FORGETTING-MARI: LLM UNLEARNING VIA MARGINAL INFORMATION REGULARIZATION

Anonymous authors
Paper under double-blind review

000

001

002003004

010 011

012

013

014

015

016

017

018

019

021

023

025

026

027 028 029

030 031 032

033

034

035

037

040

041

042

043

044

045

046

047

048

049

051

052

ABSTRACT

As AI models are trained on ever-expanding datasets, the ability to remove the influence of specific data from trained models has become essential for privacy protection and regulatory compliance. Unlearning addresses this challenge by selectively removing parametric knowledge from the trained models without retraining from scratch, which is critical for resource-intensive models such as Large Language Models (LLMs). Existing unlearning methods often degrade model performance by removing more information than necessary when attempting to "forget" specific data. We introduce Forgetting-MarI, an LLM unlearning framework that provably removes only the additional (marginal) information contributed by the data to be unlearned, while preserving the information supported by the data to be retained. By penalizing marginal information, our method yields an explicit upper bound on the unlearn dataset's residual influence in the trained models, providing provable undetectability. Extensive experiments confirm that our approach outperforms current state-of-the-art unlearning methods, delivering reliable forgetting and better preserved general model performance across diverse benchmarks. This advancement represents an important step toward making AI systems more controllable and compliant with privacy and copyright regulations without compromising their effectiveness.*

1 Introduction

As machine learning models, particularly Large Language Models (LLMs), get trained on bigger datasets containing potentially sensitive or regulated information, and as LLMs are increasingly deployed in high-stakes domains, the need to selectively remove specific data influences from these models has become critical. This requirement is driven not only by privacy regulations such as the European Union's General Data Protection Regulation (GDPR) and its "right to be forgotten," but also by practical concerns including the removal of copyrighted content, personally identifiable information, or data determined to be harmful or biased [33, 15, 11, 5, 3]. Unlearning, or removing the *influence* of specific data post hoc, is an attractive tool for achieving this information removal, especially with the high costs of retraining a model from scratch.

Existing unlearning methods often *over-forget*, removing all information linked to the data to unlearn/forget, including knowledge also legitimately supported by the data meant to be preserved. This indiscriminate approach leads to degraded model performance on tasks unrelated to the distinctive information to be forgotten.

To illustrate this distinction, consider a copyright unlearning scenario where we have an LLM pretrained on an article from *The Washington Post* and on one from *The New York Times*, but only the former is legally authorized for use. Both outlets report on an identical event, yet their articles differ in narrative style, phrasing, overlapping facts, and editorial perspective. There are two distinct unlearning objectives with this setup:

• Marginal Information Unlearning: Remove only the stylistic elements, phrasing and content unique to the *Times* article, while retaining shared factual content that also appears in the authorized *Washington Post* article.

^{*}The implementation will be released on GitHub upon acceptance of this manuscript.

• Full Information Unlearning: Erase all content associated with the *Times* article, including factual information that is independently supported by the retained *Washington Post* article.

We argue LLM unlearning naturally targets marginal unlearning, and existing LLM unlearning methods indirectly target it by balancing signals from data to unlearn and data to retain (see Appendix A.1). Indeed, the goal of unlearning is not to eradicate knowledge contained in the unlearn data, but rather to surgically remove only its marginal effect, the information not already supported by the data we are authorized to use. In this copyright scenario, the marginal effect unlearning satisfies legal requirements with minimal utility loss, whereas the full removal would unnecessarily discard information that is lawfully present in the model. This distinction motivates our proposed method, Forgetting-MarI, a direct marginal information[†] removal of the unlearned data.

054

055

056

057

060

061

062

063

064

065

066

067

068

069

071

072

073

074

075

076

077

079

081

083

084

087

880

089

090

091

092

094

095 096

106 107

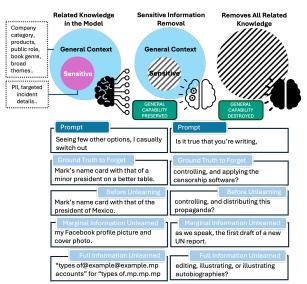


Figure 1: Comparison of sentence completions generated by Llama-3.2-1B models before and after different unlearning methods.

Figure 1 further demonstrates the difference between full-information unlearning and marginal-information unlearning (detailed experimental setup in Section 4.3). We created three models trained on ground truth prompts: one before unlearning, one after marginal information unlearning, and one after full information unlearning. Models are given the first half of a sentence (prompt) and are asked to complete it. Before unlearning, the model completes the sentences in a way that is similar to the ground truth. With marginal unlearning, the model produces different but coherent completions. With full unlearning, the model struggles to coherently complete the sentences.

1.1 OPEN CHALLENGES IN LLM UNLEARNING

Effective LLM unlearning must balance three objectives [25]. First, unlearn efficacy measures how well a model suppresses the influence of the data we want to unlearn, called the unlearn set \mathcal{D}_u . Second, utility preservation ensures the model's ability to retain performance on general tasks, and the data we are still authorized to use, called the retain set, \mathcal{D}_r . Finally, computational cost encompasses the time, memory, and carbon used during unlearning. All unlearning techniques aim to optimize these three objectives, which inherently come with tradeoffs; what differs is where and how the model parameters are updated, directly affecting their ability to balance the three. A breakdown of existing techniques and their strengths, weaknesses is shown in Table 1, with their technical details and commonality in indirect marginal unlearning in Appendix A.1.

Table 1: Comparison of LLM Unlearning Approaches

Approach	Key Methods	Strengths	Limitations
Full tuning	Loss-Rev. ¹ , GradDiff ² , KL-Dist. ³ , DPO ⁴ /NPO ⁵	Fine-grained control; preserves alignment	Risk of over-forgetting; utility drop
Weight edit	ROME ⁷ , MEMIT ⁸ , AlphaEdit ⁹ , DEPN ¹⁰	Fast; memory-light; layer-local patches	Less effective on distributed/stylistic info
Counterfactuals	IDK ¹¹ , EntityAnon. ¹² , ULMR ¹³ , SKU ¹⁴	No weight edits; easy deployment	Prompt-dependent; leakage if coverage is thin
Adaptation	Task-vec15, O316	Modular patches; base model unchanged	Adapters grow linearly; gate failures leak

[†]Numbers map to BibTeX entries: ${}^{1}[44]$, ${}^{2}[25]$, ${}^{3}[17]$, ${}^{4}[35]$, ${}^{5}[45]$, ${}^{7}[30]$, ${}^{8}[31]$, ${}^{9}[10]$, ${}^{10}[41]$, ${}^{11}[20]$, ${}^{12}[9]$, ${}^{13}[38]$, ${}^{14}[26]$, ${}^{15}[19]$, ${}^{16}[12]$.

[†]The term Marginal Information is formalized in Definition 1.1.

Despite rapid progress, LLM unlearning is still an emerging discipline with several open challenges, summarized in Table 2.

Robust unlearning & Utility Preservation: Existing LLM unlearning techniques via full-parameter fine-tuning typically treat the unlearn set \mathcal{D}_u as fully toxic, forcing the model to forget every sequence in \mathcal{D}_u regardless of overlap with the retain set \mathcal{D}_r . Examples include loss-reversal [24], gradient-difference [44], KL-ascent [17], and preference-based DPO/NPO [35, 45]. Furthermore, Even the local editors aim at precise edits (ROME, MEMIT) share this limitation [30, 31], erasing shared facts and stylistic cues and raising perplexity on \mathcal{D}_r and held-out tasks. Benchmarks (RWKU, MUSE, Eight-Method) consistently report sizable utility drops after unlearning [21, 46, 27].

Stable Continual Unlearning: As the legal landscape around data usage changes, a deployed LLM may receive hundreds or thousands of unlearn requests. Production-ready unlearning, therefore, needs to be able to repeatedly unlearn, retain utility, and keep computation and memory within a practical range. Exact methods like full retraining or shared SISA guarantee unlearning but their cost scales with both model size and request count [2, 14,

Method	Utility	Scalable	Formal
Method	Preservation	Continual	Guarantees
Full Parameter	☆	✓	X
Weight Editing	☆	✓	X
Counterfactual	×	✓	X
Adaptation	☆	☆	X
Ours	✓	✓	✓

Table 2: Comparison of families of unlearning methods based on literature evidence. Our proposed marginal effect unlearning addresses key limitations of existing approaches. (✓=yes, ✗=no, ☆=partial)

1]. Lighter updates like influence functions [16] or repeated ROME/MEMIT edits [30, 31] are cheap per removal yet accumulate inference costs and utility drift. Task-vector subtraction or adapter stacks save compute during unlearning but require storing external model adapters, also creating downstream inference costs [19, 12, 40]. Thus, continually unlearning without runaway resources or utility loss remains unsolved.

Formal Guarantees at LLM-Scale: Certified unlearning is well established for linear/kernel models [16], high-dimensional classifiers [6], and general mathematical formulation of machine unlearning [43]. However, no existing method provides guarantees that scale to autoregressive transformers with billions of parameters (7B–70B+), such as GPT or Llama. As a result, practitioners lack *reliable guarantees* of the extent to which the unlearn set remains uninferable or undetectable after common downstream operations such as compression, distillation, or adversarial probing [25].

1.2 OUR CONTRIBUTIONS

To address these challenges, we introduce *Forgetting-MarI*, a novel information-theoretic LLM unlearning framework. First, we provide a heuristic definition of marginal information (formal quantification appears in Section 2.1):

Definition 1.1 (Marginal Information (MarI)). *Marginal information is the marginal effect on model inference when adding the unlearn set to the retain set.*

The core idea of Forgetting-MarI is to penalize the model in proportion to the marginal information, and thus eliminate only the *unique* contribution of the unlearn dataset on the model's parameters and its inference abilities. This avoids erasure of shared information between the retain and unlearn sets. A key piece of our technique, therefore, is an accurate quantification of marginal information, which we detail in Section 2.1. Forgetting-MarI can be summarized by the following learning objective:

$$\min_{\substack{\text{model parameter: } \theta}} \, \ell_{\text{utility}}(\text{model}(\theta), \mathcal{D}_r) + \ell_{\text{MarI}}(\text{model}(\theta), \mathcal{D}_r, \mathcal{D}_u),$$

with $\ell_{utility}$ being a loss that aims to maintain the utility of the model and ℓ_{MarI} being the marginal information loss derived from an accurate marginal information quantification.

