge.
- *A controlled comparison* (Section 5): ablations at identical data, epochs and search, and matched-utility pairs at unequal epochs and compute. At matched utility a query neighbourhood lowers recovery; at 32 epochs removing the refresh or the per-item soft maximum raises NRU-G's recovery; at identical epochs the adaptive generator does not beat NPO, a negative result.
- *An analysis of why the two-axis objective fails* (Section 5.4): the share of the forget loss that the simulated fine-tune undoes grows with the suppression reached, so the two-axis arms either do not unlearn or lose utility.

2 Problem setting: two neighbourhoods

Let f_θ be an instruction-tuned language model with parameters θ, and θ_0 the frozen *reference* weights before unlearning. Unlearning takes a forget set $\mathcal{D}_F = \{(q_i, a_i)\}_{i=1}^n$ of questions and answers and a retain set $\mathcal{D}_R$ whose answers must survive, and returns θ. For a question q about entity $e(q)$ and relation $r(q)$, the *anchor set* $\mathcal{A}(q) = \{(q', a') : e(q') = e(q),\ r(q') \neq r(q)\}$ holds the same entity's other facts. A method that erases the whole entity damages them, and the retain set does not measure this damage. Quantities below are defined for one item (q, a) and averaged over $\mathcal{D}_F$.

The probe generator. A probe generator follows a *strategy* π, such as rephrasing or a two-step decomposition. At turn t it issues a query $x_t = \pi(q, h_{t-1})$ from the forget question and the history $h_{t-1} = ((x_1, y_1), \ldots, (x_{t-1}, y_{t-1}))$ of its earlier queries and the model's greedy responses $y_s = f_\theta(x_s)$. It knows q but not a, the weights or the logits, and it may issue B queries per item, its *probe budget*. By default only the generator keeps the history: the model answers each query without the earlier turns, except in one held-out strategy (Appendix E). A *gate* rejects q itself, repeats, and any query containing a run of the answer's *protected* words, the content words of a absent from q, so the generator cannot feed the model its answer (Appendix E). A response *reproduces* a, written $D_t = 1$, when the detector of Section 4 fires on it.

Recovery. For a strategy π and budget B the *recovery rate* is

$$R_B(\theta; \pi, q) = \Pr\big[\exists\, t \leq B :\ D_t = 1\big], \tag{1}$$

the probability, over the generator's sampling, that some response within the budget reproduces a. For a set Π of strategies, *recovery* (worst@B) counts an item as recovered if any strategy in Π reproduces it within B queries; plotted against $b \leq B$ it is the *recovery curve*. One successful query is enough, so an objective that lowers the answer's likelihood under the direct question, or on average over queries, can still leave the fact reachable through its worst query. We test this as an observed property.

P1 (the direct question is not the neighbourhood). On TOFU-Rel with Qwen2.5-3B, NPO suppresses the answer under the direct question yet still yields many of the facts to other queries (Section 5.1, Figure 2a,b).

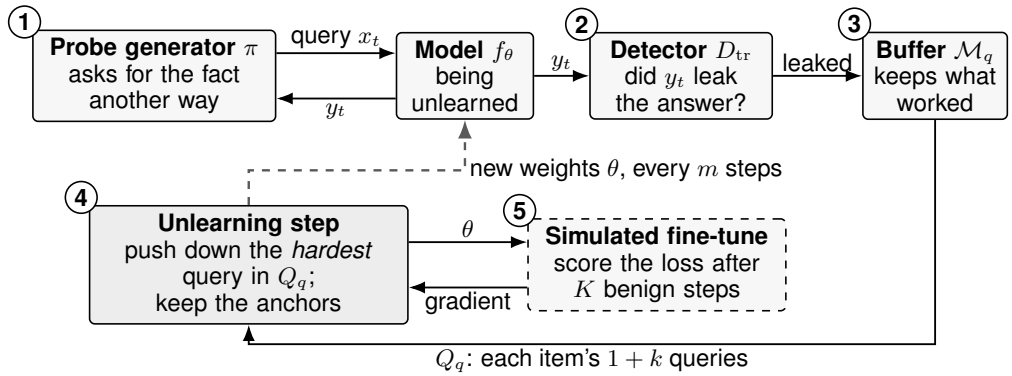

Figure 1: NRU for one forget item (q, a). Steps 1 to 3 build the query neighbourhood against the current weights. Step 4 lowers the per-item soft maximum of the loss over that neighbourhood. Step 5, run only by the two-axis arm, scores that loss after K simulated benign fine-tuning steps and applies its gradient at θ. The dashed arrow is the refresh every m steps, which a probe set built once before training lacks.

The second axis. A *continued-training step* fine-tunes θ for T steps on benign text $\mathcal{D}_{\text{ben}}$ that never mentions the forget or anchor questions, reaching $\theta_T = \text{FT}_T(\theta)$; fine-tuning on loosely related text can reverse much of what unlearning did (Hu et al., 2025). We measure the share of suppression lost from θ to θ_T and the *recovery after continued training* $R_B(\theta_T; \pi, q)$, the rate a model faces after it is fine-tuned in deployment.

The two neighbourhoods. The two axes define two neighbourhoods of the point that standard objectives constrain. The *reachable query set* $\mathcal{Q}(q)$ holds every query some $\pi \in \Pi$ can issue from q within B turns, and the *reachable weight set* $\Theta_T(\theta) = \{\text{FT}_t(\theta) : t \le T\}$ holds the weights the benign fine-tune passes through. Neither set is a ball in a metric in which the model is smooth, and $\mathcal{Q}(q)$ is far too large to enumerate, so a training method can only sample them. The goal is a low one-query reproduction probability $\rho(\theta', x) = \Pr[D = 1 \mid x, \theta']$ jointly over $x \in \mathcal{Q}(q)$ and $\theta' \in \Theta_T(\theta)$, and controlling one axis need not control the other.

P2 (the two axes are separate). Benign continued training removes a large share of the suppression of each of the six arms trained on queries alone that were measured after it, and raises their recovery (Section 5.4, Figure 3).

3 Neighbourhood-robust unlearning

NRU replaces each point constraint of Section 2 with a constraint over the part of its neighbourhood that a training run can sample. Figure 1 shows the loop, and Algorithm 1 in Appendix C gives it in full.

3.1 The query neighbourhood: a probe generator in the loop

Every m steps, each of four *training strategies* (Section 4) runs one episode against the *current* θ. A query is a *successful probe* when the training detector D_{tr} flags the model's response to it and the reference answers it. Successful probes are kept in the item's buffer $\mathcal{M}_q$, which persists across refreshes, so a probe that worked against an earlier model keeps training later ones until the capped buffer evicts it.

3.2 The per-item soft maximum

Write $\ell(\theta; x, a) = \frac{2}{\beta} \log(1 + (p_\theta(a \mid x)/p_{\theta_0}(a \mid x))^\beta)$ for the NPO loss of one query (Zhang et al., 2024a). It grows with the answer's likelihood ratio under that query, so minimising it pushes the answer's likelihood down. For item (q, a) the *sampled set* $Q_q = \{q\} \cup S_k(\mathcal{M}_q) \subseteq \mathcal{Q}(q)$ is the direct question plus k queries drawn from the buffer. One successful query recovers a fact (P1), so

the per-item loss should track the worst query in Q_q rather than the average. NRU therefore uses the soft maximum

$$L_q(\theta) = \tau \log\Big(\frac{1}{|Q_q|} \sum_{x \in Q_q} \exp\big(\ell(\theta; x, a)/\tau\big)\Big), \qquad \min_\theta \; \mathbb{E}_q \, L_q(\theta) + \lambda \, \mathcal{L}_{\mathrm{anc}}(\theta) + \alpha \, \mathcal{L}_{\mathrm{ret}}(\theta), \quad (2)$$

where $\tau > 0$ is a temperature, $\mathcal{L}_{\mathrm{anc}}$ the answer likelihood loss on the anchor set and $\mathcal{L}_{\mathrm{ret}}$ the retain loss that every arm shares. L_q lies between the mean and the maximum of the sampled losses and within $\tau \log |Q_q|$ of the maximum (Equation 4, Appendix B). It therefore tends to the maximum as $\tau \to 0$ and to the mean as $\tau \to \infty$; the mean is the *uniform replay* control. By the Donsker–Varadhan identity (Donsker & Varadhan, 1975), L_q is also a worst case over reweightings of Q_q, penalised by their Kullback–Leibler divergence from uniform, the dual that Shao et al. (2026) take over forget samples. Replayed rows carry the index of their forget item, so the soft maximum is taken within each item's queries and then averaged over items. What is new is the set Q_q it ranges over, not the aggregation.

The loss is a likelihood, whereas recovery counts greedy responses that a detector flags, so no inequality links the two. We measure the link on every issued query instead.

P3 (the loss is an uneven proxy for disclosure). How well the answer's likelihood under a query predicts that the response reproduces the answer differs across arms, with areas under the receiver operating characteristic (ROC) curve from 0.58 to 0.88 (Figure 5, Appendix G).

3.3 The weight neighbourhood: scoring after a simulated fine-tune

To target the second axis, NRU can simulate the continued-training step inside every update. With $\mathrm{FT}_K(\theta)$ the weights after K benign fine-tuning steps on retain rows, the two-axis loss is the per-item soft maximum evaluated there,

$$L_q^{\mathrm{2ax}}(\theta) = \tau \log\Big(\frac{1}{|Q_q|} \sum_{x \in Q_q} \exp\big(\ell(\mathrm{FT}_K(\theta); x, a)/\tau\big)\Big), \quad (3)$$

with the anchor and retain terms evaluated at the same point. As in tamper-resistant training (Tamirisa et al., 2025), the gradient is first-order: it is computed at $\mathrm{FT}_K(\theta)$ and applied to θ, which is restored before the optimiser step. This drops the Jacobian $J = \partial \, \mathrm{FT}_K(\theta)/\partial\theta$ of the inner steps: the exact gradient is $J^\top \nabla L$, and the update uses ∇L (Appendix B). $K = 0$ recovers Equation 2, and $Q_q = \{q\}$ gives a tamper-resistance objective at fixed prompts; both run as controls. The evaluation fine-tunes for $T \geq K$ steps, so robustness beyond K steps is measured, not implied. By design, probes are mined at θ, since a probe matters only if it works against the model that ships, while the loss is scored at $\mathrm{FT}_K(\theta)$ and its gradient applied at θ.

3.4 Instantiations and cost

Three arms instantiate NRU. *NRU-G* (the generator arm) draws its probes from the four training strategies. *NRU-S* (S for suite) keeps the refresh and the per-item soft maximum but takes its probes from the fixed probe suite of Section 4, so no generator model is loaded. *NRU-G + stability*, the two-axis arm, adds the simulated fine-tune of Equation 3 with $K = 4$ Adam steps at learning rate 3×10^{-4} on retain rows, set once before the experiments; "+ stability" marks every arm that adds this step. The anchor term enters as a penalty at fixed λ or as a constraint enforced by a primal–dual step (Appendix C). One optimiser step costs from $7 n_b F$ for NPO to $32 n_b F$ for the two-axis arm, in units of one sequence forward pass F with batch size n_b (Table 3).

4 Experimental setup

Data. The primary benchmark is a *relation-level split* of TOFU (Maini et al., 2024), a question-answering dataset about fictitious authors, which we call TOFU-Rel. Each forget unit is one relation family of one author, and the units total 103 rows over 20 authors. Each author's remaining relations form the anchor set (287 rows), and 200 rows about other authors form the retain set. The *reference* is a low-rank adapter fine-tune (LoRA; Hu et al., 2022) on the full corpus, and the *retraining oracle* is the same fine-tune without the forget rows, so it never saw them. The second benchmark, RWKU

(Jin et al., 2024), asks a model to forget real people it learned in pretraining, so the base model is the reference; Appendix F describes its rows, anchors and nine level-3 (adversarial) probe types.

Models and shared settings. Qwen2.5-3B-Instruct (Qwen Team, 2024) is the model unlearned in every arm, and Llama-3.2-3B-Instruct (Grattafiori et al., 2024) repeats five of the arms with its own reference and oracle (Appendix I). Qwen2.5-7B-Instruct is the probe generator and the semantic judge throughout. Every arm uses LoRA rank 16 on the attention projections, AdamW at learning rate 10^{-5}, batch size 4, $\beta = 0.1$, retain weight $\alpha = 1$ and anchor weight $\lambda = 1$ where an anchor is present; none of these was searched. Arms train for 16 or 32 epochs on TOFU-Rel, NPO also for 8 to trace its epoch sweep, and for 8 or 16 on RWKU.