The key contributions of our proposed method include:

- (A1) *Utility preservation*: Targeting marginal information means that only the marginal effect of the unlearn set is removed, preserving information shared with the retain set.
- (A2) Scalable and continual: Using an additive mutual-information regularizer integrates with standard gradient-based fine-tuning and naturally supports continual unlearning.
- (A3) Theoretical unlearning guarantee: Bounding marginal information yields an explicit upper bound on residual mutual information, providing provable undetectability of the unlearn set.

• (A4) Exemplary experimental performance: Experiments show that our proposed method outperforms state-of-the-art unlearning methods in a wide range of unlearning tasks.

2 UNLEARNING: MARGINAL INFORMATION

Forgetting-MarI relies on a novel quantification of *marginal information* that (i) vanishes when the unlearn set \mathcal{D}_u adds *no new* information beyond the retain set \mathcal{D}_r , and (ii) increases as \mathcal{D}_u contributes information absent from \mathcal{D}_r , recovering the full information in \mathcal{D}_u as \mathcal{D}_r vanishes. We propose a mutual information (MI)-based quantification that satisfies these properties.

2.1 QUANTIFYING AND UNLEARNING MARGINAL EFFECTS

Fix a language model p_{θ} (with parameter θ) over a finite vocabulary V and a length $T \geq 1$. For $y \in V^T$, let $p_{\theta}(\cdot \mid y_{< t})$ be the next-token distribution. For a subset $s \subseteq V^T$, let μ_s be the uniform law on s and define its averaged next-token marginals $(p_{\theta})_t^s(v) := \mathbb{E}_{Y \sim \mu_s} \left[p_{\theta}(v \mid Y_{< t}) \right]$ for $t \in [T], v \in V$. Write $p^r := \{(p_{\theta})_t^r\}_{t \in [T]}, \ p^u := \{(p_{\theta})_t^u\}_{t \in [T]}.$ For $d := r \cup u, \ p_t^d = \alpha \ p_t^r + (1 - \alpha) \ p_t^u, \ \alpha := \frac{|r|}{|r| + |u|} \in (0, 1)$. Let $T^* \sim \text{Uniform}([T])$ and $T^* = T^*$ and set $T^* = T^*$. Then the mutual information between T^* and T^* is defined as

$$I(X_{\text{MarI}}; Z) := \frac{1}{T} \sum_{t=1}^{T} \text{JSD}(p_t^d, p_t^r).$$

$$\tag{1}$$

Here, we denote the Jensen-Shannon divergence as $JSD(p,q) := \frac{1}{2}D_{KL}(p||m) + \frac{1}{2}D_{KL}(q||m)$ with $m := \frac{p+q}{2}$ and $D_{KL}(p||q) := \sum_v p(v) \log \frac{p(v)}{q(v)}$. By construction, the information or distribution represented by d can be decomposed into the contribution of $r \cap d = r$ and the marginal contribution of $d \setminus r = u$. The distribution contributed by r through the model p_θ is p^r . The distributional contribution from the addition of u through p_θ is the distributional difference between p^r and p^d .

By construction, the quantification of the marginal effect is small if p^r is close to p^d , because such proximity suggests that the information content in u has already been largely represented by r. Conversely, the quantification will be large if p^r differs significantly from p^d , indicating that u contributes substantial new information w.r.t. p_θ and induces a model output distribution shift. Therefore, defining this marginal effect quantification boils down to differentiating p^r from p^d for any $r \in \mathcal{D}_r$ and $d = r \cup u$ with arbitrary $u \in \mathcal{D}_u$.

A natural way to *quantify* this difference is via a binary detection problem. Consider a binary detection problem using the construction above:

$$X_t := X \big|_{T^* = t} \sim \begin{cases} p_t^d, & Z = 0, \\ p_t^r, & Z = 1, \end{cases} \quad \mathbb{P}[Z = 0] = \mathbb{P}[Z = 1] = \frac{1}{2}.$$
 (2)

If $p^r = p^d$, even an optimal classifier does no better than flip a coin. If there is distributional shift, it can detect the difference. A sharp information—theoretic upper bound on the Bayes accuracy, denoted by $P_{\rm acc}$ and defined below in Proposition 2.1, is the following:

Proposition 2.1 (Detection accuracy upper bounded by mutual information). For (X_{MarI}, Z) with prior $\pi = \mathbb{P}[Z=1]$,

$$P_{\rm acc} \ = \ \mathbb{E}\left[\max\{P(Z=0\,|\,X_{\rm MarI}),\, P(Z=1\,|\,X_{\rm MarI})\}\right] \ \leq \ 1 - H_2^{-1}\big(H_2(\pi) - I(X_{\rm MarI};Z)\big),$$

where $H_2(\cdot)$ is the binary entropy and H_2^{-1} denotes the inverse of H_2 restricted to $[0, \frac{1}{2}]$.

Proof in Appendix B.1. Here $P(Z \mid X_{\mathrm{MarI}})$ denotes the Bayes-optimal posterior between retain r and union d. Note that $I(X_{\mathrm{MarI}}; Z) \in [0, H_2(\pi)]$ satisfies: (1) $I(X_{\mathrm{MarI}}; Z) = 0$ when $p^d = p^r$; (2) $I(X_{\mathrm{MarI}}; Z)$ grows with their divergence, approaching $H_2(\pi)$ as \mathcal{D}_r vanishes. Since a gap $p^d \neq p^r$ occurs precisely when u induces model confidence shifts not explained by r, Proposition 2.1 gives $I(X_{\mathrm{MarI}}; Z)$ an intuitive meaning as the detectability of the marginal effect (Definition 1.1).

Definition 2.1 (MI-based marginal information loss). With (X_{MarI}, Z) as in equation 2, define

$$\ell_{\text{MarI}}(\theta, r, u) := I(X_{\text{MarI}}; Z).$$

Thus, Forgetting-MarI solves

$$\min_{\theta} \ell_{KL}(\theta, r) + \ell_{MarI}(\theta, r, u), \tag{3}$$

where $\ell_{KL}(\theta,r) := D_{KL}(p^r(\theta) \| p^r(\theta_0))$ is the KL divergence between the updated model (parameter θ) and the frozen original model (parameter θ_0) on r, and ℓ_{MarI} is as above, motivated by Prop. 2.1. Algorithm 1 describes an efficient LLM implementation.

Remark 2.1 (Alternative quantification). Marginal information measures the shift from $p_{\theta}(r)$ to $p_{\theta}(d)$. Alternatively, one may use $\ell'_{\text{MarI}}(\theta, r, u) := D_{\text{KL}}(p^d \parallel p^r)$ or $D_{\text{KL}}(p^d \parallel p^r(\theta_0))$. But mutual information has the advantage of (1) stability (boundedness), (2) interpretability (Proposition 2.1), and (3) continuous unlearning (evolving reference $m = \frac{p+q}{2}$). See Appendix B.2 for details.

2.2 Marginal Information & Perplexity-Based Detectors

We provide theoretical guarantees for the unlearning performance of Forgetting-MarI against white-box copyright detectors that rely on model confidence (perplexity / cross-entropy). Let

$$S_{\theta}(x,y) = \frac{1}{T} \sum_{t=1}^{T} \left(-\log p_{\theta}(x_t \mid y_{< t}) \right)$$

be the standard cross-entropy (per-token negative log-likelihood). State-of-the-art detectors [3, 47, 29] flag the membership of x in training by testing whether $S_{\theta}(x,x)$ is suspiciously low. We adopt the notation from Section 2.1: sequences $r,u\in V^T$, next-token marginals p_t^r,p_t^u , their mixture $p_t^d=\alpha\,p_t^r+(1-\alpha)\,p_t^u$ with $\alpha=\frac{|r|}{|r|+|u|}$, and the mutual information $I(X_{\text{MarI}};Z)$.

The next result shows that, given a set of sequences to forget, denoted by u, Forgetting-MarI guarantees that there is a set of sequences in the retain set, denoted by r, such that the score $S_{\theta}(u,u)$ is close to $S_{\theta}(u,r)$. In other words, a high model confidence implied by $S_{\theta}(u,u)$ is possibly due to the existence of r because one would get the same score for u if feeding the model r instead of u.

Theorem 2.1 (MarI controls the self-perplexity gap). Fix $u = (u_1, \dots, u_T) \in V^T$ and assume the pathwise non-vanishing condition $\min\{p_t^u(u_t), p_t^r(u_t)\} \ge \gamma \in (0, 1]$ for all t. Then

$$\left| S_{\theta}(u, u) - S_{\theta}(u, r) \right| \leq \frac{2\sqrt{2}}{\gamma(1 - \alpha)} \sqrt{I(X_{\text{MarI}}; Z)}.$$

See Appendix B.3 for the proof. In particular, when $I(X_{\text{MarI}}; Z)$ goes to zero, the score gap above vanishes, and the perplexity/log-likelihood detectors lose discriminative power after unlearning.

Finally, we show that MarI directly controls the score gap even for a neighborhood of u rather than only u itself. Proof can be found in Appendix B.4.

Theorem 2.2 (MarI controls neighborhood-perplexity gap). Draw $U := \{U_t\}_{t=1}^T$ with $U_t \sim p_t^u$ independently across $t \in [T]$ and suppose $\max_{t,x} \left\{ \frac{p_t^u(x)}{p_t^r(x)} \vee \frac{p_t^r(x)}{p_t^u(x)} \right\} \leq M$. Let $C := \max_{t,x: p_t^u(x) > 0} \left[\log \frac{p_t^r(x)}{p_t^u(x)} \right]^2 < \infty$. Then, for any $\varepsilon > 0$, with probability at least $1 - 2 \exp(-T\varepsilon^2/(2C))$,

$$\left| S_{\theta}(U, u) - S_{\theta}(U, r) \right| \le \left(\log M \right) \frac{M}{M - 1} \frac{\sqrt{2}}{1 - \alpha} \sqrt{I(X_{\text{MarI}}; Z)} + \varepsilon.$$

3 ALGORITHM DESIGN

Algorithm 1 presents a pseudo-code for Forgetting-MarI. In theory, the *token-wise* MarI (Equation 1) provides strong guarantees when sentences in r and u are *homogeneous* in length

271

272

273

274

275

276

277

278

279

281

282

283

284

286

289

290

291

293

295

296

297

298

299

300

301 302

303

304

305

306

307

308

310

311

312

313

314

315

316

317

318

319

320

321 322 323 and token-wise context. In practice, however, token-wise MarI Loss can be noisy under heterogeneous batches. To address this, we also provide a *pooled* ("flattened") estimator that first averages across token positions (and batch) to form $\bar{p}^s = \frac{1}{T} \sum_t p_t^s$, $s \in \{r, u, d\}$, then computes the pooled MarI Loss $I(\bar{X}_{\text{MarI}}; Z) = \text{JSD}\left(\bar{p}^d, \bar{p}^r\right)$. By the data-processing inequality, $I(\bar{X}_{\text{MarI}}; Z)$ is a *variational lower bound* to the token-wise MarI, $I(X_{\text{MarI}}; Z)$. Such a pooled version aims to stabilize marginal information quantification by filtering the position-heterogeneous noise and emphasizing the dominant distribution shift.

Algorithm 1: Forgetting-MarI.

```
Algorithm 1: Forgetting-MarI
      Require: Retain dataset \mathcal{D}_r = \{x_i^r\}_{i=1}^{|\mathcal{D}_r|}, unlearn dataset \mathcal{D}_u = \{x_j^u\}_{j=1}^{|\mathcal{D}_u|}; pretrained model f_{\theta_0}; learning rate \eta; batch sizes B_r, B_u; epochs E; trade-off \lambda \in (0, 1);
               hetero \in \{False, True\}.
      Ensure: Unlearned parameters \theta_E.

 Initialize θ ← θ<sub>0</sub>; build dataloaders L<sub>r</sub>, L<sub>u</sub>.

        2: for e = 1 to E do
                   Shuffle \mathcal{L}_r, \mathcal{L}_u; S \leftarrow \min\{|\mathcal{L}_r|, |\mathcal{L}_u|\}
                   for s = 1 to S do
                        Minibatches: r \leftarrow \mathcal{L}_r(s), u \leftarrow \mathcal{L}_u(s).
                         \begin{array}{l} \textbf{Model logits \& probabilities:} \\ L^r \leftarrow f_{\theta}(r) \in \mathbb{R}^{B_r \times T \times |V|}, \ p^r_{b,t} \leftarrow \text{softmax}(L^r[b,t,:]) \in [0,1]^{B_r \times T \times |V|}; \end{array} 
                              L_0^r \leftarrow f_{\theta_0}(r) \in \mathbb{R}^{B_r \times T \times |V|}, \ p_{b,t}^r(\theta_0) \leftarrow \operatorname{softmax}(L_0^r[b,t,:]) \in [0,1]^{B_r \times T \times |V|};
                              L^u \leftarrow f_{\theta}(u) \in \mathbb{R}^{B_u \times T \times |V|}, \ p_{b,t}^u \leftarrow \operatorname{softmax}(L^u[b,t,:]) \in [0,1]^{B_u \times T \times |V|};
        9:
                                p^d \leftarrow \text{concat}_{\text{batch}}([p^r, p^u]) \in [0, 1]^{(B_r + B_u) \times T \times |V|}.
      11:
                         Utility Loss:
      12:
                              If hetero=True
                                 \bar{p}^r \leftarrow \tfrac{1}{B_r*T} \textstyle\sum_{b,t} p^r_{b,t}, \, \bar{p}^r_0 \leftarrow \tfrac{1}{B_r*T} \textstyle\sum_{b,t} p^r_{b,t}(\theta_0),
      13:
                                  \ell_{\text{KL}} \leftarrow \sum_{v \in V} \bar{p}^r(v) \log(\frac{\bar{p}^r(v)}{\bar{p}_n^r(v)})
      14:
                               Else:
                                  p_t^r \leftarrow \tfrac{1}{B_r} \sum_{b=1}^{B_r} p_{b,t}^r, \, p_t^r(\theta_0) \leftarrow \tfrac{1}{N_r} \sum_{b=1}^{B_r} p_{b,t}^r(\theta_0),
      17:
                                  \ell_{\text{KL}} \leftarrow \frac{1}{T} \sum_{t} \left( \sum_{v \in V} p_{t}^{r}(v) \log(\frac{p_{t}^{r}(v)}{p_{t}^{r}(\theta_{0})(v)}) \right)
                        MarI Loss:
      18:
                              If hetero=True:
      19:
                                 P(v,0) \leftarrow \frac{1}{2(B_r + B_u) * T} \sum_{b,t} p_{b,t}^d, \ P(v,1) \leftarrow \frac{1}{2B_r * T} \sum_{b,t} p_{b,t}^r, \ P(v) = \sum_z P(v,z),
      21:
                                 \ell_{\text{MarI}} \leftarrow \sum_{z \in \{0,1\}} \sum_{v \in V} P(v, z) \log \frac{P(v, z)}{P(v) P(z)}
      22:
                                  P_t(v,0) \leftarrow \frac{1}{2(B_r + B_u)} \sum_b \, p_{b,t}^d, \, P_t(v,1) \leftarrow \frac{1}{2B_r} \sum_b \, p_{b,t}^r, \, P_t(v) = \sum_z P_t(v,z),
      23:
                        \ell_{\text{MarI}} \leftarrow \frac{1}{T} \sum_{t} \left( \sum_{z \in \{0,1\}} \sum_{v \in V} P_t(v,z) \log \frac{P_t(v,z)}{P_t(v)P_t(z)} \right).
Total loss: \ell_{\text{total}} = (1 - \lambda) \ell_{\text{KL}} + \lambda \ell_{\text{MarI}}.
      24:
      25:
                         Update: \theta \leftarrow \theta - \eta \nabla_{\theta} \ell_{total}.
                  end for
      27.
      28: end for
      29: return \theta_E \leftarrow \theta
```

In practice, one should use hetero=False when r and u are homogeneous (similar lengths and aligned token-wise contexts), to leverage the more accurate tokenwise signal. Use hetero=True for large dataset with random batching and unaligned sequences, to obtain a more stable pooled marginal information signal.

A full theoretical/empirical comparison (and mixtures of both signals) is deferred to future work. In our experiments, $I(X_{\mathrm{MarI}};Z)$ and pooled $I(\bar{X}_{\mathrm{MarI}};Z)$ achieve comparable unlearning under different hyperparameters. Therefore, Section 4 reports results under one label. We include a head-to-head comparison between the two approaches of marginal information in Appendix C.

4 EXPERIMENTS

We evaluated Forgetting-MarI across the GPT and Llama family models on two real-world datasets: *Harry Potter and the Prisoner of Azkaban* [22] and *Careless People: A Cautionary Tale of Power, Greed, and Lost Idealism* [42], chosen for contrasting genre and pretraining prevalence (Harry Potter is abundant in pretraining; *Careless People* is likely scarce).[‡] As a full-parameter unlearning approach, we benchmark against full-parameter fine-tuning baselines: gradient ascent (GA) [44], KL-gradient ascent (KL-GA) [17], gradient difference (GD) [24], and direct preference optimization (DPO) [35]. Table 3 summarizes these baselines alongside other unlearning approaches. Detailed discussion of the latter is deferred to Appendix A.1.

Our numerical results demonstrate four key advantages: (1) Balanced unlearning: after hyperparameter tuning for all methods, Forgetting-MarI best matches the retain—unlearn trade-off, preserving performance on \mathcal{D}_r while removing \mathcal{D}_u ; (2) Stable continual unlearning: in terms of both regularization parameter selection and model performance during training, Forgetting-MarI is the most stable; (3) General capacity: models unlearned with Forgetting-MarI retain general model capabilities; (4) Theoretical guarantee: perplexity/confidence-based detectors can no longer detect unlearning data following unlearning via Forgetting-MarI, as guaranteed by Theorems 2.2 and 2.1.

Algorithm	Utility Preservation	Unlearning
Gradient Ascent	0	GA
Gradient Difference	GDe	GA
KL-Gradient Ascent	KL-divergence	GA
DPO	GDe	Preference
Forgetting MarI	KL-divergence	MarI

GA = Gradient Ascent, GDe = Gradient Descent

Table 3: The table summarizes the utility and unlearning types of methods included in the comparison experiments.

[‡]All experiments ran on a node with four NVIDIA A100 GPUs (80 GB HBM2 each).

4.1 EXPERIMENTAL SETUP

To simulate real-world scenarios where forget and retain content is semantically intertwined, we construct the unlearn set \mathcal{D}_u and retain set \mathcal{D}_r by selecting *alternating sentences* from each dataset. We use a separate validation set to monitor utility. For *Careless People*, we additionally use Reddit stories [36]. Experiments share the following protocol: (i) fine-tune on $\mathcal{D}_u \cup \mathcal{D}_r$ to obtain a baseline; (ii) apply each unlearning method to remove the influence of \mathcal{D}_u ; (iii) train a *gold-standard* unlearn baseline on \mathcal{D}_r only, simulating a model that never saw \mathcal{D}_u . An optimally unlearned model should match the performance of this gold-standard. Model performance on the unlearn, retain, and validation sets were assessed using next-token prediction accuracy. General model performance was assessed using benchmarks from Eleuther's LM Evaluation Harness [13].

4.2 Forgetting Harry Potter

Setting We fine-tune a pretrained GPT-2 Large [34] model on excerpts of *Harry Potter and the Prisoner of Azkaban* [22] following a similar approach of Eldan et al. [9]. We curated equally-sized (50/50) retain and unlearn datasets with high contextual overlap to maximize their correlation.

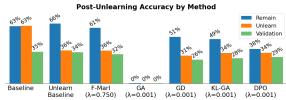


Figure 2: Next-token prediction accuracies for different unlearning methods compared to the baselines.