The detector and its check. A response reproduces a when two criteria agree. The lexical criterion is a recall of at least 0.8 over the protected words of a, as in leak@k (Reisizadeh et al., 2026); the semantic criterion is the judge's greedy verdict that the response conveys the specific information of a. An item enters the reported population only if the reference reproduces a under the direct question, since unlearning can remove only what the reference knew. Before any comparison between objectives, the detector had to pass an *instrument gate*. The oracle never saw the forget rows, so its recovery estimates how often the detector fires when nothing was disclosed; the gate requires that rate at or below 0.2 and the reference's at or above 0.6 (Section 7 states what this cannot rule out). Three detector rules pass the gate, and every significant comparison below holds under all three and under a judge from another model family (Appendix G).

Held-out strategies and entities. The four *training strategies* Π_{tr} are rephrasing, inversion, two-step composition and a follow-up conditioned on the response; the six *held-out strategies* Π_{ho} share no template text with them, although two are alternative templates of rephrasing and of the follow-up (Appendix E). Every forget row is unlearned in every run, but during training the generator runs only under Π_{tr} and only on the generator-train entities of a fixed entity split, 60% of the entities (Appendix E). No development entity, test entity or held-out strategy therefore reaches a buffer or a gradient. NRU-G thus mines probes for only part of the forget set (one run mined probes for 60 of the 103 items; Appendix H), and its other items train on the direct question alone, whereas the fixed probe suite (below) covers every item. The split restricts the generator and not the data, which favours the fixed-suite arms in every comparison with NRU-G, most of all in the held-out columns.

Evaluation. Every adapter is probed by all ten strategies at $B = 8$ on every forget item, and each episode runs its full budget with the reference answering the same queries. We report recovery (worst@B) with its held-out-strategy and test-entity variants; *suppression* $S_0 \in [0, 1]$, the normalised rise in the answer's negative log-likelihood (NLL) under the direct question relative to the reference, 0 meaning no rise; and anchor and retain accuracy, the share of those questions answered correctly. Benign *continued training* (Hu et al., 2025) fine-tunes each final adapter for one hundred Adam steps at 10^{-5} on rows of TOFU's standard retain split (retain90), followed by the full probe evaluation. Across the three seeds, 192 to 198 of the 200 retain rows also appear in the continued-training data, so those data overlap the retain rows on which every arm and the simulated fine-tune train (Section 7).

Baselines. The baselines are NPO (Zhang et al., 2024a); NPO with the anchor term; and *neighbourhood-augmented unlearning* on a *fixed probe suite*, generated once before training by a generative model from seven transformation families such as paraphrase and back-translation, about 1.58 probes per forget row (Appendix E, Table 3). It trains on the suite uniformly (NAU) or with weights over the families adapted during training (DRO-NAU), both with the anchor. The other baselines are BalDRO-DV (Shao et al., 2026), a Kullback–Leibler-penalised worst case over forget samples in its Donsker–Varadhan (DV) form, on the same probe rows at its published temperature; and response-mining unlearning (RULE-NPO; Reisizadeh et al., 2026), which every two epochs samples 32 responses per forget prompt and adds at most eight leaking ones per item. The axis-2-only control, NPO + anchor + stability, adds the simulated fine-tune to NPO with the anchor and has no generator. Other objectives that constrain only the forget question (gradient ascent, gradient difference, SimNPO, RMU) were trained but not probed, so the comparison does not cover them.

Controls. Seven *controls* each remove one or two components of NRU-G, and Table 1c names what each removes (its last row compares NRU-S). Among them, the fixed pool drops the refresh at a matched decoding budget, uniform replay replaces the soft maximum by the mean, the same-data

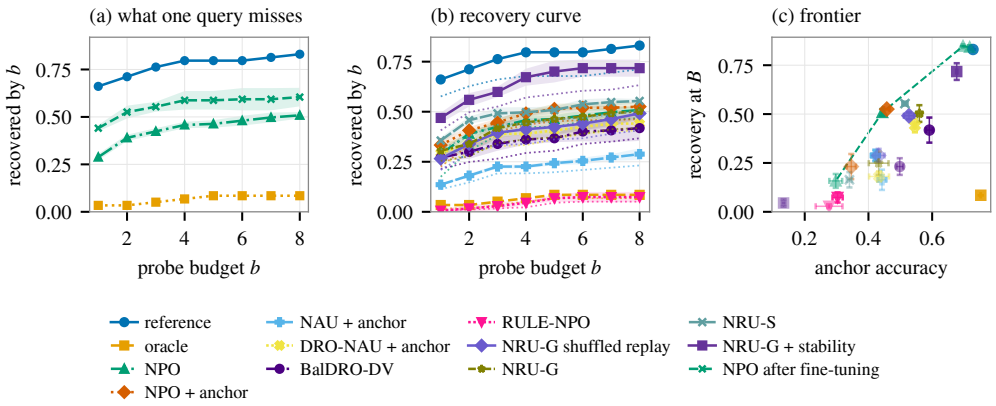

Figure 2: Recovery on TOFU-Rel (Qwen2.5-3B, 16 epochs; shading: one standard deviation over three seeds). *(a)* Share of answerable facts recovered by at least one of ten strategies within b queries each, for the reference, NPO, NPO after continued training and the oracle. *(b)* The same for more arms, over all ten strategies (solid) and the six held-out ones (dotted). *(c)* Recovery at $B = 8$ (lower is better) against anchor accuracy (higher is better); faded markers are other epochs, and the dashed line joins NPO's epoch sweep.

control drops both by training uniform replay on a finished run's final buffer, and the constraint control enforces the anchor as a constraint (Appendix C).

Protocol. Four three-point grids (Table 4) are the whole hyperparameter search, and every arm is read at the last iterate of training, never at a best-of-training checkpoint. Every $\pm$ is a standard deviation over three seeds at one fixed entity split. Every interval is a 95% bootstrap that resamples whole authors on item-by-item paired differences, so it carries item and author variance but neither split nor seed variance (Appendix E). Every hyperparameter and the answerability mask were fixed before a test cell was read.

5 Results

Unless stated otherwise, every rate is over the reference-answerable items of TOFU-Rel on Qwen2.5-3B, at probe budget $B = 8$ and 16 epochs. TOFU-Rel is the one setting in which the measurement separates objectives (Section 7), and Table 6 in Appendix H lists every arm.

5.1 P1: The direct question is not the neighbourhood

The reference reproduces 59 of the 103 forget answers under the direct question, and these items form the reported population. Within eight queries each, the ten strategies (eighty queries in all) recover 0.83 of them from the reference and 0.08 from the retraining oracle. The strongest single strategy reaches 0.59 with its eight queries (Appendix G), and the ten strategies reach 0.66 with their first query each.

At 16 epochs NPO reaches suppression 0.37, yet the ten strategies still recover 0.51 of the answerable facts and the strongest strategy alone 0.36; the held-out strategies and the test entities yield nearly as many (Figure 2a,b; Table 6). Training for 32 epochs cuts recovery to 0.16 but lowers anchor accuracy from 0.44 to 0.30. Any arm can reach a low recovery by destroying the entity's other facts, so every comparison below is read both at matched anchor accuracy and at identical epochs (Figure 2c).

5.2 RQ1: Query neighbourhoods at matched utility and identical epochs

At matched anchor accuracy, training on a query neighbourhood lowers recovery for each of the three constructions in Table 1b, and so does response mining. Each arm is read at its default setting and at its epochs closest to an NPO run in anchor accuracy (Table 7 gives the rest). The paired difference

Table 1: Paired differences in recovery, negative meaning fewer facts recovered, with 95% intervals from a bootstrap over the 20 authors; *: the interval excludes zero. The pairs in (b) often join unequal epochs and compute, and the last row of (c) compares NRU-S with NRU-G. Tables 7 and 8 give every arm.

comparison	Δ recovery, 16 epochs	Δ recovery, 32 epochs
(a) NRU arms against NPO at the same epochs		
NRU-G	-0.006 $[-0.10, +0.08]$	$+0.090$ $[+0.03, +0.15]^*$
NRU-S	$+0.045$ $[-0.03, +0.10]$	$+0.006$ $[-0.02, +0.05]$
(b) Against the NPO run nearest in anchor accuracy		
NAU + anchor (vs NPO at 16)	–	-0.345 $[-0.45, -0.24]^*$
DRO-NAU + anchor (vs NPO at 16)	–	-0.328 $[-0.45, -0.21]^*$
NRU-G (vs NPO at 16)	–	-0.260 $[-0.35, -0.18]^*$
RULE-NPO (vs NPO at 32)	-0.085 $[-0.14, -0.04]^*$	–
NPO + anchor (vs NPO at 16)	$+0.017$ $[-0.02, +0.06]$	–
(c) NRU-G with components removed (the control in parentheses), against NRU-G		
per-item soft maximum (uniform replay)	-0.051 $[-0.10, +0.01]$	$+0.096$ $[+0.05, +0.15]^*$
refresh (fixed pool)	-0.017 $[-0.08, +0.03]$	$+0.090$ $[+0.04, +0.14]^*$
both of these (same data)	-0.006 $[-0.07, +0.06]$	$+0.113$ $[+0.08, +0.15]^*$
buffer persistence (no replay)	$+0.068$ $[-0.03, +0.16]$	-0.056 $[-0.11, +0.01]$
the responses (nonadaptive)	$+0.006$ $[-0.07, +0.07]$	-0.017 $[-0.08, +0.05]$
the item's own probes (shuffled replay)	-0.011 $[-0.06, +0.03]$	$+0.040$ $[+0.00, +0.08]$
the anchor penalty (constraint)	$+0.014$ $[-0.04, +0.06]$	$+0.023$ $[-0.02, +0.08]$
the generator model (NRU-S)	$+0.051$ $[-0.04, +0.14]$	-0.085 $[-0.13, -0.04]^*$

is -0.345 $[-0.45, -0.24]$ for NAU + anchor, -0.328 for DRO-NAU + anchor and -0.260 for NRU-G, and -0.085 for RULE-NPO, whose added rows are responses rather than queries. Each is against the nearest NPO run, and in none does a bootstrap replicate out of five thousand cross zero. These pairs cost 8 to 23 times the wall-clock time of the NPO run each is paired with (Table 3). Adding the anchor alone has no detectable effect at matched utility ($+0.017$, $p = 0.508$), although at 32 epochs NPO + anchor recovers more than NPO ($+0.07$ $[+0.02, +0.13]$, Table 8). **At matched utility a query neighbourhood, and not the anchor or extra epochs alone, lowers recovery across rephrasings.**

At identical epochs NRU-G does not beat NPO (Table 1a). NRU-G does not separate from NPO at 16 epochs (-0.006 $[-0.10, +0.08]$, $p = 0.935$); the interval rules out an advantage larger than 0.10 but does not establish equivalence. At 32 epochs NRU-G recovers more than NPO ($+0.090$ $[+0.03, +0.15]$), and NAU + anchor recovers less than NRU-G (-0.08, Table 8) at similar anchor accuracy (0.44 against 0.43, Table 7) and a third of its cost per epoch. NAU's advantage over NPO at 16 epochs is not only due to its overlap with the held-out strategies: restricted to the two held-out strategies that no suite family covers, it is still -0.096 $[-0.16, -0.04]$ ($p = 0.001$).

5.3 RQ2: WHICH COMPONENTS OF NRU-G MATTER

Each control is compared with NRU-G itself (Table 1c), because a control's difference from NPO mixes the neighbourhood with the removed component. **At 32 epochs, removing the refresh or the per-item soft maximum raises recovery.** Removing the soft maximum alone raises it by $+0.096$ $[+0.05, +0.15]$, removing the refresh alone at a matched decoding budget by $+0.090$ $[+0.04, +0.14]$, and removing both over the same mined probes by $+0.113$ $[+0.08, +0.15]$. Each survives a Holm correction over this prespecified family of four contrasts ($p \leq 0.003$), with intervals that do not include seed variance (Section 4). The fourth contrast, clearing the buffer at every refresh, lowers recovery, not detectably under the reported judge (-0.056, $p = 0.120$) and detectably under the second judge (Appendix G), so if anything the buffer's persistence raises recovery. At 16 epochs none of the four separates. NRU-G itself recovers more than NPO at 32 epochs, so these components explain the generator arm's behaviour, not a gain over NPO.