Utility Preservation Figure 2 reports the *best* accuracies achieved for each unlearning method after tuning their regularization parameters (e.g., with the best λ marked below the methods). The reported accuracies correspond to the stopping criterion when validation accuracy fell by 3% or more. Forgetting-MarI shows the best balance of unlearning and utility preservation, as it best matches the retain-unlearn trade-off of the gold-standard unlearn baseline.

Stability of continual unlearning Figure 3 highlights two notions of robustness:

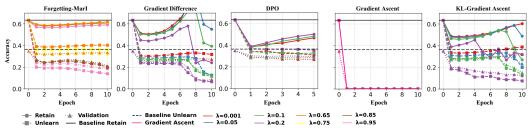


Figure 3: Training curves for each method with varying choices of the regularization parameter λ . Forgetting-MarI exhibits smooth monotone behavior, while the other methods show oscillation or utility collapse.

(1) Training robustness over epochs. Forgetting-MarI descends steadily to its optimum. GA and GD overshoot and bounce, KL-GA diverges after 5-6 epochs, DPO plateaus prematurely.

(2) Robustness against regularization tuning. Forgetting-MarI shows a monotone and smooth utility-unlearning trade-off when adjusting the regularization parameter. In comparison, GD and KL-GA display unstable oscillations with different choices of λ .

Both training and regularization robustness are necessary for a practical use of unlearning techniques. Practitioners do not have access to ground truth baselines and have limited time to select/determine the best parameter or training epoch to stop at, so stability is essential. Forgetting-MarI is the most stable technique during unlearning, mak-

(WikiText)	GF 12-LG	Dasenne	Baseline	r-iviari
bits/byte ↑	0.841	0.925	0.905	0.910
byte-pplx ↓	1.792	1.898	1.873	1.879
word-pplx ↓	22.61	30.80	28.66	29.19

Subset	GPT2-LG	Baseline	Unlearn Baseline	F-MarI
All (MMLU) ↑	0.232	0.232	0.236	0.230
Humanities ↑	0.246	0.243	0.247	0.243
Other ↑	0.242	0.242	0.242	0.241
Social Sci. ↑	0.217	0.222	0.221	0.215
STEM ↑	0.217	0.217	0.227	0.216

Table 5: Performance on the MMLU hierarchical evaluation.

ing it the safest choice in practice. We refer the reader to Appendix D Figure 10 for a smaller learning rate training curve, which further demonstrates the stability of continual unlearning.

Preserved general model capacity To assess the general model capacity after unlearning, we compare the unlearned model with the baseline, the unlearn baseline, and GPT2-LG (the model before finetuning on the retain set) on the following benchmarks: WikiText and MMLU [13]. We repeated the experiment with a realistic 10/90 unlearn/retain split to simulate practical scenarios where only

a small portion of the data requires removal. The training curves for this setup is in Figure 9. Forgetting-MarI exhibits small declines on a few capability metrics (Tables 4 and 5) compared with the baseline or the unlearn baseline, while matching them on *most* tested metrics. Overall, the method largely preserves general model capacity on standard evaluations. Additional results on PIQA, ARC, and HellaSwag are reported in Table 7 (Appendix D).

4.3 FORGETTING CARELESS PEOPLE

Setting We fine-tune Meta's Llama-3.2-1B [32] on Careless People: A Cautionary Tale of Power, Greed, and Lost Idealism [42], following [28]. Beyond the setup in Sec. 4.1, we also test an uncorrelated scenario: use 2025 Reddit stories as \mathcal{D}_u and the other half of Careless People as validation. Both datasets were published after the Llama model's release, ensuring \mathcal{D}_u is not in the model pre-training.

Utility preservation Figure 4 shows the best results for each unlearning method on both experiments. Consistent with results in Sec. 4.2, Forgetting-MarI most closely matches the unlearn baseline, achieving

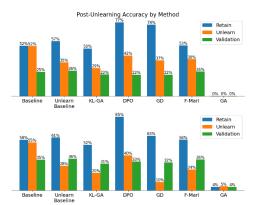


Figure 4: Next-token prediction accuracy for correlated (top) and uncorrelated (bottom) test sets.

most closely matches the unlearn baseline, achieving comparable unlearn accuracy while maintaining a similar retain and validation accuracy in both settings:

(i) Correlated experiment: GD and DPO both overtrain on \mathcal{D}_r , while KL-GA struggles to remove the information in \mathcal{D}_u while maintaining performance on \mathcal{D}_r . Moreover, all other methods show a gradual loss of performance on the test set. (ii) Uncorrelated experiment: With less informational overlap between \mathcal{D}_u and \mathcal{D}_r , other unlearning methods more easily remove \mathcal{D}_u . However, GD and DPO still overtrain on \mathcal{D}_r while either over- or undershooting the desired unlearning. KL-GA quickly overshoots the desired amount of unlearning.

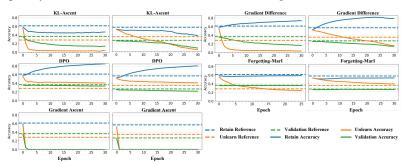


Figure 5: Next-token prediction accuracy during training for both experiments. Epoch 0 represents the model before unlearning. Horizontal lines represent the "gold standard" unlearn baseline model, the model trained only on \mathcal{D}_r . Left/right columns for each method show correlated/uncorrelated test results respectively.

Stability of continual unlearning Figure 5 shows the unlearning performance of each method over course of unlearning, where the curves for each method correspond to the experiment with the best performing regularization parameter for each method.

Forgetting-MarI is the

best at smoothly approximating the unlearn baseline. The methods based on gradient ascent, GA, GD, and KL-GA, all over-penalize \mathcal{D}_u due to the utility-destroying nature of gradient ascent. DPO, meanwhile, never matches the unlearn baseline in accuracy on \mathcal{D}_u and over-trains on \mathcal{D}_r . Across both sets of experiments, Forgetting-MarI minimally drops the validation accuracy, where the validation curve remains largely unchanged.

We note that there is the possibility that, in theory, one could find a perfect balance between gradient ascent and utility regularization, leading to a stable balance between unlearning and utility preservation, using one of the other methods. However, such a balance seems practically unattainable due to the unlearning instability over time and the lack of monotonicity in the choice of λ for methods based on gradient ascent.

Metric	Baseline	Unlearn Baseline	F-MarI
	Cor	related	
ARC-Easy, acc	0.594	0.599	0.582
PIQA, acc	0.713	0.706	0.706
	Unce	orrelated	
ARC-Easy, acc	0.595	0.628	0.556
PIQA, acc	0.709	0.726	0.700

Table 6: Evaluation results from the ARC-Easy and PIQA benchmarks for both the correlated and uncorrelated experiments.

Preserved general model capacity Table 6 shows the performance of Llama-3.2-1B unlearned with Forgetting-MarI on the ARC-Easy and PIQA [13] benchmarks. Similar to Sec. 4.2, the model maintains relatively little performance drop compared to the baseline and the unlearn baseline, with the largest gap being observed on the ARC-Easy benchmark for the uncorrelated experiment.

4.4 EMPIRICAL VERIFICATION OF THEORETICAL GUARANTEES

Theorems 2.2 and 2.1 imply that, after Forgetting-MarI, the mutual information between logits and the "seen/unseen" bit Z is negligible; hence *any* confidence-based test (perplexity, cross-entropy, log-likelihood, etc.) should fail to separate forgotten from genuinely unseen text. We therefore evaluate Forgetting-MarI with the strongest confidence-based detector before/after unlearning.

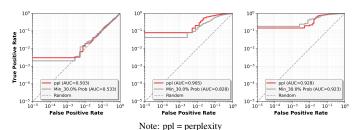


Figure 6: Detector performance for a GPT2-LG model without unlearning (left), unlearned with Forgetting-MarI (middle), and the golden-standard unlearn baseline (bottom).

Copyrighted-text detection methods fall into two families: (i) white-box detectors [3, 47, 29, 39], which use (tail or reference-model) perplexity to infer training membership; (ii) black-box detectors [23, 4, 8, 7, 18], which rely on string-level similarity without logits. Because Forgetting-MarI (and most

threat models) allow weight access, white-box tests are strictly harder to defeat; we thus focus on them. (See Appendix D for method details and additional results.)

We run the current SOTA white-box detector of Shi et al. [39] on: (i) the model finetuned on $\mathcal{D}_r \cup \mathcal{D}_u$, (ii) the *gold-standard* unlearn baseline trained only on \mathcal{D}_r , and (iii) model (i) after Forgetting-MarI. We report ROC-AUC: low values indicate the detector believes the model was trained with \mathcal{D}_u , high values indicate opposite. As shown in Fig. 6, the ROC-AUC after Forgetting-MarI closely matches the unlearn baseline, indicating effective removal of \mathcal{D}_u 's influence, as predicted by theory.

5 CONCLUSION

This work presents Forgetting-MarI, a novel approach to LLM unlearning that improves upon existing state-of-the-art unlearning methods while providing rigorous theoretical guarantees. Our experimental results across multiple benchmarks confirm the practical effectiveness of our proposed technique, while our theoretical analysis establishes formal bounds on the unlearning process and convergence properties.

The combination of strong empirical results and theoretical foundations represents a significant advancement in machine unlearning for LLMs. However, several important directions remain for future research. First, while our theoretical guarantees provide valuable insights into the method's behavior, there is still a gap between theoretical bounds and the practical performance we observed. For example, it is still unknown how Forgetting-MarI finds the unearthed baseline with certainty. Bridging this gap could lead to tighter analysis and potentially improved algorithms.

Second, our work highlights the importance of parameter selection in unlearning effectiveness. Developing principled approaches for optimal parameter tuning, especially with theoretical guidance, remains an open challenge that could significantly enhance the practicality of unlearning methods. Additionally, future work could explore the scalability of our approach to even larger models and datasets, investigate our method's robustness across different model architectures and domains, and the principles of our approach could be applied to models trained on other data modalities.

As LLMs continue to grow in capability and deployment, developing reliable and theoretically grounded unlearning methods becomes increasingly important for responsible AI development and deployment. Forgetting-MarI is an important step towards that end. §

[§]During the preparation of this work, the authors used large language model ChatGPT by OpenAI to refine the language and enhance readability. After using this tool or service, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

REFERENCES

- [1] Lenaic Bourtoule, Vijay Chandrasekaran, Christopher A. Choquette-Choo, Haoran Jia, Aidan L. Travers, Zhiwei Steven Zhang, Nicholas Carlini, Florian Tramer, Dan Boneh, and Úlfar Erlingsson. Machine unlearning. In *IEEE Symposium on Security and Privacy*, pp. 141–159, 2021. doi: 10.1109/SP40001.2021.00003.
- [2] Yinzhan Cao and Junfeng Yang. Towards machine unlearning. In *IEEE Symposium on Security and Privacy*, pp. 463–480, 2015. doi: 10.1109/SP.2015.37.
- [3] Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, Alina Oprea, and Colin Raffel. Extracting training data from large language models. In 30th USENIX Security Symposium, 2021. URL https://www.usenix.org/conference/usenixsecurity21/presentation/carlini. Extended version with larger models.
- [4] Xiangyu Chang, Ruiqi Zhang, and Yuanzhi Li. Name-Cloze membership inference attacks for large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 2023. URL https://aclanthology.org/2023.emnlp-main.789.
- [5] Amy B Cyphert. Generative AI, plagiarism, and copyright infringement in legal documents. *Minn. JL Sci. & Tech.*, 25:49, 2023.
- [6] Zhongqi Ding, Suresh Venkatasubramanian, and Blake Woodworth. Certified data removal in high-dimensional settings. *arXiv* preprint arXiv:2505.07640, 2024.
- [7] Jiajun Dong, Wenjuan Han, Tongshuang Wu, and Xiang Ren. Consistency speaks volumes: A black-box test for memorization in LLMs. *arXiv preprint arXiv:2403.08999*, 2024. URL https://arxiv.org/abs/2403.08999.
- [8] João Duarte, Diego Esteves, and André F. T. Martins. DE-COP: Detecting copyrighted passages in large language models via multi-choice paraphrase tests. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024. URL https://arxiv.org/abs/2404.07111.
- [9] Ronen Eldan and Mark Russinovich. Who's Harry Potter? approximate unlearning in LLMs. *arXiv preprint arXiv:2310.02238*, 2023.
- [10] Junfeng Fang, Houcheng Jiang, Kun Wang, Yunshan Ma, Jie Shi, Xiang Wang, Xiangnan He, and Tat-Seng Chua. AlphaEdit: Null-space constrained model editing for language models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=HvSytvg3Jh.
- [11] Joshua Freeman, Chloe Rippe, Edoardo Debenedetti, and Maksym Andriushchenko. Exploring Memorization and Copyright Violation in Frontier LLMs: A Study of the New York Times v. OpenAI 2023 Lawsuit. arXiv preprint arXiv:2412.06370, 2024.
- [12] Chongyang Gao, Lixu Wang, Chenkai Weng, Xiao Wang, and Qi Zhu. Practical unlearning for large language models. *arXiv preprint arXiv:2407.10223*, 2024.
- [13] Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac'h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. The language model evaluation harness, 07 2024.
- [14] Atoosa Ginart, Leon Guan, Gregory Valiant, and James Zhu. Making AI forget you: Data deletion in machine learning. In *Advances in Neural Information Processing Systems*, pp. 3518–3531, 2019.
 - [15] Michael M Grynbaum and Ryan Mac. The Times sues OpenAI and Microsoft over AI use of copyrighted work. *The New York Times*, 27, 2023.

- 540 [16] Ruiqi Guo, Tom Goldstein, Awni Hannun, et al. Certified data removal from machine learning models. In *International Conference on Artificial Intelligence and Statistics*, pp. 1516–1525, 2020.
 - [17] Yihuai Hong, Yuelin Zou, Lijie Hu, Ziqian Zeng, Di Wang, and Haiqin Yang. Dissecting fine-tuning unlearning in large language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 3933–3941, 2024.
 - [18] Zihan Hu, Yuo Xu, Lavanya Shukla, and Milad Nasr. VeilProbe: Automated black-box membership detection for large language models. *arXiv preprint arXiv:2406.01840*, 2024. URL https://arxiv.org/abs/2406.01840.
 - [19] Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. In *Proceedings of the 11th International Conference on Learning Representations (ICLR 2023)*, 2023.
 - [20] Yoichi Ishibashi and Hidetoshi Shimodaira. Knowledge sanitization of large language models. In *Proceedings of the 2024 Conference of the North American Chapter of the ACL (NAACL 2024)*, 2024.
 - [21] Zhuoran Jin, Pengfei Cao, Chenhao Wang, and et al. RWKU: Benchmarking real-world knowledge unlearning for Large Language Models. In *NeurIPS Datasets and Benchmarks*, 2024.
 - [22] J.K. Rowling. *Harry Potter and the Prisoner of Azkaban*. Harry Potter. Scholastic, 2013. ISBN 9780545582933. URL https://books.google.com/books?id=uUp5mAEACAAJ.
 - [23] Maria Karamolegkou, Emily Dinan, and Angela Fan. Leaks in the library: Prompting large language models to reveal copyrighted text. In *Findings of EMNLP*, 2023. URL https://aclanthology.org/2023.findings-emnlp.456.
 - [24] Bo Liu, Qiang Liu, and Peter Stone. Continual learning and private unlearning. In *Proceedings* of the 1st Conference on Lifelong Learning Agents (CoLLAs), pp. 243–254. PMLR, 2022.
 - [25] Sijia Liu, Yuanshun Yao, Jinghan Jia, Stephen Casper, Nathalie Baracaldo, Peter Hase, Yuguang Yao, Chris Yuhao Liu, Xiaojun Xu, Hang Li, Kush R. Varshney, Mohit Bansal, Sanmi Koyejo, and Yang Liu. Rethinking machine unlearning for large language models, 2024. URL https://arxiv.org/abs/2402.08787.
 - [26] Zheyuan Liu, Guangyao Dou, Zhaoxuan Tan, Yijun Tian, and Meng Jiang. Towards safer large language models through machine unlearning. In *Findings of the Association for Computational Linguistics (ACL 2024)*, 2024.
 - [27] Aengus Lynch, Phillip Guo, Aidan Ewart, Stephen Casper, and Dylan Hadfield-Menell. Eight methods to evaluate robust unlearning in large language models. *arXiv preprint arXiv:2402.16835*, 2024.
 - [28] Pratyush Maini, Zhili Feng, Avi Schwarzschild, Zachary C. Lipton, and J. Zico Kolter. TOFU: A task of fictitious unlearning for LLMs, 2024. URL https://arxiv.org/abs/2401.06121.
 - [29] Rahul Maini, Nicholas Carlini, Kobbi Nissim, and Shafi Goldwasser. When one example fails, a thousand may succeed: Dataset-level membership inference for large language models. In *International Conference on Machine Learning (ICML)*, 2024. URL https://arxiv.org/abs/2403.01234.
 - [30] Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in gpt. In Advances in Neural Information Processing Systems 35 (NeurIPS 2022), 2022.
 - [31] Kevin Meng, Arnab Sen Sharma, Alex Andonian, Yonatan Belinkov, and David Bau. Massediting memory in a transformer. In *Proceedings of the 11th International Conference on Learning Representations (ICLR 2023)*, 2023.

- 594
 595
 [32] Inc. Meta Platforms. meta-llama/Llama-3.2-1B · Hugging Face huggingface.co. https://huggingface.co/meta-llama/Llama-3.2-1B, 2024.
 - [33] Cade Metz. OpenAI says New York Times lawsuit against it is 'without merit'. *The New York Times (Digital Edition)*, pp. NA–NA, 2024.
 - [34] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
 - [35] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *Advances in Neural Information Processing Systems 36 (NeurIPS 2023)*, pp. 53728–53741, 2023.
 - [36] Reddit.r/WritingPrompts.https://www.reddit.com/r/WritingPrompts/, 2025.
 - [37] Fangxin Shi, Yifan Zhou, Nicolas Papernot, and Milad Nasr. Min-k%: A simple yet strong membership inference score for large language models. arXiv preprint arXiv:2402.12345, 2024. URL https://arxiv.org/abs/2402.12345.
 - [38] Shaojie Shi, Xiaoyu Tan, Xihe Qiu, Chao Qu, Kexin Nie, Yuan Cheng, Wei Chu, Yinghui Xu, and Yuan Qi. Ulmr: Unlearning large language models via negative response and model parameter average. In *Proceedings of the 2024 EMNLP: Industry Track*, pp. 755–762. Association for Computational Linguistics, 2024.
 - [39] Weijia Shi, Anirudh Ajith, Mengzhou Xia, Yangsibo Huang, Daogao Liu, Terra Blevins, Danqi Chen, and Luke Zettlemoyer. Detecting pretraining data from large language models, 2023.
 - [40] Lixu Wang, Chongyang Gao, Chenkai Weng, and Qi Zhu. Layer-Aggregated OOD detection for continual LLM unlearning. *arXiv preprint arXiv:2409.01234*, 2024.
 - [41] Xinwei Wu, Junzhuo Li, Minghui Xu, Weilong Dong, Shuangzhi Wu, Chao Bian, and Deyi Xiong. DEPN: Detecting and editing privacy neurons in pretrained language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 2875–2886. Association for Computational Linguistics, 2023.
 - [42] S. Wynn-Williams. Careless People: A Cautionary Tale of Power, Greed, and Lost Idealism. Flatiron Books, 2025. ISBN 9781250391247. URL https://books.google.com/books?id=oHdJEQAAQBAJ.
 - [43] Shizhou Xu and Thomas Strohmer. Machine unlearning via information theoretic regularization. *arXiv preprint arXiv:2502.05684*, 2025.
 - [44] Yuanshun Yao, Xiaojun Xu, and Yang Liu. Large language model unlearning. In *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*, 2024.
 - [45] Ruiqi Zhang, Licong Lin, Yu Bai, and Song Mei. Negative preference optimization: From catastrophic collapse to effective unlearning. *arXiv preprint arXiv:2404.05868*, 2024.
 - [46] Xiangrui Zhang, Ruiqi Zhang, Yuan Zhao, and Suhang Wang. Machine unlearning six-way evaluation (MUSE) for language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 2024.
 - [47] Xiaomin Zhang, Fali Wang, Wenpeng Yin, and Suhang Wang. Domain-calibrated pretraining data detection for large language models. *arXiv preprint arXiv:2405.09876*, 2024. URL https://arxiv.org/abs/2405.09876.