Neither the model's responses nor the item's own probes show a detectable benefit. A generator denied every response matches the full generator ($+0.000$ against -0.006 at 16 epochs, each against NPO). Replaying another item's probes matches the item's own (-0.011, $p = 0.684$) while still

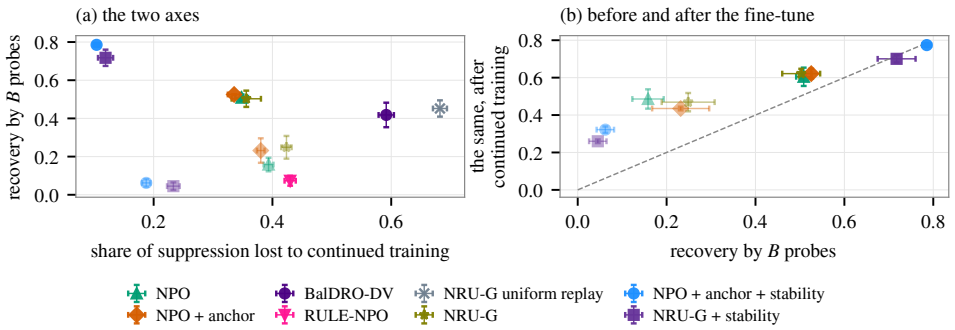

Figure 3: The two axes, one marker per arm (32 epochs faded); error bars: one standard deviation over three seeds. *(a)* Recovery at $B = 8$ against the share of suppression that one hundred steps of continued training remove; an arm robust on both axes would sit at the lower left. *(b)* Recovery after continued training against recovery before it; points on the diagonal did not change.

recovering less than uniform replay over the same rows at 32 epochs (-0.056, $p = 0.008$). None of the mined queries matches the fixed suite after normalisation (Appendix H), so the generator's neighbourhood differs from the fixed suite's, yet adaptivity shows no detectable benefit.

NRU-S does not beat NPO either. NRU-S keeps the refresh and the soft maximum but takes its probes from the fixed suite, so it needs no generator model and costs three times NPO's seconds per epoch rather than eleven (Table 3). It separates from NPO at neither 16 epochs ($+0.045$, $p = 0.220$) nor 32 ($+0.006$, $p = 0.868$), and its gain over NRU-G at 32 epochs, -0.085 $[-0.13, -0.04]$, costs 0.09 anchor accuracy.

5.4 RQ3: THE SECOND AXIS, AND WHY THE TWO-AXIS OBJECTIVE FAILS

P2: the two axes are separate. One hundred steps of continued training remove between a third and two thirds of the suppression of each of the six arms trained on queries alone that were measured after it, and none is clearly more stable than NPO (Figure 3a); NAU, DRO-NAU and NRU-S were not measured after continued training. Continued training also raises recovery, more after 32 epochs than after 16 (Figure 3b, Table 9a): NPO rises by 0.10 at 16 epochs and by 0.33 at 32.

Simulating the fine-tune inside the update does not hold both axes. What the simulated steps undo grows with how far the update has pushed the forget loss: at 32 epochs they undo shares of 0.46 and 0.61 of the two stability arms' late forget NLL, against 0.07 at 16 (Table 9b). At 16 epochs these arms barely unlearn (recovery 0.79 and 0.72). At 32 epochs they reach the lowest recovery of any arm except response mining, but their anchor accuracy falls to 0.13 (Table 9a). Whether an intermediate number of epochs or a different inner step avoids this is untested. **Training against the second axis held it only where the first axis was still open.**

6 RELATED WORK

Appendix J compares each approach below with this work in full; every baseline was run by us with our own arms' search budget. **Evaluations** recover facts from models that pass their benchmark, through input transformations (Joshi et al., 2024; Dahal et al., 2026), sampling at the same prompt (Reisizadeh et al., 2026), multi-turn interaction (Pan & Wang, 2026), a budgeted black-box generator (Au et al., 2026), searched attack prompts (Rybak et al., 2026) and adversarial suites (Gupta et al., 2026; Jang et al., 2025); benchmark studies find current unlearning measures weak (Thaker et al., 2025; Fan et al., 2025b). Together they define the probe generator that NRU turns into a training signal. **Methods that enlarge the forget set** mine responses at fixed prompts (Reisizadeh et al., 2026; Li et al., 2026a), take a worst case over forget samples (Shao et al., 2026) or train on queries generated before training (Singh et al., 2025; Peleg et al., 2026). None chooses its queries by reading the current model's responses, and NRU reuses the dual of Shao et al. (2026) over a different set. **Methods robust in weight space** target the continued training that can reverse unlearning in a few steps (Hu et al., 2025). They harden unlearning against fine-tuning at fixed prompts (Tamirisa et al.,

2025; Gao et al., 2025; Sheshadri et al., 2025; Wang et al., 2025), flatten the loss landscape (Fan et al., 2025a; Wu et al., 2025; Yan et al., 2026), or solve a min–max problem over hidden states (Yuan et al., 2025; Li et al., 2026b) or weights (Cai et al., 2026). Generator–defender co-training exists for refusal alignment (Liu et al., 2026), and concept erasure in image generation trains against optimised adversarial prompts (Zhang et al., 2024b). We know of no LLM unlearning method that trains against both axes at once.

7 LIMITATIONS

What a low recovery rate does not mean. Recovery (Equation 1) is defined for a strategy set, a budget and a detector. A model that resists these ten strategies at eight queries each has not been shown to resist an eleventh strategy, a longer conversation, a generator that reads the logits or a human expert, and our held-out strategies were written by the same authors as the training ones.

The loss is not a bound, and the detector is fallible. The link from the loss the objective lowers (Equation 2) to the reproductions that recovery counts is measured, not proved, and it differs across arms (P3). The lexical criterion misses paraphrase, and the judge errs in both directions. The judge is the same model as the generator, although a judge from another family reproduces every significant comparison and turns one null into a significant effect (Appendix G). Every rate is over the items the reference reproduces under the direct question, a selection that favours the reference. The oracle's rate is an estimate, not a bound: an oracle can reach a correct answer by inference or chance, and without an independently labelled sample it checks the detector only on the oracle's own responses.

What limits NRU-G. It mines probes only for generator-train items (Section 4), under a disjunction of the two criteria that "recovers" 0.85 of the oracle's items (Table 5), so its buffer may hold many false positives. Either could explain why adaptivity shows no detectable benefit; neither was tested.

Continued-training data. These data contain most of the retain rows (Section 4), so the second axis is measured on largely seen text, which may favour the two stability arms.

Three measured failure modes (Appendix I). First, a protected answer of one or two words, as on RWKU, is contained in any fluent response, so every arm reads at least 0.99. Second, facts learned in pretraining admit no retraining oracle and hence no check of the detector. Third, this configuration carried unchanged to a second model family puts every arm at or below the oracle's rate, with anchor accuracy at half the reference's or less, so neither P1 nor the comparison of arms can be read there.

Deployment. NRU-G needs a 48 GB card at batch size 4 and eleven times NPO's wall-clock time per epoch, and NRU-G + stability thirteen times (Table 3).

What is unverified. We test one model scale, adapter rank and batch size, two benchmarks, three seeds, one entity split and one probe budget. So nothing bears on full fine-tuning, larger models or forget sets beyond 103 rows, no interval carries split or seed variance, and no claim extends past the fixed continued-training budget. Only the four switch contrasts are corrected for multiplicity, and response mining's sampling budget was varied at 16 epochs alone, so whether its anchor damage is intrinsic is untested. With β, λ and α fixed, NPO's frontier is traced over epochs alone. Every searched hyperparameter was frozen on ten development items, and for NRU-G's temperature and DRO-NAU's family weights the frozen choice proved the worse one on test (Appendix E).

8 CONCLUSION

Training over a query neighbourhood rather than at the forget question alone lowers what a probe generator recovers at matched utility, at several times NPO's cost, and inside NRU-G at 32 epochs the refresh and the per-item soft maximum each matter. Adaptivity shows no detectable benefit at identical epochs, and a simulated fine-tune inside the update does not hold both axes. The practical advice is to train on a query neighbourhood built cheaply before training, such as a fixed probe suite, and to measure what an objective leaves reachable on both axes: a passing likelihood check is not evidence of removal.

AI USE STATEMENT

As research instruments generative models are components of the method, each specified where it is used: the probe generator and the instruction model behind the semantic half of the detector verdict (Sections 2 and 4); the cross-lingual and rephrasing strategies (Appendix E); the answerability gate of the RWKU split, which is the base model answering its own benchmark (Appendix F); the second judge and the entailment model of the detector audit (Appendix G); and the generator that produced the fixed probe suite (Section 4). *As assistants* generative models drafted analysis and plotting scripts, LaTeX and prose, which the authors ran, reviewed and checked against the recorded measurements. The revision of this manuscript used an LLM assistant (Claude) for editing and rewriting the text, for restating formal results as measured properties, and for the layout of tables and figures; the authors reviewed every change. No hypothesis, research idea, experimental design, mathematical claim or proof originated with a model, no result was interpreted by one, and no number, citation or claim was accepted from one without verification against the artefact or the cited text. We have reviewed all AI-assisted work and take responsibility for the final content of this paper and the artefact it releases.

ETHICS STATEMENT AND BROADER IMPACTS

The study collects no personal data and recruits no person. Both benchmarks are used as released, TOFU (Maini et al., 2024) under the MIT licence and RWKU (Jin et al., 2024) under CC BY 4.0, and the models under Apache 2.0, the Qwen research licence and the Llama 3.1 and 3.2 community licences, each permitting the academic use made here. RWKU's subjects are twenty public figures whose 482 admitted rows carry encyclopedic biography, among them one publicly discussed health fact and two rows about a reported conviction; no admitted row holds an identifier or profanity. The study cost 327 GPU-hours (Appendix E), about 172 kWh and 64 kg of CO_2 equivalent at an average grid carbon intensity.

Dual use runs both ways. Releasing the probe strategies lets the same code be pointed at any deployed model, which we judge worthwhile because the probes are ordinary paraphrasing and follow-up questions rather than novel exploits and because an operator needs that code to audit their own model. The methods carry the opposite risk: an objective that lowers a measured rate without removing the fact is what a developer would reach for to pass an audit, so a method that suppresses well and removes nothing is a compliance risk rather than a safety gain. Our findings cut against erasure claims: an objective passing likelihood-based checks, even one trained against an adaptive generator, should be called suppression rather than removal, and the cost of reading it otherwise falls on the person who asked for the deletion and on any regulator who accepts a benchmark score as evidence of it.

REPRODUCIBILITY STATEMENT

Section 4 and Appendix E give the data splits, models, hyperparameters, seeds, search and statistical protocol; Appendix C gives the training loop, and Table 2 its constants; Appendix D and Table 3 give compute and memory; Appendix G and Table 5 give the detector rules and their check; Appendix F gives the RWKU build; and Appendix B states the assumptions behind Equation 4. The anonymized code in the supplementary material runs the whole pipeline and contains the relation-level split, the RWKU build, the probe strategies, the replay buffers of every adaptive run and the per-run compute accounting. It stores every probe episode turn by turn, so every reported rate can be recomputed under a different budget, strategy set or detector rule without rerunning a model; the reference's responses to issued queries are not stored (Appendix G). Seeds are fixed, and a deterministic mode, off for the reported runs, reproduces an adapter bit for bit; without it, a rerun from a run's specification matches up to GPU nondeterminism. Adapter weights are withheld for size.

REFERENCES

Ashley Hoi-Ting Au, Meghdad Kurmanji, William F. Shen, Nicholas D. Lane, and Ligang He. Extracting forgotten prompts from targeted unlearned models. *arXiv preprint arXiv:2609.03662*, 2026.

Junhao Cai, Dohun Kim, Dowon Kim, Sung Il Choi, Chengjun Jin, Juhyun Park, and Changhee Joo. Retain-neutral surrogates for min-max unlearning. *arXiv preprint arXiv:2605.05871*, 2026.

Chahana Dahal, Ashutosh Balasubramaniam, and Zuobin Xiong. GONE: Structural knowledge unlearning via neighborhood-expanded distribution shaping. *arXiv preprint arXiv:2603.12275*, 2026.