A APPENDIX OF SECTION 1

A.1 DETAILS OF LLM UNLEARNING METHODS: IMPLICIT MARGINAL INFORMATION UNLEARNING

Recent surveys highlight four broad families of LLM-unlearning techniques, each making a different compromise between *unlearn efficacy*, the ability to remove information from a model, *utility preservation*, how well the model performs on the remaining data, and *computational cost*, the resources expended to perform the unlearning [25]. A heuristic commonality of the techniques is their implicit/indirect target of marginal unlearning: all the methods tend to detect and thereby remove only the marginal effect of adding an "unlearn set" (\mathcal{D}_u) , the dataset that is meant to be forgotten, to a "retain set" (\mathcal{D}_r) , the dataset that the model should remember, on the given model.

Full parameter fine-tuning: These techniques train and perform weight updates on the whole model. Gradient ascent (or "loss reversal") [44] is the most straight-forward unlearning technique. It directly maximizes the cross-entropy on \mathcal{D}_u , effectively penalizing the model performance on the unlearn set. This type of unlearning was shown to lead to an overall decrease in model performance, so Gradient Difference [25] was developed to balance unlearning while maintaining general model performance. Gradient Difference maximizes the cross-entropy loss on \mathcal{D}_u while continuing to minimize the loss on \mathcal{D}_r :

$$\min_{\theta} \underbrace{\mathbb{E}_{x \in \mathcal{D}_r} \ell(\theta; x)}_{\text{utility}} - \lambda \underbrace{\mathbb{E}_{x \in \mathcal{D}_u} \ell(\theta; x)}_{\text{loss reversal}},$$
$$\ell(\theta; x) = \text{CE}\left(p_{\theta}(\cdot \mid x_{< t}), x_t\right).$$

Here $\lambda > 0$ balances utility preservation and unlearning. Intuitively, gradient descent is applied on \mathcal{D}_r while gradient ascent is applied on \mathcal{D}_u .

Follow-up studies revealed that, even when balanced with gradient descent, this global ascent signal is too coarse: it suppresses the target examples but also degrades correlated yet legitimate content [17]. To overcome this challenge, variants have aimed to improve both sides of the problem. For utility preservation, past work has shown that distillation-style regularization with a Kullback–Leibler (KL) divergence penalty outperforms gradient descent on \mathcal{D}_r in keeping the updated model close to the original without over-training on the retain set. For unlearning, alignment-style variants such as Direct Preference Optimization (DPO) and Negative Preference Optimization (NPO) replace the unlearn objective with more specific preference-based objectives, slowing catastrophic performance collapse [35, 45]. However, such preference-supervised methods can be difficult to generalize to unlearn at a large scale.

Finally, from the perspective of *marginal* unlearning, these full-parameter objectives act as *indirect* proxies for the marginal effect of adding \mathcal{D}_u to \mathcal{D}_r : they rely on carefully balancing ascent on \mathcal{D}_u and descent (or KL regularization) on \mathcal{D}_r . In practice, such proxies can be neither the most effective nor the most efficient at isolating the unique contribution of \mathcal{D}_u without erasing information shared with \mathcal{D}_r .

Weight editing and partial tuning: In an effort to perform unlearning more efficiently, this line of methods focuses on selectively altering only a subset of a model's parameters rather than retraining the entire network. Such "model-surgery" methods perform rank-constrained updates at one or a few layers. Rank-One Model Editing (ROME) edits a single MLP weight with a closed-form rank-1 patch [30]. In particular, it modifies only the weights causally responsible for one token sequence in the unlearn set:

$$\min_{\Delta W} \|\Delta W\|_F^2 \quad \text{s.t.} \quad W_{l^*} h_{l^*}(x) + \Delta W h_{l^*}(x) = v_{\text{new}}.$$

Here, $\Delta W := \frac{\left(v_{\text{new}} - v_{\text{old}}\right)h^{\top}}{\|h\|_{2}^{2}}$, l^{*} is the layer most influenced by the unlearn sample or prompt x, $h_{l^{*}}(x)$ is the activation and $W_{l^{*}}$ is the weight matrix of layer l^{*} , $v_{old} := W_{l^{*}}h_{l^{*}}(x)$, and finally v_{new} is the alternative answer we want to replace v_{old} by. Mass Editing Memory in a Transformer (MEMIT) [31] extends this idea to thousands of facts simultaneously and stacks many $(h^{i}_{l^{*}}, v^{i}_{l^{*}})$ pairs. AlphaEdit furthers the idea by projecting edits into the null space of preserved knowledge, with the aim to improve robustness in sequential settings, ensuring minimal disruption to previously learned information. Detecting and Editing Privacy Neurons (DEPN) [41] masks the gradients of

neurons identified as contributing the most to the prediction of privacy-related content. In general, weight editing and partial tuning techniques are fast, but they are limited to short factual associations and struggle with stylistic or distributed knowledge.

Finally, the above weight-editing and partial-tuning methods share a common *indirect marginal unlearning* proxy: they infer marginal information by targeting parameters most influenced by the unlearn set, while largely ignoring parameters most influenced by the retain set. This can help isolate some marginal information signal, but again risks overlooking deep interactions between r and u.

Curating counterfactuals: Instead of directly unlearning all or part of the model, another approach is to substitute the parametric knowledge of the unlearn set with benign knowledge. Broadly, this class of methods can be characterized by:

$$\min_{\theta} \ \underbrace{\mathbb{E}_{x \in \mathcal{D}_r} [\ell(\theta; x)]}_{\text{retain utility}} + \lambda \underbrace{\mathbb{E}_{x \in \mathcal{D}_{\text{neg}}} [\ell(\theta; x)]}_{\text{counterfactual prompts}},$$

where \mathcal{D}_{neg} contains prompts or contexts designed to *neutralize* the influence of the unlearn set, $\ell(\theta; x)$ is the same cross entropy loss as before, and $\lambda > 0$ balances unlearning against utility.

"I don't know" [20] trains the model on question-answer pairs that map sensitive questions to a safe refusal (e.g. "I don't know"), teaching the model to decline queries about the unlearn set. Entity anonymization [9] replaces sensitive entities with anonymized placeholders and trains the model on the rewritten placeholders to scrub identifiable information from the model. Unlearning Large Language Models via Negative Response and Model Parameter Average (ULMR) [38] constructs adversarial "negative" prompts, trains on the paired responses, and then averages the updated weights with the base model to dampen overshoot. Selective Knowledge-negation Unlearning (SKU) first mines harmful or copyrighted contexts via red-teaming, then injects counterfactuals that negate them [26]. Such approaches are easy to deploy but depend heavily on prompt engineering and high-quality counterexamples.

From a marginal information proxy perspective, the curating counterfactuals approach aims to first penalize model utility related to the unlearn set by replacing the original model capability on the unlearn set with a lower-utility capacity on the counterfactuals, then rescue the utility related to the retain set using the utility preservation term, and finally balance the two to indirectly find the marginal information and penalize it.

Model adaptation: These methods train something external to the model and then use that externally trained adapter to update the model itself. A common instantiation is the task-vector framework: let p_{θ_0} be the original model and p_{θ_u} the same model fine-tuned on the unlearn set \mathcal{D}_u . The element-wise difference $\Delta\theta := \theta_u - \theta_0$ is treated as an encoding of the deleted knowledge and direct-subtraction methods [19] form the unlearned model as $p_{\theta_0 - \Delta\theta}$. Orthogonality offers an alternative geometric control. O^3 [12] trains one orthogonal LoRA adapter per removal request and learns a contrastive out-of-distribution (OOD) gate that activates the corresponding adapter at inference time. Orthogonality limits interference between requests, but the approach incurs two key costs: (1) the number of adapters (and hence memory) grows linearly with the number of unlearning requests, and (2) any mismatch between model behavior and the assumed linear/inner-product structure in weight space can undermine both unlearning guarantees and downstream utility.

From a marginal-information viewpoint, model-adaptation methods isolate the contribution of \mathcal{D}_u by (i) subtracting the unlearn-induced task vector from the retain base, $\theta_0 - \Delta \theta$, or (ii) enforcing orthogonality between components aligned with retain and the unlearn signals and then penalize the isolated component. Both approaches can be considered as proxy of marginal information, though with strong arithmetic or geometric assumptions.

The proposed method, Forgetting-MarI, belongs to the full-parameter fine-tuning category. It applies a "marginal information" penalty that suppresses only the influence of the unlearn set while leaving the shared information, which is supported by the retain data, largely intact.

B APPENDIX OF SECTION 2

B.1 Proof of Proposition 2.1

Proof. To start, define the Bayes error as

$$P_e := \mathbb{E}_{X_{\text{MarI}}} \Big[\min \{ P(Z = 0 \mid X_{\text{MarI}}), P(Z = 1 \mid X_{\text{MarI}}) \} \Big] = 1 - P_{acc}.$$

In addition, for each x, let $p(x) := P(Z = 1 \mid X_{\text{MarI}} = x) \in [0,1]$ be the conditional probability of $\{Z = 1\}$ given $\{X_{\text{MarI}} = x\}$. Then it follows from Z being binary that $H(Z \mid X_{\text{MarI}} = x) = H_2(p(x))$. Denote $m(X_{\text{MarI}}) := \min \left\{ P(Z = 0 \mid X_{\text{MarI}}), P(Z = 1 \mid X_{\text{MarI}}) \right\}$. Since H_2 is concave, it follows from Jensen's inequality that

$$\begin{split} H(Z \mid X_{\text{MarI}}) &= \mathbb{E}_{X_{\text{MarI}}} \big[H_2(p(X_{\text{MarI}})) \big] \\ &= \mathbb{E}_{X_{\text{MarI}}} \big[H_2(m(X_{\text{MarI}})) \big] \\ &\leq H_2 \big(\mathbb{E}_{X_{\text{MarI}}} [m(X_{\text{MarI}})] \big) \\ &= H_2(P_e). \end{split}$$

where the second equality holds due to the fact that $H_2(p) = H_2(1-p)$. Now, since $I(X_{\text{MarI}}; Z) = H(Z) - H(Z \mid X_{\text{MarI}})$ and $H(Z) = H_2(\pi)$, we obtain

$$H_2(\pi) - I(X_{\text{MarI}}; Z) = H(Z \mid X_{\text{MarI}}) \le H_2(P_e).$$

Since $P_e \in [0, \frac{1}{2}]$ and H_2 is strictly increasing on this interval, by applying the inverse H_2^{-1} , we have

$$P_e \geq H_2^{-1}(H_2(\pi) - I(X_{\text{MarI}}; Z)).$$

Finally, by $P_{acc} = 1 - P_e$, we have

$$P_{acc} \le 1 - H_2^{-1}(H_2(\pi) - I(X_{\text{MarI}}; Z)).$$

This proves the stated inequality. The particular case $\pi = \frac{1}{2}$ follows from $H_2(\frac{1}{2}) = 1$.

It remains to show that the upper bound is tight. Indeed, fix an arbitrary $I \in [0, H_2(\pi)]$. Choose $p^\star \in \left[\frac{1}{2}, 1\right]$ such that $H_2(p^\star) = H_2(\pi) - I$. Construct $P_{Z|X_{\mathrm{MarI}}}$ such that $P(Z = 1 \mid X_{\mathrm{MarI}}) \in \{p^\star, 1 - p^\star\}$ with probabilities chosen to match the prior π . Then $H(Z \mid X_{\mathrm{MarI}}) = H_2(p^\star)$ and $I(X_{\mathrm{MarI}}; Z) = I$, while the Bayes error satisfies $P_e = \min\{p^\star, 1 - p^\star\} = H_2^{-1}(H(Z) - I)$. Hence, equality holds in the bound.

B.2 WHY MUTUAL INFORMATION RATHER THAN KL DIVERGENCE

One might consider penalizing a directional KL divergence between the "to-unlearn" and "to-retain" distributions. Instead, we regularize the *mutual information* between the model output and a binary indicator of sensitive content, which is equal to the Jensen-Shannon divergence as shown in Section 2. Here, we show that mutual information offers several advantages over one—way or two-way KL divergence:

- Flexibility for utility and continual unlearning. The reference m in Jensen-Shannon divergence is the *mixture* of the two conditionals and evolves with training; we do not assume a fixed "gold" model. This yields a pure unlearning regularizer that can be combined with any utility term (e.g., $\ell_{\mathrm{KL}}(\theta,r)$) and naturally supports continual/online updates.
- Stable training signal. $I(\hat{X}; Z) \leq H_2(\pi) \leq \log 2$ for binary Z, so the gradients remain well-behaved even when supports differ, unlike one—way KL which can be unbounded on support mismatch.
- Downstream robustness via data processing. For any downstream representation or task $T = g(\hat{X})$, the data-processing inequality gives $I(T; Z) \leq I(\hat{X}; Z)$. Thus, suppressing $I(\hat{X}; Z)$ at the model output (or an internal layer) upper-bounds leakage throughout the pipeline.

In contrast, a directional KL requires committing to a *fixed* target (encoding a specific utility assumption) and can be unstable or unbounded when supports are disjoint. That said, if an ideal frozen reference is indeed mandated, a one—way KL to that reference is a reasonable alternative.

B.3 Proof of Theorem 2.1

Here, we provide the proof for Theorem 2.1:

Proof. By the mean value theorem, for each t there exists $\xi_t \in [\min\{p_t^u(u_t), p_t^r(u_t)\}, 1] \subseteq [\gamma, 1]$ such that

$$\begin{aligned} \left| \log p_t^u(u_t) - \log p_t^r(u_t) \right| &= \frac{\left| p_t^u(u_t) - p_t^r(u_t) \right|}{\xi_t} \\ &\leq \frac{\left| p_t^u(u_t) - p_t^r(u_t) \right|}{\gamma} \\ &\leq \frac{\left\| p_t^u - p_t^r \right\|_1}{\gamma} \\ &= \frac{2 \left\| p_t^u - p_t^r \right\|_{TV}}{\gamma}. \end{aligned}$$

Averaging over t,

$$|S_{\theta}(u,u) - S_{\theta}(u,r)| \leq \frac{2}{\gamma} \frac{1}{T} \sum_{t=1}^{T} ||p_{t}^{u} - p_{t}^{r}||_{TV}.$$

Apply Lemma B.3 followed by Lemma B.2 and Jensen's inequality:

$$\frac{1}{T} \sum_{t} \|p_{t}^{u} - p_{t}^{r}\|_{TV} = \frac{1}{1 - \alpha} \frac{1}{T} \sum_{t} \|p_{t}^{d} - p_{t}^{r}\|_{TV} \leq \frac{\sqrt{2}}{1 - \alpha} \sqrt{\frac{1}{T} \sum_{t} \text{JSD}(p_{t}^{d}, p_{t}^{r})}.$$

Combining yields the claim.

B.4 Proof of Theorem 2.2

We start with the following three lemmata that are needed for the proof of Theorem 2.2:

Lemma B.1 (Point-wise KL bound). Let p,q be two probability distributions over a finite set V such that $\frac{p(x)}{q(x)} \in [1,M]$ for every $x \in V$ for some constant M > 1. Then for every $x \in V$

$$p(x) \log \frac{p(x)}{q(x)} \le (\log M) \frac{M}{M-1} \left[p(x) - q(x) \right]. \tag{4}$$

Proof. Fix $x \in V$ and set $y := \frac{p(x)}{q(x)} \in [1, M]$. Inequality equation 4 is equivalent to

$$y\log y \le \frac{M}{M-1}(\log M)(y-1), \qquad \forall y \in [1, M]. \tag{4}$$

For y>1 let $g(y):=\frac{y\log y}{y-1}$ and set $g(1):=\lim_{y\to 1^+}g(y)=1$. We show that g is strictly increasing on [1,M]. Indeed, compute $g'(y)=\frac{(y-1)-\log y}{(y-1)^2}$. Since $\log y< y-1$ for all y>1, we have g'(y)>0; thus, g is strictly increasing. Because g is increasing and $y\in [1,M]$, we have

$$g(y) \leq g(M) = \frac{M \log M}{M - 1}.$$

Multiplying both sides by y-1 yields equation 4, which is precisely equation 4 after reinstating y=p(x)/q(x). Therefore, equation 4 holds for every $x \in V$. This completes the proof.

Lemma B.2. (Total Variation is controlled by Jensen-Shannon Divergence) For any two probability measures p, q on a finite set, we have

$$||p - q||_{TV} \le \sqrt{2 \operatorname{JSD}(p, q)},$$

where $JSD(p,q) := \frac{1}{2}D_{KL}(p||m) + \frac{1}{2}D_{KL}(q||m)$, $m := \frac{p+q}{2}$ and $D_{KL}(p||q) := \sum_{v} p(v) \log \frac{p(v)}{q(v)}$, denotes the Jensen–Shannon divergence.

Proof. Let $m=\frac{p+q}{2}$. Pinsker's inequality gives $\|p-m\|_1^2 \leq 2 D_{\mathrm{KL}}(p\|m)$ and analogously for q. Hence

$$JSD(p,q) \ge \frac{1}{4} \left[\|p - m\|_1^2 + \|q - m\|_1^2 \right] = \frac{1}{8} \|p - q\|_1^2,$$

because $p-m=\frac{p-q}{2}$ and $q-m=-\frac{p-q}{2}$. Since $||p-q||_{TV}=\frac{1}{2}||p-q||_1$, it follows that $||p-q||_{TV}^2 \leq 2 \text{ JSD}(p,q)$.

Lemma B.3 (Exact TV scaling under mixture). If $p^d = \alpha p^r + (1 - \alpha)p^u$ with $\alpha \in (0, 1)$, then

$$||p^u - p^r||_{TV} = \frac{1}{1 - \alpha} ||p^d - p^r||_{TV}.$$

Proof. $p^d - p^r = (1 - \alpha)(p^u - p^r)$. Taking ℓ_1 -norms and dividing by 2 yields the identity.

Now, we are ready to prove Theorem 2.2:

Proof. Define $Y_t := \log \frac{p_t^r(U_t)}{p_t^u(U_t)}$, so that

$$S_{\theta}(U,u) - S_{\theta}(U,r) = \frac{1}{T} \sum_{t=1}^{T} Y_{t}.$$

Since $U_t \sim p_t^u$, $\mathbb{E}[Y_t] = \sum_x p_t^u(x) \log \frac{p_t^r(x)}{p_t^u(x)} = -D_{\mathrm{KL}}(p_t^u || p_t^r)$, hence

$$\mathbb{E}[S_{\theta}(U, u) - S_{\theta}(U, r)] = -\frac{1}{T} \sum_{t=1}^{T} D_{KL}(p_t^u || p_t^r).$$

Now, by the assumption $\max_{t,x} \max \left\{ \frac{p_t^u(x)}{p_t^r(x)}, \frac{p_t^r(x)}{p_t^u(x)} \right\} \leq M$, we have $p_t^r(x) > 0$ for $p_t^u(x)$ -a.e. x for all t. Therefore, for all t, we have $\log \frac{p_t^r(x)}{p_t^u(x)} < \infty$ and taking the maximum over $t \in [T]$, we obtain $C := \max_{t, \ x: \ p_t^u(x) > 0} \left[\log \frac{p_t^r(x)}{p_t^u(x)}\right]^2 < \infty$. It then follows from the definition of Y_t that $|Y_t| \leq \sqrt{C}$ a.s.. Hoeffding's inequality for independent bounded variables yields, for any $\varepsilon > 0$,

$$\mathbb{P}\left(\left|\frac{1}{T}\sum_{t=1}^{T}Y_{t} - \mathbb{E}\frac{1}{T}\sum_{t=1}^{T}Y_{t}\right| \geq \varepsilon\right) \leq 2\exp\left(-\frac{T\varepsilon^{2}}{2C}\right).$$

Using $|a| - b| \le |a - b|$ for $b \ge 0$, we have

$$\left|S_{\theta}(U,u) - S_{\theta}(U,r)\right| \leq \frac{1}{T} \sum_{t=1}^{T} D_{\mathrm{KL}}(p_t^u || p_t^r) + \varepsilon.$$

with probability at least $1 - 2\exp(-\frac{T\varepsilon^2}{2C})$.