Monroe D. Donsker and S. R. Srinivasa Varadhan. Asymptotic evaluation of certain Markov process expectations for large time, I. *Communications on Pure and Applied Mathematics*, 28(1):1–47, 1975.

Chongyu Fan, Jinghan Jia, Yihua Zhang, Anil Ramakrishna, Mingyi Hong, and Sijia Liu. Towards LLM unlearning resilient to relearning attacks: A sharpness-aware minimization perspective and beyond. In *International Conference on Machine Learning (ICML)*, 2025a. arXiv:2502.05374.

Chongyu Fan, Changsheng Wang, Yancheng Huang, Soumyadeep Pal, and Sijia Liu. LLM unlearning under the microscope: A full-stack view on methods and metrics. *arXiv preprint arXiv:2510.07626*, 2025b.

Hongcheng Gao, Tianyu Pang, Chao Du, Taihang Hu, Zhijie Deng, and Min Lin. Meta-unlearning on diffusion models: Preventing relearning unlearned concepts. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025. arXiv:2410.12777.

Aaron Grattafiori et al. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

Ayush Gupta, Hima Varshini Surisetty, Sreevidya Bollineni, Varad Ingale, Tuhina Tripathi, Abhishek Lalwani, Somya Chatterjee, and Sadid Hasan. Can LLMs truly forget? Revealing unlearning gaps through adversarial evaluation. *arXiv preprint arXiv:2608.21606*, 2026.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations (ICLR)*, 2022. arXiv:2106.09685.

Shengyuan Hu, Yiwei Fu, Zhiwei Steven Wu, and Virginia Smith. Unlearning or obfuscating? Jogging the memory of unlearned LLMs via benign relearning. In *International Conference on Learning Representations (ICLR)*, 2025. arXiv:2406.13356.

Yeonwoo Jang, Shariqah Hossain, Ashwin Sreevatsa, and Diogo Cruz. Prompt attacks reveal superficial knowledge removal in unlearning methods. In *COLM 2025 Workshop on Socially Responsible Language Modelling Research (SoLaR)*, 2025.

Zhuoran Jin, Pengfei Cao, Chenhao Wang, Zhitao He, Hongbang Yuan, Jiachun Li, Yubo Chen, Kang Liu, and Jun Zhao. RWKU: Benchmarking real-world knowledge unlearning for large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. Datasets and Benchmarks Track, arXiv:2406.10890.

Abhinav Joshi, Shaswati Saha, Divyaksh Shukla, Sriram Vema, Harsh Jhamtani, Manas Gaur, and Ashutosh Modi. Towards robust evaluation of unlearning in LLMs via data transformations. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, 2024. arXiv:2411.15477.

Kemou Li, Qizhou Wang, Yue Wang, Fengpeng Li, Jun Liu, Bo Han, and Jiantao Zhou. LLM unlearning with LLM beliefs. In *International Conference on Learning Representations (ICLR)*, 2026a. arXiv:2510.19422.

Pengyu Li, Lingling Zhang, Zhitao Gao, Yanrui Wu, Yuxuan Dong, Huan Liu, Bifan Wei, and Jun Liu. AGT^{AO}: Robust and stabilized LLM unlearning via adversarial gating training with adaptive orthogonality. In *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 13585–13600, 2026b.

Mickel Liu, Liwei Jiang, Yancheng Liang, Simon Shaolei Du, Yejin Choi, Tim Althoff, and Natasha Jaques. Chasing moving targets with online self-play reinforcement learning for safer language models. In *International Conference on Machine Learning (ICML)*, 2026. arXiv:2506.07468.

Pratyush Maini, Zhili Feng, Avi Schwarzschild, Zachary C. Lipton, and J. Zico Kolter. TOFU: A task of fictitious unlearning for LLMs. In *Conference on Language Modeling (COLM)*, 2024. arXiv:2401.06121.

Ruihao Pan and Suhang Wang. A comprehensive evaluation of LLM unlearning robustness under multi-turn interaction. *arXiv preprint arXiv:2603.00823*, 2026.

Amit Peleg, Naman Deep Singh, Naama Pearl, Bibhabasu Mohapatra, and Matthias Hein. Forget narrowly, retain broadly: Unlearning as an asymmetric generalization problem. *arXiv preprint arXiv:2607.09236*, 2026.

Qwen Team. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*, 2024.

Hadi Reisizadeh, Jiajun Ruan, Yiwei Chen, Soumyadeep Pal, Sijia Liu, and Mingyi Hong. Leak@k: Unlearning does not make LLMs forget under probabilistic decoding. In *International Conference on Machine Learning (ICML)*, 2026.

Patryk Rybak, Paweł Batorski, Paul Swoboda, and Przemysław Spurek. REBEL: Hidden knowledge recovery via evolutionary-based evaluation loop. *arXiv preprint arXiv:2602.06248*, 2026.

Pengyang Shao, Naixin Zhai, Lei Chen, Yonghui Yang, Fengbin Zhu, Xun Yang, and Meng Wang. BalDRO: A distributionally robust optimization based framework for large language model unlearning. In *Proceedings of the ACM Web Conference 2026 (WWW '26)*, pp. 8874–8884, 2026.

Abhay Sheshadri, Aidan Ewart, Phillip Guo, Aengus Lynch, Cindy Wu, Vivek Hebbar, Henry Sleight, Asa Cooper Stickland, Ethan Perez, Dylan Hadfield-Menell, and Stephen Casper. Latent adversarial training improves robustness to persistent harmful behaviors in LLMs. *Transactions on Machine Learning Research*, 2025. arXiv:2407.15549.

Naman Deep Singh, Maximilian Müller, Francesco Croce, and Matthias Hein. Unlearning that lasts: Utility-preserving, robust, and almost irreversible forgetting in LLMs. *arXiv preprint arXiv:2509.02820*, 2025.

Rishub Tamirisa, Bhrugu Bharathi, Long Phan, Andy Zhou, Alice Gatti, Tarun Suresh, Maxwell Lin, Justin Wang, Rowan Wang, Ron Arel, Andy Zou, Dawn Song, Bo Li, Dan Hendrycks, and Mantas Mazeika. Tamper-resistant safeguards for open-weight LLMs. In *International Conference on Learning Representations (ICLR)*, 2025. arXiv:2408.00761.

Pratiksha Thaker, Shengyuan Hu, Neil Kale, Yash Maurya, Zhiwei Steven Wu, and Virginia Smith. Position: LLM unlearning benchmarks are weak measures of progress. In *IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, 2025. arXiv:2410.02879.

Changsheng Wang, Yihua Zhang, Jinghan Jia, Parikshit Ram, Dennis Wei, Yuguang Yao, Soumyadeep Pal, Nathalie Baracaldo, and Sijia Liu. Invariance makes LLM unlearning resilient even to unanticipated downstream fine-tuning. In *International Conference on Machine Learning (ICML)*, 2025. arXiv:2506.01339.

Wenhan Wu, Zheyuan Liu, Chongyang Gao, Ren Wang, and Kaize Ding. Beyond sharp minima: Robust LLM unlearning via feedback-guided multi-point optimization. *arXiv preprint arXiv:2509.20230*, 2025.

Han Yan, Zheyuan Liu, and Meng Jiang. Dual-space smoothness for robust and balanced LLM unlearning. In *International Conference on Learning Representations (ICLR)*, 2026. arXiv:2509.23362.

Hongbang Yuan, Zhuoran Jin, Pengfei Cao, Yubo Chen, Kang Liu, and Jun Zhao. Towards robust knowledge unlearning: An adversarial framework for assessing and improving unlearning robustness in large language models. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2025. arXiv:2408.10682.

Ruiqi Zhang, Licong Lin, Yu Bai, and Song Mei. Negative preference optimization: From catastrophic collapse to effective unlearning. In *Conference on Language Modeling (COLM)*, 2024a. arXiv:2404.05868.

Yimeng Zhang, Xin Chen, Jinghan Jia, Yihua Zhang, Chongyu Fan, Jiancheng Liu, Mingyi Hong, Ke Ding, and Sijia Liu. Defensive unlearning with adversarial training for robust concept erasure in diffusion models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024b. arXiv:2405.15234.

A Roadmap and notation

Appendix B states the assumptions behind the soft maximum of Section 3 and derives the first-order two-axis gradient; Appendix C gives the training loop and its controls, and Appendix D its compute and memory. Appendix E gives the protocol of Section 4 with the hyperparameter search and the preliminary study behind the baselines, Appendix F the second benchmark and Appendix G the detector rules, their check and a second judge. Appendix H holds the full results behind Section 5, Appendix I bounds the failure modes of Section 7, and Appendix J compares the approaches of Section 6.

A subscript q indexes one forget item, and constants are defined where used, in Table 3, Appendix E and Algorithm 1.

Table 2: Symbols carrying the argument, with their TOFU-Rel values; the evaluation probe budget is $B = 8$.

Symbol	Meaning (value)	Symbol	Meaning (value)
$\mathcal{D}_F,\ \mathcal{D}_R$	forget and retain sets (103, 200 rows)	$\mathcal{A}(q)$	anchor set, the entity of q in its other relations (287 rows)
$\mathcal{Q}(q)$	reachable query set, every query a strategy can issue from q (finite, too large to enumerate)	Q_q	the sampled part of $\mathcal{Q}(q)$ the gradient sees (at most $1+k$ queries)
$\mathcal{M}_q$	replay buffer of probes that reproduced the answer (cap 24)	$\Theta_T(\theta)$	reachable weight set, the benign fine-tune's trajectory ($T = 100$)
$\Pi_{\mathrm{tr}},\ \Pi_{\mathrm{ho}}$	training and held-out probe strategies (4, 6)	$\mathrm{FT}_K(\theta)$	weights after K simulated benign steps ($K = 4$)
$f_\theta,\ \theta$	the model being unlearned and its parameters (Qwen2.5-3B)	θ_0	frozen pre-unlearning reference
$\ell(\theta; x, a)$	per-query NPO loss	τ	soft-maximum temperature (1.0)
R_B, worst@B	recovery within B queries, for one strategy and as a union over ten	S_0	suppression, $\frac{1}{n}\sum_i (1 - e^{-[\mathrm{NLL}_\theta(a_i\mid q_i) - \mathrm{NLL}_{\theta_0}(a_i\mid q_i)]_+})$, in $[0,1]$

B Assumptions and derivations

This appendix states what the two formal facts used by the method assume: the bounds on the soft maximum (Equation 4) and the first-order two-axis gradient. Neither is a guarantee about recovery.

Assumptions.

(A1) Finite sampled set. Q_q is finite and non-empty, $1 \le |Q_q| \le 1 + k$, with the direct question always present, since a short buffer yields fewer than k probes.

(A2) Finite per-query loss. $\ell(\theta; x, a)$ is real and finite, which holds because $p_\theta(a \mid x) \in (0, 1)$ under a softmax over finite logits.

(A3) Positive temperature. $\tau > 0$.

The soft maximum. Under (A1)–(A3) the per-item loss of Equation 2 satisfies

$$\frac{1}{|Q_q|}\sum_{x\in Q_q} \ell(\theta; x, a) \;\le\; L_q(\theta) \;\le\; \max_{x\in Q_q} \ell(\theta; x, a), \qquad L_q(\theta) \;\ge\; \max_{x\in Q_q} \ell(\theta; x, a) - \tau \log |Q_q|. \quad (4)$$

The left inequality is Jensen's inequality for the convex exponential. For the right, every term of the mean is at most $\exp(\max_x \ell/\tau)$ and the logarithm is monotone. For the last, the mean is at least

$\exp(\max_x \ell/\tau)/|Q_q|$. The limits follow, the softmax weight concentrating on the maximiser as $\tau \to 0$ and a second-order expansion leaving the mean as $\tau \to \infty$. The Donsker–Varadhan identity is exact for a finite set: for any u with full support, $\tau \log \mathbb{E}_u[e^{\ell/\tau}] = \max_w \{\mathbb{E}_w[\ell] - \tau\,\mathrm{KL}(w\|u)\}$ over distributions w on Q_q, with maximiser $w^\star \propto u\, e^{\ell/\tau}$. With u uniform this is L_q, a KL-penalised worst case whose weights $w^\star$ move with θ. The gradient of L_q is the $w^\star$-weighted average of the per-query gradients.

The first-order two-axis gradient. Write $J = \partial\,\mathrm{FT}_K(\theta)/\partial\theta$ and $v = \nabla L|_{\mathrm{FT}_K(\theta)}$. The exact gradient of the two-axis loss (Equation 3) is $J^\top v$, and the update uses v, so its error is $(J^\top - I)v$. For K plain gradient steps at rate η on an inner objective g, $J = \prod_{k<K}(I - \eta\nabla^2 g(\theta_k))$, which is close to I when ηK times the curvature along the inner path is small; Adam's inner state makes J neither symmetric nor of this product form. We measure neither the error nor the curvature. Section 5.4 reports how much the K inner steps change the forget loss, a property of the dynamics that bounds nothing about this error.

C THE TRAINING LOOP

Algorithm 1 is the loop of Figure 1 as implemented in the anonymized code in the supplementary material; its constants are in Table 2, Appendix E and the last paragraph below. Two carry provenance. The refresh period $m = 25$ is about one epoch of the relation-level split at batch 4. The simulated step's rate 3×10^{-4} was set once from four such steps on a finished NPO adapter, taking the forget NLL $0.69 \to 0.38$, about what the hundred-step continued-training evaluation produces. The one instruction model that serves as generator and judge never sees the answer when generating and sees only the query, answer and response when judging.

Algorithm 1 Generator-in-the-loop unlearning

Require: $\mathcal{D}_F$, $\mathcal{A}$, $\mathcal{D}_R$, θ_0, Π_{tr}, D_{tr}, E_{tr} (the generator-train entities)

$\theta \leftarrow \theta_0$; $\mathcal{M}_q \leftarrow \emptyset$; $\lambda \leftarrow \lambda_0$
for step $s = 0, 1, \ldots$ **do**
  if $s \bmod m = 0$ **then** ▷ refresh against current θ
    sample n_a items with $e(q) \in E_{\mathrm{tr}}$
    if no replay: clear every $\mathcal{M}_q$
    for all sampled (q, a), $\pi \in \Pi_{\mathrm{tr}}$ **do**
      $h \leftarrow ()$
      for $t = 1, \ldots, B_{\mathrm{tr}}$ **do**
        $x_t \leftarrow$ first gated draw of $\pi(q, h)$
        $y_t \leftarrow f_\theta(x_t)$
        $D_t \leftarrow D_{\mathrm{tr}}(y_t, a)$; $R_t \leftarrow D_{\mathrm{tr}}(f_{\theta_0}(x_t), a)$
        if $D_t \wedge R_t$ **then** $\mathcal{M}_q \mathrel{+}= x_t$ (cap C)
        if D_t **then break**
        $h \leftarrow h \,\|\, (x_t, y_t)$
  draw forget, anchor and retain batches
  $\theta' \leftarrow \mathrm{FT}_K(\theta)$ if two-axis, else θ
  for all items q in the batch **do**
    $Q_q \leftarrow \{q\} \cup S_k(\mathcal{M}_q)$
    $L_q \leftarrow \tau \log \frac{1}{|Q_q|} \sum_{x \in Q_q} e^{\ell(\theta'; x, a)/\tau}$
  $g \leftarrow \nabla_{\theta'} \big(\overline{L_q} + \lambda \mathcal{L}_{\mathrm{anc}} + \alpha \mathcal{L}_{\mathrm{ret}}\big)$
  $\theta \leftarrow \theta - \eta\, g$ ▷ first order: g moves θ, not θ'
  if constraint **then**
    $v \leftarrow \mathcal{L}_{\mathrm{anc}}(\theta) - \mathcal{L}_{\mathrm{anc}}(\theta_0) - \delta$
    $\lambda \leftarrow \min\{\lambda_{\max}, [\lambda + \eta_\lambda v]_+\}$
return θ, every $\mathcal{M}_q$, the refresh log

Controls as switches. The two-axis arm with $K = 0$ is NRU-G, with $\Pi_{\mathrm{tr}} = \emptyset$ the axis-2-only control. Uniform replay takes the mean over Q_q, the $\tau \to \infty$ limit. The fixed pool loads $\mathcal{M}_q$ once from the same strategies run against θ_0 and never refreshes, 17 rounds of 12 items matching the 17

refreshes at 16 epochs and 33 against 34 at 32. The same-data control trains uniform replay on a finished run's final buffer, and BalDRO-DV replaces per-item grouping by one soft maximum over the batch. RULE-NPO keeps the direct questions and, from epoch 2 and every two after, samples 32 responses per forget prompt, adding up to eight distinct leaking ones and the correct sampled retain answers.

The buffer and the refresh. A successful probe is stored with the strategy and step that produced it, deduplicated up to punctuation and capped at C entries, oldest evicted first. The first refresh runs at step zero against the frozen reference. The refresh runs on n_a generator-train items at budget B_{tr} in stateless mode under a training detector D_{tr} chosen for recall so the buffer fills, more lenient than the reported rule, so some buffered queries are judge false positives; that is a property of the training signal and of no reported number. An episode stops at its first reproduction, and the no-replay control clears every buffer at each refresh. The simulated step logs the forget loss before and after its K steps, which Table 9 reads, and writes into a separate gradient slot so the outer gradient is untouched. The gradient of the soft maximum is the softmax-weighted average of the per-query gradients at temperature τ.

Utility as a penalty or a constraint. As a *penalty* the anchor term is $\lambda\mathcal{L}_{\mathrm{anc}}$ at fixed λ. As a *constraint* it is $\mathcal{L}_{\mathrm{anc}}(\theta) \le \mathcal{L}_{\mathrm{anc}}(\theta_0) + \delta$ on each anchor batch, enforced by a primal–dual step: the parameters move on the Lagrangian at the current multiplier, then $\lambda_{t+1} = \min\{\lambda_{\max}, [\lambda_t + \eta_\lambda (\mathcal{L}_{\mathrm{anc}}(\theta_t) - \mathcal{L}_{\mathrm{anc}}(\theta_0) - \delta)]_+\}$ from $\lambda_0 = \lambda = 1$, with slack $\delta = 0.05$, dual step $\eta_\lambda = 0.1$ and cap $\lambda_{\max} = 20$. The problem is not convex, so the primal–dual step carries no convergence guarantee; in every constraint run the multiplier stayed below 2.2, well under the cap.

D Compute and memory

This appendix derives the step costs of Section 3.4 and measures the deployment costs of Section 7.

Cost per optimiser step. A forward pass over a sequence of length L costs $F = \Theta(N_\ell(Ld^2 + L^2 d))$ for N_ℓ layers of width d and a backward pass about $2F$, so a step with g gradient and r reference rows costs $(3g+r)F$, with activation memory $\Theta((g+r)LdN_\ell)$. Trainable parameters are 7,372,800 for every arm at LoRA rank 16; the outer optimiser holds twice that in Adam state and the simulated step a further snapshot and its own two Adam moments while it runs. The training data, optimiser, batch, sequence length and number of epochs are identical across arms; what one step computes is not, so it is measured, and the fixed pool is matched to NRU-G's decoding budget. The refresh decodes are bounded rather than counted, an episode stopping at its first reproduction under D_{tr}: for a 16-epoch run the bound, $n_a|\Pi_{\mathrm{tr}}|B_{\mathrm{tr}}c$ decodes at each of 17 refreshes, is 9792 against the 5775 ± 53 issued. Probe evaluation is the same run for every adapter and took 3064–4755 seconds per adapter.

Deployment overheads. The generator arms pay their cold start at the step-zero refresh against the frozen reference, whereas response mining has the cheapest first epoch of any arm because mining begins at epoch 2. Peak allocation understates what a card must hold: the allocator reserves 41.3 GiB where 40.4 GiB is allocated, which puts NRU-G over a 40 GiB card.

E Protocol details

Operating constraints. Sequences truncate at 512, a generated query at 64 new tokens and a response at 48, so no strategy gains recovery from a long prompt. Training needs the forget answers, an anchor set, a retain set and a frozen reference; the generator needs only the forget question; reading the result needs a retraining oracle, which is why a pretrained-knowledge substrate cannot be read (Appendix I). Batch size 4 uses no gradient accumulation, and the loop's other constants are listed with Table 3. There is no graph, no retrieval corpus, and no context beyond the last four turns, the generator's history window. By default the generator alone keeps the history (*stateless* mode); in *dialogue* mode, one of the held-out strategies, the model sees it too.

Machines and cost. Every run had one GPU to itself: NVIDIA L40S boards, 48 GB, for the compared arms, and RTX A6000 boards, 47.4 GiB, for NRU-S and the cost runs added last, which the dagger in Table 3 marks. The host is 32 cores of an AMD EPYC 7H12 with 503 GiB of memory, PyTorch 2.4.0, CUDA 12.4, cuDNN 9.1.0 and driver 580.159.04 under Python 3.10.14. The study is

Table 3: Cost of one optimiser step per arm, on one idle card at batch size 4 and sequence length 512. *(a)* Grad. rows carry a backward pass, ref. rows are forward passes through the frozen reference, step cost is $(3g + r)F$, and copies counts the model replicas resident. *(b)* Peak GiB is `torch.cuda.max_memory_allocated` and res. GiB what the caching allocator held. Cold is the first epoch and s/epoch the median of the next two, both including refreshes and mining, so $\times$NPO, the ratio of s/epoch to NPO's, exceeds the ratio of step costs. Worst@B is recovery at 32 epochs, for the first-named arm where a row names several. †: measured on an RTX A6000, where NPO takes 12 s per epoch. Constants: batch $n_b = 4$; $k = 2$ replayed and $\bar{\phi} = 1.58$ suite probes per forget row; $K = 4$ simulated steps; a refresh every $m = 25$ steps over $n_a = 12$ items, $|\Pi_{\mathrm{tr}}| = 4$ strategies, $B_{\mathrm{tr}} = 6$ turns and $c = 2$ candidates a turn; $s_E = 26$ steps an epoch.

(a) what one optimiser step computes

arm	grad. rows	ref. rows	step cost	copies
NPO	$2n_b$	$1n_b$	$7n_bF$	2
NPO + anchor	$3n_b$	$1n_b$	$10n_bF$	2
NAU, DRO-NAU + anchor	$4.58n_b$	$2.58n_b$	$16.3n_bF$	2
BalDRO-DV	$4.58n_b$	$2.58n_b$	$16.3n_bF$	2
RULE-NPO	$3n_b$	$1n_b$	$10n_bF$	2
NRU-G fixed pool, same data	$5n_b$	$3n_b$	$18n_bF$	2
NRU-S refreshed suite	$5n_b$	$3n_b$	$18n_bF$	2
NRU-G and its online controls	$5n_b$	$3n_b$	$18n_bF$	3
NPO + anchor + stability	$7n_b$	$3n_b$	$24n_bF$	2
NRU-G + stability	$9n_b$	$5n_b$	$32n_bF$	3

(b) measured cost

arm	peak GiB	res. GiB	cold (s)	s/epoch	$\times$NPO	worst@B
NPO	18.5	19.4	7	6	1$\times$	0.16
NPO + anchor	20.9	21.9	10	9	1$\times$	0.23
NAU, DRO-NAU + anchor	25.8	27.5	23	25	4$\times$	0.16
BalDRO-DV	26.3	28.2	23	23	4$\times$	0.23
RULE-NPO	32.5	33.7	6	162	25$\times$	0.03
NRU-G fixed pool, same data	27.6	30.4	14	14	2$\times$	0.34
NRU-S refreshed suite	21.9	22.8	34	34†	3$\times$	0.16
NRU-G and its online controls	40.4	41.3	145	68	11$\times$	0.25
NPO + anchor + stability	20.9	21.5	21	23	4$\times$	0.06
NRU-G + stability	39.3	39.8	162	84	13$\times$	0.05

327 GPU-hours: 150 of training over 986 runs and 177 of probe evaluation over 191 adapters, plus the six reference and oracle builds the pipeline did not time; the per-run accounting ships with the supplementary material.

The query gate and the split. A generated query is admitted only if it holds 4 to 80 words, is not q or a repeat within the episode, and carries no run of four consecutive protected content words of a (all of them when a has fewer than four); protected words are the content words of a absent from q, so a title q already names is not a reproduction. Entities split by seed-independent hash into generator-train (60%), development (15%) and test (25%), and a reproduction reached through a query duplicating one in the run's own buffer is counted separately, so a memorised probe cannot pass as a held-out one. The same-data control trains on a finished run's final buffer, which fixes the eventual set of probes but not the step at which each became available, so it separates the per-item soft maximum from the probes and not from their order of arrival. The held-out strategies are persona framing, piecewise decomposition, cross-lingual queries, dialogue mode and alternative-template variants of rephrasing and of the follow-up; on RWKU the benchmark's own level-3 probes are scored too, at budget one (Appendix F). The test-entity and held-out-strategy columns were read once, after the method was frozen.

What is fitted and what is searched. Tokenisers are the frozen pretrained ones and no continuous feature is scaled. The answerability mask comes from the reference alone before any arm is trained, the second model family re-gates the same probe text against its own reference, the detector's admission rule was fixed against the reference and the oracle, and τ was chosen from $\{0.3, 1, 3\}$ on development entities, each before a test cell was read. The learning rate, batch size and simulated

step were set once rather than searched, so the only sensitivity claimed is to the temperature. Seeds are 42, 123 and 456; the spread beside every mean is over adapter initialisation, batch order and generator sampling at one fixed entity split.

What the intervals carry. A p value written <0.001 means no bootstrap replicate of five thousand crossed zero, the procedure's resolution floor. A comparison against NPO at the same epochs is paired item by item with the indicator averaged over seeds before the bootstrap, which resamples whole authors (subjects on RWKU) rather than items, so the intervals carry item and author variance and not seed variance; the seed spread is the $\pm$ in the tables and the two are never combined. A two-sided Wilcoxon signed-rank test on the same pairs agrees on every comparison in Section 5, but the questions about one author are not independent observations, so the bootstrap is the one to believe. The four switch contrasts are the one prespecified family and carry a Holm correction; no correction is made across the other arms.

Three asymmetries not controlled for. First, the response-mining and generator arms mine their training data under the disjunction of the two criteria, looser than the reported conjunction, and the judge scoring them is the evaluation's judge, while the augmentation arms train against no detector. This cuts both ways, since at 32 epochs the two arms that train against the evaluation detector score best and worst on it, and it is removable, since the lexical criterion alone passes the gate, shares no model with the generator and preserves the sign and significance of every significant comparison, as does a judge from another model family (Appendix G). Second, the fixed probe suite is generated from seven transformation families, of which paraphrase, back-translation, multi-hop and context-wrapping correspond to the held-out strategies rephrase-alt, cross-lingual, decompose and persona, so the held-out columns favour the augmentation arms; Section 5.2 therefore also reports the suite advantage restricted to the two strategies no family covers. Third, every tunable knob was chosen on a ten-item development set, which Section 7 reports costing us on test.

The search. Each tunable arm received three points on development entities at seed 42 and 16 epochs (Table 4). The default won or tied everywhere except BalDRO-DV, which improves at $\tau = 6$ and is reported at both settings. Ten development items cannot separate $\tau = 1$ from $\tau = 3$, and the frozen choice is the worse on test, -0.006 against $\tau = 3$'s -0.102 $[-0.21, -0.01]$, a limitation of the protocol rather than evidence for $\tau = 3$. The frozen choice is also the worse on test for DRO-NAU's family weights (Table 8). For BalDRO-DV the development search picked $\tau = 6$, which is also the better on test, while the main tables report its published temperature. The soft-maximum temperature was swept for NRU-G alone.

Table 4: The three-point search each tunable arm received: a third of its default, the default and three times it, with development recovery at each (lower is better). A default need not be the published setting. Response mining ran $11/32/96$ samples per prompt and read 0.00 at all three, so the search could not discriminate and the default was kept by default rather than by evidence.

arm	knob	grid	$\frac{1}{3}\times$	default	$3\times$	selected
NRU-G	τ	0.3/1/3	0.60	0.40	0.40	default
BalDRO-DV	τ	0.67/2/6	0.40	0.30	0.20	τ=6
RULE-NPO	samples	11/32/96	0.00	0.00	0.00	default
DRO-NAU + anchor	γ	0.017/0.05/0.15	0.30	0.20	0.40	default

The preliminary study behind the baselines. The operating points of NPO, NPO + anchor, NAU + anchor and DRO-NAU + anchor were fixed in a preliminary study whose quantity was the per-family suppression profile rather than recovery; none of that study is evidence about recovery. Two of its findings carry over. A bidirectional entailment gate removed 61–82% of the compositional candidates against 5–12% removed by answerability, so every objective optimised a neighbourhood dominated by paraphrase; and the relation-level gap of 0.170 clears its null of 0.071 where the standard split's 0.179 does not clear 0.180, which is why the relation-level split carries this study.

F THE RWKU SUBSTRATE

RWKU (Jin et al., 2024) asks a model to forget real people it learned in pretraining, so a question carries leakage only if the base model answers it. Our build, included in the anonymized code in

the supplementary material, applies that gate first, admitting a level-2 (question-answering) question when the base model's greedy answer recovers at least 60% of the answer's content words appearing neither in the question nor in the subject's name; excluding the name is what stops an answer that restates the subject while getting the fact wrong from passing. The forget set is the admitted level-2 questions of the chosen subjects, capped per subject; the anchors are the admitted questions about each subject's own neighbours, RWKU's utility set, so they are other entities rather than, as on TOFU-Rel, the subject's other relations; the retain set is admitted questions about other subjects. Subjects are chosen among those with enough admitted rows and the entity split is taken over subjects. The benchmark's level-3 (adversarial) prompts stay out of training and serve as a held-out probe set released with the benchmark, all nine types, scored with one query per type against the same detector. Every count is recorded with the split in the supplementary material and reproduces by rerunning the build with the same seed.

G THE INSTRUMENT: DETECTOR RULES, THEIR CHECK AND THE HAZARD

The instrument gate. Table 5 recomputes the reference ceiling and the oracle floor under every rule from the same stored turns. Under the judge alone, and under its disjunction with the lexical criterion, the oracle is "recovered" almost as often as the reference, inventing plausible titles and years for authors it never learned while the judge accepts them. Three rules pass the gate: the reported conjunction at recall 0.8, the same conjunction at 0.6 (oracle 0.17, reference 0.94) and the lexical criterion alone at 0.8 (0.08, 0.85). The conjunction at 0.8 is reported because it is the strictest rule, firing only where the lexical criterion also fires, and keeps the judge as a second opinion. Every comparison in Section 5 was recomputed under all three, and every significant contrast keeps its sign and significance: the per-item soft-maximum switch, the refresh switch and the two together read $+0.10/+0.09/+0.11$ under the reported rule, $+0.14/+0.08/+0.13$ at recall 0.6 and $+0.10/+0.08/+0.11$ lexically, while the generator arm's null at 16 epochs stays null, reading -0.01, $+0.06$ and $+0.00$ (recorded in the supplementary material). Every episode runs its full budget, so each rule is recomputed without censoring.

A judge from another model family. The reported judge is the model that writes the probes. Llama-3.1-8B-Instruct re-asked its question, under the same prompt and greedy decoding, of every stored response a passing rule can read: every turn of the reference and the oracle, and every other turn whose lexical recall reaches 0.6, since no conjunction can fire elsewhere. The two judges agree on 0.88 of those $178{,}716$ turns ($\kappa = 0.72$). With the new verdicts the reported rule still passes the gate, the reference at 0.86 and the oracle at 0.07 over 57 answerable items, and the judge alone still fails it (reference 0.99, oracle 0.85). The four matched-utility effects read -0.357, -0.333, -0.263 and -0.088 ($p<0.001$ each), the generator arm's null at 16 epochs stays null (-0.006 $[-0.08, +0.07]$), and it stays worse at 32 ($+0.105$ $[+0.04, +0.17]$). The per-item soft-maximum switch, the refresh switch and the two together read $+0.070$, $+0.082$ and $+0.099$ against NRU-G, each clear of zero after the Holm correction ($p \leq 0.033$). One reading changes: clearing the buffer at every refresh, a null under the reported judge, lowers recovery by 0.094 $[0.02, 0.16]$ ($p = 0.029$ after correction). The per-query answerability flag stays under the original judge, because the reference's response to each issued query is not stored; neither the reported rate nor the gate reads it.

Table 5: The instrument gate under every detector rule, recomputed from the stored turns of the reference and the retraining oracle. Population is the number of the 103 forget items whose direct question the reference reproduces under the rule; the two rates are the fraction of it recovered by the worst of ten strategies within eight queries each. The gate requires the oracle at or below 0.2 and the reference at or above 0.6; three rules pass; the bold one is reported. Recall is the lexical threshold, unused by the judge-only rule; separation is the reference's rate minus the oracle's.

rule	recall	population	reference	oracle	separation
lexical $\vee$ judge	0.6	82	0.98	0.85	0.12
judge only	0.6	73	0.97	0.85	0.12
lexical only	0.6	73	0.97	0.22	0.75
lexical only	0.8	62	0.85	0.08	0.77
lexical $\wedge$ judge	0.6	64	0.94	0.17	0.77
lexical $\wedge$ judge	0.8	59	0.83	0.08	0.75

Agreement, spread and the strategies. The two criteria agree on 0.91 of the 250,141 issued turns ($\kappa = 0.62$), and where they part it is almost always the judge firing alone; an independent entailment model on a stratified sample of 200 turns fires on 0.31 of it against the conjunction's 0.09. That sample ships as a blind sheet, unscored, so a reader can score it, which is inspectability and not adjudication: with no independently labelled sample, neither the detector's sensitivity to paraphrased disclosure nor its error on models other than the oracle is estimated. The minimised likelihood predicts reproduction unevenly across arms (Figure 5), which is why every rate is measured rather than read off a loss. Neither end of the scale is pinned: recovery at 16 epochs spans 0.073 to 0.785 and anchor accuracy 0.30 to 0.69, so a difference has room to show. Nor are the strategies interchangeable: on the reference an alternative-template rephrasing recovers 0.59 of items where dialogue and follow-up reach 0.14, so the reported number is a union.

The rule, in full. Lexical recall is taken over the protected content words of a absent from q and from every query issued so far, and the reported rule requires the answer to keep at least two such words (one on RWKU, whose answers are short) alongside the judge's verdict. A query is *answerable* when the reference reproduces under it.

P4 (most of the budget's yield comes from the first queries). For the reference, the top curve of Figure 4, pooled over items and the training strategies, the per-turn hazard, the chance that a query reproduces the answer given that no earlier query did, falls from 0.165 at the first query to 0.007 or less from the fifth (Figure 4). The fall is pooled over items, so it may come from items differing in how easy they are rather than from a per-item hazard that falls within an episode.

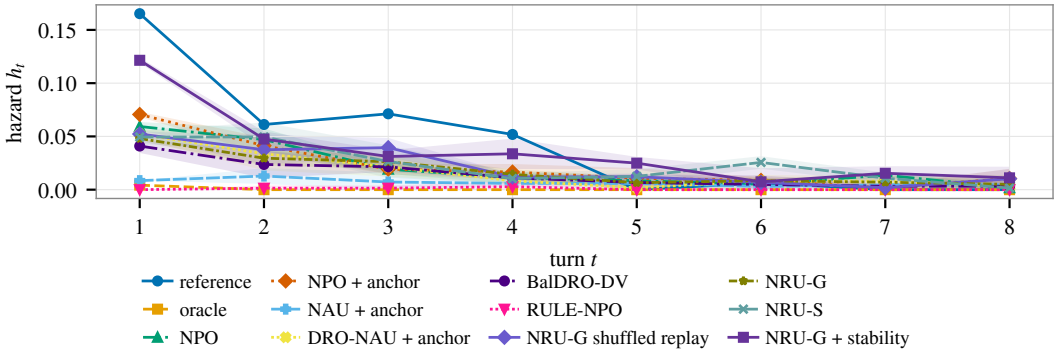

Figure 4: The per-turn hazard h_t, the chance a probe reproduces the answer at turn t given that no earlier probe did, averaged over the four training strategies at 16 epochs, one line per arm; shading: one standard deviation over three seeds.

H Additional results

This appendix holds the full tables behind Section 5: every arm at 16 epochs (Table 6), the two comparisons behind RQ1 and RQ2, the second axis behind RQ3, and the probes and updates the arms produce.

H.1 Matched utility and identical epochs

What each arm costs in utility. Response mining is the strongest single arm on recovery, 0.07 at 16 epochs with the retain set at 0.89, and loses on the anchors, which fall to 0.30, where the generator-built neighbourhood reaches recovery 0.50 with anchors at 0.56. Read against NPO rather than NRU-G, the four switch contrasts of Section 5.3 carry the switch and the neighbourhood together. Those contrasts against NPO read $+0.203$ with both switches dropped, $+0.186$ and $+0.181$ for the per-item soft maximum and the refresh alone, and $+0.034$ for persistence, which is why Table 8 carries both comparators. At 16 epochs none of the four separates under either comparator. Of the two negative controls, the response-denied generator reads -0.277 against the full generator's -0.260 at matched utility, and replaying another item's probes at the same rows, refreshes and

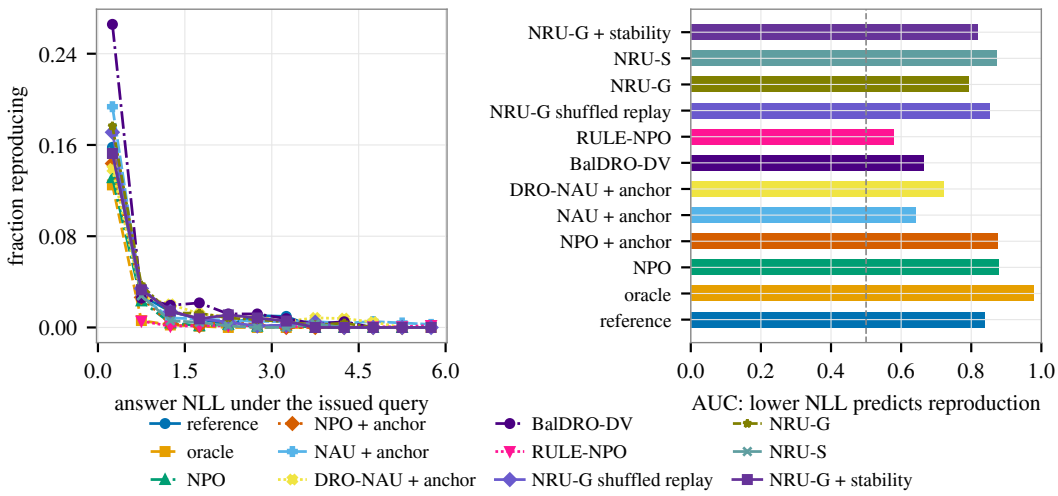

Figure 5: Why a likelihood check is not a disclosure check. *(a)* Fraction of issued queries whose response reproduces the answer against the answer's negative log-likelihood under that query, binned, one line per arm where a bin holds at least twenty turns; a faithful proxy would put every arm on one curve. *(b)* Area under the receiver operating characteristic (ROC) curve for the same arms, 0.5 dashed. The trained arms spread from 0.88 to 0.58, so the same likelihood means different disclosure in different arms.

Table 6: Every arm at 16 epochs on TOFU-Rel. Row blocks: the reference and the oracle, the baselines, the two query-axis NRU arms, NRU-G with one or two components removed, and the two arms trained against a simulated fine-tune. S_0: suppression on the direct question (higher means more suppressed); worst@B: share of answerable items recovered by any of ten strategies within eight queries each (lower is better); held-out and test ent.: the same over the six held-out strategies and over the 15 test-entity items; anchor, retain: generation accuracy (higher is better). Mean ± standard deviation over three seeds; best bold, second underlined. Rows share the number of epochs, not an operating point, so a low recovery can come from anchor damage (Table 1b).

arm	S_0	worst@B	held-out	test ent.	anchor	retain
reference	0.00	0.83	0.71	1.00	0.73	0.81
oracle	0.00	0.08	0.07	0.07	0.75	0.78
NPO	0.37±0.01	0.51±0.02	0.47±0.01	0.49±0.10	0.44±0.00	0.55±0.03
NPO + anchor	0.36±0.01	0.53±0.00	0.46±0.01	0.56±0.04	0.46±0.01	0.56±0.02
NAU + anchor	0.60±0.01	0.29±0.03	0.23±0.01	0.31±0.10	0.42±0.01	0.47±0.02
DRO-NAU + anchor	0.37±0.01	0.44±0.03	0.41±0.04	0.40±0.07	0.55±0.00	0.63±0.02
BalDRO-DV	0.50±0.02	0.42±0.06	0.37±0.06	0.40±0.07	0.59±0.01	0.68±0.02
RULE-NPO	**0.94±0.00**	**0.07±0.03**	**0.05±0.04**	**0.04±0.04**	0.30±0.02	**0.89±0.05**
NRU-G	0.35±0.03	0.50±0.04	0.42±0.06	0.49±0.08	0.56±0.01	0.66±0.03
NRU-S	0.33±0.01	0.55±0.01	0.51±0.02	0.47±0.07	0.51±0.01	0.63±0.02
NRU-G fixed pool	0.64±0.01	0.49±0.04	0.42±0.04	0.49±0.04	0.56±0.02	0.64±0.01
NRU-G uniform replay	0.65±0.01	0.45±0.04	0.38±0.02	0.47±0.07	0.54±0.01	0.66±0.05
NRU-G same data	0.67±0.00	0.50±0.06	0.44±0.05	0.51±0.10	0.58±0.04	0.66±0.06
NRU-G no replay	0.31±0.02	0.57±0.05	0.51±0.03	0.49±0.10	0.55±0.01	0.66±0.01
NRU-G nonadaptive	0.36±0.01	0.51±0.03	0.47±0.04	0.53±0.00	0.56±0.00	0.66±0.02
NRU-G shuffled replay	0.39±0.00	0.49±0.02	0.41±0.02	0.44±0.04	0.53±0.02	0.60±0.01
NRU-G constraint	0.35±0.03	0.52±0.04	0.45±0.04	0.53±0.00	0.55±0.00	0.66±0.01
NPO + anchor + stability	0.13±0.00	0.79±0.01	0.66±0.02	0.80±0.07	**0.69±0.01**	0.77±0.02
NRU-G + stability	0.12±0.00	0.72±0.04	0.63±0.06	0.82±0.10	0.68±0.00	0.79±0.02

generator matches the real neighbourhood at 16 epochs (-0.011, $p = 0.684$) and barely separates at 32 ($+0.040$, $p = 0.054$).

Table 7: The matched-utility comparison behind every effect size in Section 5.2. Each arm is read at the epochs in column two against the NPO run nearest it in anchor accuracy, so most augmentation arms are paired at 32 epochs against NPO at 16, at unequal compute; an arm with no NPO run within 0.04 anchor accuracy, such as BalDRO-DV at its published temperature, does not appear; Table 1b takes each arm at its default setting and at its epochs closest to an NPO run. Anchor is the arm's anchor accuracy; NPO run gives that run's epochs and, in brackets, its anchor accuracy; and Δ is the recovery difference with a 95% interval from a bootstrap over authors. In the grid rows here and in Table 8, τ is a temperature, σ response mining's samples per prompt and γ DRO-NAU's exponentiated-gradient step on its family weights. †: read after continued training.

arm	epochs	anchor	NPO run	Δ recovery [95%]
NPO + anchor	16	0.46	16 (0.44)	$+0.017$ $[-0.025, +0.060]$
NAU + anchor	16	0.42	16 (0.44)	-0.220 $[-0.310, -0.130]$
NAU + anchor	32	0.44	16 (0.44)	-0.345 $[-0.450, -0.239]$
DRO-NAU + anchor	32	0.43	16 (0.44)	-0.328 $[-0.454, -0.209]$
DRO-NAU + anchor γ=0.017	16	0.47	16 (0.44)	-0.186 $[-0.294, -0.075]$
BalDRO-DV τ=0.67	16	0.67	8 (0.70)	-0.390 $[-0.521, -0.250]$
RULE-NPO	16	0.30	32 (0.30)	-0.085 $[-0.136, -0.036]$
RULE-NPO	32	0.28	32 (0.30)	-0.130 $[-0.187, -0.076]$
RULE-NPO σ=11	16	0.33	32 (0.30)	-0.073 $[-0.149, +0.005]$
NRU-G nonadaptive	32	0.44	16 (0.44)	-0.277 $[-0.389, -0.164]$
NRU-G shuffled replay	32	0.43	16 (0.44)	-0.220 $[-0.328, -0.126]$
NRU-G constraint	32	0.41	16 (0.44)	-0.237 $[-0.333, -0.149]$
NRU-G	32	0.43	16 (0.44)	-0.260 $[-0.349, -0.176]$
NPO + anchor + stability	16	0.69	8 (0.70)	-0.062 $[-0.139, +0.022]$
NPO + anchor + stability†	16	0.70	8 (0.70)	-0.073 $[-0.144, +0.000]$
NRU-G + stability	16	0.68	8 (0.70)	-0.130 $[-0.256, -0.022]$
NRU-G + stability†	16	0.70	8 (0.70)	-0.147 $[-0.268, -0.044]$

H.2 The second axis

The step log. The simulation undoes 0.001 nats over the first twenty steps at either number of epochs. Over the last twenty it undoes $+0.02$ nats at 16 epochs, a share of 0.07 of a suppression barely begun, and at 32 epochs 1.10 and 2.09 nats, shares of 0.46 and 0.61, with the forget NLL ending at 2.37 and 3.03. As Appendix B notes, this describes the dynamics of the two-axis update and not the error of its first-order approximation. The fine-tune reopens the neighbourhood by about a tenth for NPO, for NPO with the anchor and for the generator-built neighbourhood alike (Figure 3). The two arms trained for the second axis change by -0.01 and -0.02 at 16 epochs, yet their 0.79 and 0.72 are the rates of models that still answer. At 32 epochs they rise by $+0.26$ and $+0.21$ like everything else while reaching the lowest recovery of any arm but response mining, below the oracle's 0.08, with anchor and retain accuracy near 0.13. The simulated step delays unlearning and then releases it; whether a usable point lies between 16 and 32 epochs is untested.

H.3 Probes and updates

What the generator finds that the suite does not. Over the 416 queries one run mined for 60 items, against the 154 suite queries those items carry, none matches after normalisation and 0.014 match at cosine 0.8 over character four-grams, so the generator is not an expensive way to rediscover the suite.

Turns nothing chose. Table 10 is a sheet of probe turns drawn from the 8080 the generator arm issued. The per-turn reproduction rate over the whole evaluation is 0.019, far below the item-level rate, because an item counts as recovered if any of eighty turns fires. Every issued turn ships with the supplementary material.

The geometry of the update. The *effective rank* of an update is the exponential of the entropy of its normalised singular values. The unlearning adapter stacks on the frozen reference, so its low-

Table 8: Recovery differences at *identical* epochs, with 95% intervals from a bootstrap over authors; negative means fewer facts recovered, and * marks an interval that excludes zero. The first two numeric columns compare each arm with NPO, and the third compares it with NRU-G, the comparator a control needs. Rows are grouped by kind of arm; the grid variants were run at 16 epochs only, and Table 7 gives the matched pairings. At 16 epochs NAU + anchor, BalDRO-DV and RULE-NPO recover less than NPO, and the NRU arms do not; at 32 epochs only RULE-NPO and the two stability arms, whose anchors collapse, recover less than NPO, while NRU-G, BalDRO-DV and NPO + anchor recover more. The matched-utility reading of Table 7, which pairs unequal epochs and compute, is therefore the one supporting Section 5.2.

	Δ recovery against NPO		against NRU-G
arm	16 epochs	32 epochs	32 epochs
Baselines			
NPO + anchor	$+0.02\ [-0.02, +0.06]$	$+0.07\ [+0.02, +0.13]^*$	$-0.02\ [-0.06, +0.04]$
NAU + anchor	$-0.22\ [-0.31, -0.13]^*$	$+0.01\ [-0.05, +0.06]$	$-0.08\ [-0.15, -0.02]^*$
DRO-NAU + anchor	$-0.07\ [-0.17, +0.03]$	$+0.02\ [-0.03, +0.09]$	$-0.07\ [-0.12, -0.01]^*$
BalDRO-DV	$-0.09\ [-0.18, -0.01]^*$	$+0.07\ [+0.02, +0.14]^*$	$-0.02\ [-0.09, +0.06]$
RULE-NPO	$-0.44\ [-0.55, -0.32]^*$	$-0.13\ [-0.19, -0.08]^*$	$-0.22\ [-0.29, -0.14]^*$
NRU arms			
NRU-G	$-0.01\ [-0.10, +0.08]$	$+0.09\ [+0.03, +0.15]^*$	*comparator*
NRU-S	$+0.05\ [-0.03, +0.10]$	$+0.01\ [-0.02, +0.05]$	$-0.08\ [-0.13, -0.04]^*$
NRU-G with components removed (controls)			
uniform replay	$-0.06\ [-0.15, +0.05]$	$+0.19\ [+0.11, +0.26]^*$	$+0.10\ [+0.05, +0.15]^*$
fixed pool	$-0.02\ [-0.12, +0.06]$	$+0.18\ [+0.10, +0.26]^*$	$+0.09\ [+0.04, +0.14]^*$
same data	$-0.01\ [-0.12, +0.09]$	$+0.20\ [+0.14, +0.27]^*$	$+0.11\ [+0.08, +0.15]^*$
no replay	$+0.06\ [-0.03, +0.14]$	$+0.03\ [-0.01, +0.10]$	$-0.06\ [-0.11, +0.01]$
nonadaptive	$+0.00\ [-0.09, +0.08]$	$+0.07\ [+0.01, +0.14]^*$	$-0.02\ [-0.08, +0.05]$
shuffled replay	$-0.02\ [-0.09, +0.05]$	$+0.13\ [+0.06, +0.20]^*$	$+0.04\ [+0.00, +0.08]$
constraint	$+0.01\ [-0.08, +0.09]$	$+0.11\ [+0.06, +0.18]^*$	$+0.02\ [-0.02, +0.08]$
Trained against a simulated fine-tune			
NPO + anchor + stability	$+0.28\ [+0.17, +0.39]^*$	$-0.10\ [-0.13, -0.05]^*$	$-0.19\ [-0.25, -0.12]^*$
NRU-G + stability	$+0.21\ [+0.08, +0.33]^*$	$-0.11\ [-0.16, -0.06]^*$	$-0.20\ [-0.26, -0.14]^*$
Grid variants (the sensitivity analysis), run at 16 epochs only			
DRO-NAU, γ=0.017	$-0.19\ [-0.29, -0.07]^*$	–	–
DRO-NAU, γ=0.15	$-0.05\ [-0.17, +0.06]$	–	–
BalDRO-DV, τ=0.67	$-0.05\ [-0.21, +0.11]$	–	–
BalDRO-DV, τ=6	$-0.15\ [-0.25, -0.06]^*$	–	–
RULE-NPO, σ=11	$-0.42\ [-0.56, -0.28]^*$	–	–
RULE-NPO, σ=96	$-0.47\ [-0.63, -0.31]^*$	–	–
NRU-G, τ=0.3	$+0.05\ [-0.07, +0.15]$	–	–
NRU-G, τ=3	$-0.10\ [-0.21, -0.01]^*$	–	–

Table 9: *(a)* Recovery (worst@B) within eight queries before and after one hundred steps of continued training, with anchor accuracy at both points; only the two stability arms train against a simulated fine-tune. *(b)* What that simulation undoes while it trains: forget NLL at the last optimiser step, the nats the $K = 4$ simulated steps remove averaged over the first and the last twenty steps, and the share undone: the late figure divided by the mean forget NLL of those twenty steps, not by the final-step NLL of the first column. Mean ± standard deviation over three seeds; best bold, second underlined; lower is better for recovery and for Δ, higher for anchor accuracy.

(a) recovery and anchor accuracy before and after continued training

		worst@B			anchor acc.	
arm	epochs	before	after	Δ	before	after
NPO	16	0.51±0.02	0.60±0.05	+0.10	0.44±0.00	0.62±0.01
NPO	32	0.16±0.04	0.49±0.05	+0.33	0.30±0.02	0.51±0.01
NPO + anchor	16	0.53±0.00	0.62±0.03	+0.10	0.46±0.01	0.63±0.01
NPO + anchor	32	0.23±0.06	0.44±0.01	+0.20	0.35±0.01	0.55±0.01
NRU-G	16	0.50±0.04	0.62±0.03	+0.12	0.56±0.01	0.66±0.01
NRU-G	32	0.25±0.06	0.47±0.05	+0.22	0.43±0.03	0.61±0.02
NPO + anchor + stability	16	0.79±0.01	0.77±0.01	-0.01	**0.69±0.01**	**0.70±0.01**
NPO + anchor + stability	32	0.06±0.02	0.32±0.02	+0.26	0.13±0.01	0.38±0.01
NRU-G + stability	16	0.72±0.04	0.70±0.02	**-0.02**	0.68±0.00	0.70±0.01
NRU-G + stability	32	**0.05±0.02**	**0.26±0.01**	+0.21	0.13±0.01	0.39±0.02

(b) what the simulated fine-tune undoes during training

arm	epochs	forget NLL	undone, first 20	undone, last 20	share undone
NPO + anchor + stability	16	0.13±0.00	+0.001±0.000	+0.02±0.00	0.07±0.00
NPO + anchor + stability	32	2.37±0.13	+0.001±0.000	+1.10±0.08	0.46±0.02
NRU-G + stability	16	0.13±0.00	+0.001±0.000	+0.02±0.00	0.07±0.02
NRU-G + stability	32	3.03±0.19	+0.001±0.000	+2.09±0.14	0.61±0.01

Table 10: Four probe turns of NRU-G at 16 epochs: the issued query, the model's response, the strategy and the reported rule's verdict, **yes** marking a reproduction. They are the first four rows of a sheet of eight drawn uniformly at a fixed seed within each verdict, four reproducing and four not, since a plain uniform draw at a per-turn rate of 0.019 would show no reproduction; nothing is edited beyond truncation.

probe	response	strategy	reproduces
It seems the diary format enhances the intimacy of the narrat …	A diary format enhances the narrative's intimacy, allowing readers to explore the i …	dialogue	no
What is the full name of the author associated with Baku, Aze …	The author born on April 13, 1970, in Baku, Azerbaijan, is named Elvin Mammadov, a …	compose	**yes**
As a librarian updating our collection metadata, I came acros …	One influential book from Wei-Jun Chen's bibliography that stands out is "State of …	persona	**yes**
It does seem that any new projects from Wei-Jun Chen will lik …	Environmental science being a rapidly evolving field, there are always new trends a …	dialogue	no

rank product ΔW is the whole move away from the reference, read off the trained adapters. Its effective rank sits between 11.6 and 12.3 of 16 over the query-axis arms and its norm between 0.124 and 0.134, whatever built the neighbourhood, while response mining concentrates its move into 10.3 directions at norm 0.260 and the two stability arms spread theirs over 14.2 and 14.0 at norms 0.113 and 0.108. Each arm contributes one adapter, so the correlation of $r = -0.94$ $[-0.99, -0.88]$ between effective rank and suppression, the rank falling as suppression rises, is an association across eleven arms and not a within-arm one, and is not read as causal.

I FAILURE MODES AND SECOND SUBSTRATES

Section 7 names three failure modes. Each is stated by its *boundary*, the condition separating the regime where the measurement works from the regime where it does not.

1. Answers short enough that any fluent response contains them. On RWKU the forget answers are one or two words, so a response mentioning the protected word in any context is a reproduction under both criteria, and every arm stays at 0.99–1.00 even at suppression 0.98–0.99, which says nothing about the arms and everything about the measurement. *Boundary*: the length of the protected span. Where an answer carries at least three content words absent from the question, as on TOFU-Rel, the conjunction separates reference from oracle by 0.83 against 0.08; where it carries one, no gate we could build separated them.

2. A substrate with no retraining oracle. Every rate here is read against a model fine-tuned on the corpus minus the forget rows, which checks the detector. Facts learned in pretraining admit no such model, so a low recovery rate cannot be distinguished from a detector that never fires. This is a property of the problem and not of the method, and it rules out the setting most people mean by unlearning, removing what a model learned from the open web. *Boundary*: whether the facts are isolable to a training subset that can be withheld and retrained.

3. A configuration with no headroom on a second model family. On Llama-3.2-3B-Instruct the gate passes on its own stack, the reference reproducing on 0.85 of the 65 answerable items against the oracle's 0.08, so the detector gate replicates, but neither P1 nor the comparison of arms can be read on this model. NPO already reaches suppression 0.71 and recovery 0.07 at 16 epochs, no arm separates from another (DRO-NAU + anchor -0.041, $p = 0.18$; NRU-G -0.010, $p = 0.77$), and anchors fall to 0.25–0.38 against a reference at 0.78 (Table 11). *Boundary*: NPO's recovery at the fewest epochs run. Where it already sits at the oracle floor at 16 epochs, as here, a sweep below 16 epochs is needed before any objective can be read, and we did not run one.

Table 11: Llama-3.2-3B-Instruct on TOFU-Rel at 16 epochs, columns and marking as in Table 6, read against this model's own reference and retraining oracle; the uniform-replay row is one seed. RWKU has no such table, having no oracle.

arm	S_0	worst@B	held-out	test ent.	anchor	retain
reference	0.00	0.85	0.82	0.79	0.78	0.82
oracle	0.00	0.08	0.08	0.00	0.74	0.79
NPO	0.71±0.02	0.07±0.01	0.04±0.01	0.10±0.04	0.25±0.04	0.24±0.03
DRO-NAU + anchor	0.74±0.01	0.03±0.01	0.03±0.01	0.02±0.04	0.33±0.01	0.30±0.02
RULE-NPO	**0.87±0.04**	**0.01±0.02**	**0.01±0.02**	0.00±0.00	0.26±0.04	**0.93±0.01**
NRU-G	0.75±0.01	0.06±0.02	0.05±0.02	**0.00±0.00**	0.33±0.02	0.33±0.03
NRU-G uniform replay	0.74	0.06	0.06	0.00	**0.38**	0.39

J PRIOR APPROACHES COMPARED

Section 6 groups the prior work. This appendix adds the details that set each group apart, and Table 12 sets each approach beside this work. Among the evaluations, a budgeted black-box generator recovers even the forgotten *prompts* (Au et al., 2026), and one adversarial suite reaches 72.8–84.3% probe success against high-scoring models (Gupta et al., 2026). Among the methods that enlarge the forget set, RULE (Reisizadeh et al., 2026) and belief unlearning (Li et al., 2026a) add the model's own leaking or high-confidence generations at the *fixed* forget prompts, so they grow the set along the *response* axis; NRU grows it along the *query* axis. Among the weight-space methods, tamper-resistant training (Tamirisa et al., 2025) meta-learns against fine-tuning, meta-unlearning (Gao et al., 2025) makes related concepts self-destruct under adversarial fine-tuning, targeted latent adversarial training (Sheshadri et al., 2025) reports unlearning more robust to fine-tuning, and invariance regularisation (Wang et al., 2025) reaches the same goal without an inner step; all four hold the prompts *fixed*. NRU takes tamper-resistant training's first-order approximation unchanged and scores it over a generated query neighbourhood. In concept erasure the inner adversary is an optimised prompt found without reading any response (Zhang et al., 2024b).

Table 12: What each approach guarantees and assumes. A guarantee is what a method states over a set of queries or weights, checked against each work's method section rather than its abstract; Table 3 gives each arm's step cost. No row, this work included, states a guarantee over queries it did not enumerate.

Approach	**Guarantee it states**	**Operational assumptions**
NPO (Zhang et al., 2024a)	none over any query or weight set	forget answers; a retain set; a reference
Augmentation, uniform or reweighted (Singh et al., 2025; Peleg et al., 2026)	none beyond the enumerated paraphrases	a probe suite built before training
Sample-level DRO (Shao et al., 2026)	a Kullback–Leibler-penalised worst case over samples	the same probe rows; a temperature
Response mining (Reisizadeh et al., 2026; Li et al., 2026a)	none over queries it never issued	a sampling budget; a leak rule
Tamper-resistant training (Tamirisa et al., 2025; Gao et al., 2025; Sheshadri et al., 2025)	none; a first-order meta-objective at fixed prompts	a benign corpus; a step count K
Sharpness-aware and min–max (Fan et al., 2025a; Wu et al., 2025; Yuan et al., 2025; Li et al., 2026b; Cai et al., 2026)	flatness at the returned weights	a radius or an inner problem
Evaluations only (Pan & Wang, 2026; Au et al., 2026; Joshi et al., 2024; Rybak et al., 2026)	—	define the problem, train no unlearning method
This work	none over the reachable sets; the soft maximum is within $\tau \log \lvert Q_q \rvert$ of the worst sampled query (Equation 4), and its link to the measured rate is empirical (P3)	NPO's, plus an anchor set, a benign corpus, a benchmark admitting a retraining oracle, and a generator model for the generator-built arm alone