Now, for each t, let $A_t = \{x: p_t^u(x) \ge p_t^r(x)\}$. Then by Lemma B.1, we have

$$D_{\mathrm{KL}}(p_t^{\,u} \| p_t^{\,r}) \, \leq \, \kappa(M) \, \sum_{x \in A_t} \left(p_t^{\,u}(x) - p_t^{\,r}(x) \right) \, \leq \, \kappa(M) \, \| p_t^{\,u} - p_t^{\,r} \|_{TV}.$$

Averaging in t gives

$$\frac{1}{T} \sum_{t=1}^{T} D_{\mathrm{KL}}(p_t^u \| p_t^r) \leq \kappa(M) \frac{1}{T} \sum_{t=1}^{T} \| p_t^u - p_t^r \|_{TV}.$$

Finally, it follows from Lemma B.3 and Lemma B.2 that

$$\frac{1}{T} \sum_{t=1}^T \| p_t^u - p_t^r \|_{TV} \ = \ \frac{1}{1-\alpha} \, \frac{1}{T} \sum_{t=1}^T \| p_t^d - p_t^r \|_{TV} \ \le \ \frac{\sqrt{2}}{1-\alpha} \, \frac{1}{T} \sum_{t=1}^T \sqrt{\mathrm{JSD}(p_t^d, p_t^r)}.$$

By Jensen's inequality, $\frac{1}{T}\sum_t \sqrt{\mathrm{JSD}(p_t^d, p_t^r)} \leq \sqrt{\frac{1}{T}\sum_t \mathrm{JSD}(p_t^d, p_t^r)}$. Combining the displays proves the claim with $I(X_{\mathrm{MarI}}; Z) = \frac{1}{T}\sum_t \mathrm{JSD}(p_t^d, p_t^r)$.

C APPENDIX OF SECTION 3

C.1 POSITION-WISE VS. POOLED MARI: EMPIRICAL COMPARISON

We empirically compare the token/position-wise MarI, $I(X_{\text{MarI}}; Z) = \frac{1}{T} \sum_{t=1}^{T} I(X_t; Z)$, with the pooled ("flattened") MarI, $I(\bar{X}_{\text{MarI}}; Z)$, on our heterogeneous dataset. As predicted by the data–processing inequality, $I(\bar{X}_{\text{MarI}}; Z) \leq I(X_{\text{MarI}}; Z)$, so the position-wise estimator produces a stronger marginal-information signal. Nevertheless, by appropriately tuning the trade-off parameter γ (weighting MarI vs. utility), both estimators attain comparable forget–utility trade-offs.

However, we can also observe the influence of the heterogeneity of dataset and random batch sampling. In particular, in Figure 7, the position-wise estimator exhibits higher variance on heterogeneous batches (varying lengths, topics, and token alignments). Furthermore, Figure 8 shows that, with fixed γ (e.g., $\gamma=0.9$), the position-wise MarI tends to over-unlearn relative to the gold unlearn baseline. Intuitively, it can over-penalize idiosyncratic, position-specific fluctuations rather than true marginal effects.

A deeper theoretical and experimental analysis of these behaviors, hybrids that combine both signals, is left for future work. In our experiments, because the text is heterogeneous in both length and context, we use random mini-batches and the pooled estimator by default.

D APPENDIX OF SECTION 4

D.1 More Robustness Test Results

In Figure 9 we show that 10/90 split unlearning training curve, which supplements the 50/50 split unlearning training curve in Figure 3 of Section 4.

In Figure 10 we show the training curve using a smaller learning rate than the result shown in Figure 3.

D.2 SUPPLEMENTARY GENERAL MODEL CAPACITY TEST RESULTS

Table 7 illustrates comprehensive evaluation results across multiple benchmark test.

E APPENDIX OF SECTION 4.4: DETECTION TESTS

E.1 Detector Methods

Here, we provide a more detailed introduction to the current study of copyright content detectors for LLMs so that readers better understand the numerical study in section 4.4. The current study of copyrighted text detectors can be roughly separated into two lines of work:

- White-box methods: Perplexity outlier and reference model perplexity outlier [3], domain normalized minimum k-percentage [47], and data-set level inference [29].

 The above methods largely share the same idea of constructing a statistic (or a vector of statistics) that indicates the probability that a model has seen a given sentence or not. It bases the probability on how confidently the model predicts the true output. The idea is based on the intuition that a model that has seen the sentence during training will have high confidence when trying to complete
- Black-box methods: Direct regurgitation probes [23], Name-cloze membership inference [4], DE-COP: multi-choice preference [8], Output-consistency measures [7], and VeilProbe [18]. Black-box methods, which do not have access to the model parameters and therefore the output logits or prediction distributions, often use either edit distance (a.k.a. Levenshtein distance) or some token embedding model (e.g. a small transformer) to quantify the distance or similarity between a model's output and a reference string, then generate statistics of the similarity between the two.

Task	Metric	GPT-2	Baseline	Unlearn Baseline	F-MarI
ARC-Easy	acc	0.53 ± 0.01	0.46 ± 0.01	0.48 ± 0.01	0.45 ± 0.01
	acc_norm	0.47 ± 0.01	0.42 ± 0.01	0.43 ± 0.01	0.43 ± 0.01
ARC-Challenge	acc	0.22 ± 0.01	0.23 ± 0.01	0.22 ± 0.01	0.22 ± 0.01
	acc_norm	0.25 ± 0.01	0.27 ± 0.01	0.27 ± 0.01	0.26 ± 0.01
PIQA	acc	0.70 ± 0.01	0.66 ± 0.01	0.66 ± 0.01	0.66 ± 0.01
	acc_norm	0.69 ± 0.01	0.65 ± 0.01	0.65 ± 0.01	0.66 ± 0.01
Hellaswag	acc	0.36 ± 0.00	0.36 ± 0.00	0.36 ± 0.00	0.35 ± 0.00
	acc_norm	0.45 ± 0.00	0.43 ± 0.00	0.43 ± 0.00	0.42 ± 0.00
MMLU - humanities - other - social sciences - stem	acc acc acc acc	0.23 ±0.00 0.25 ±0.01 0.24 ±0.01 0.22 ±0.01 0.22 ±0.01	0.23 ± 0.00 0.24 ± 0.01 0.24 ± 0.01 0.22 ± 0.01 0.22 ± 0.01	0.24 ± 0.00 0.25 ± 0.01 0.24 ± 0.01 0.22 ± 0.01 0.23 ± 0.01	0.23 ±0.00 0.24 ±0.01 0.24 ±0.01 0.22 ±0.01

Task	Metric	GPT-2	Baseline	Unlearn Baseline	F-MarI
WikiText	bits/byte	0.841	0.925	0.905	0.910
	byte-pplx	1.792	1.898	1.873	1.879
	word-pplx	22.61	30.80	28.66	29.19

Table 7: Comprehensive evaluation results across multiple benchmarks for the 10/90 split test with GPT2-LG baselines. When the true underlying stopping criteria is unknown, we choose to stop when the detector fails. This corresponds to F-Mari with lambda = 0.95 and epoch = 10.

Black-box methods are weaker detectors than white-box methods since they do not have access to a model's internals. Since our method assumes access to the model parameter, we tested our method against the current SotA white-box method, the minimum k-percent method [37], to demonstrate the effectiveness of our unlearning in real-world applications.

E.2 MULTIPLE DETECTION TEST RESULTS

Here, we provide more detailed results in addition to Figure 6 in Section 4.4.

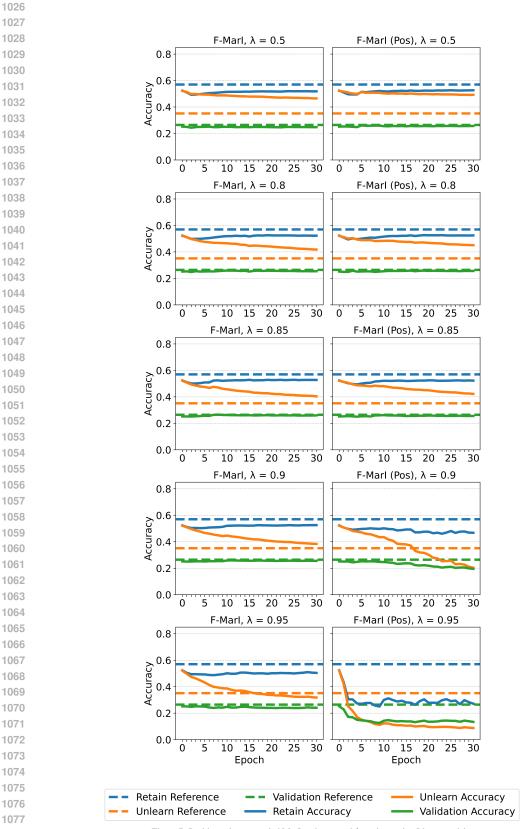


Figure 7: Position-wise vs. pooled MarI under several λ settings using Llama models.

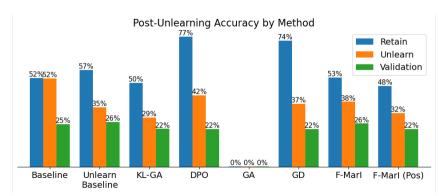


Figure 8: With a fixed trade-off ($\lambda=0.9$) on LLama models, position-wise MarI is noisier on heterogeneous data and over-unlearns compared to the unlearn baseline.

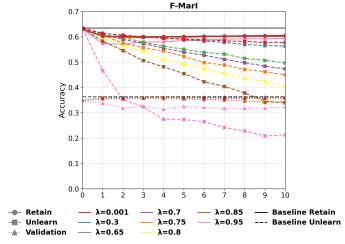
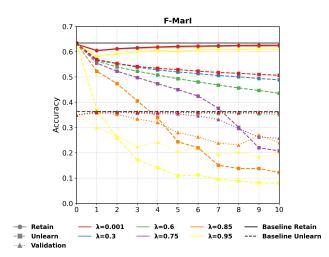
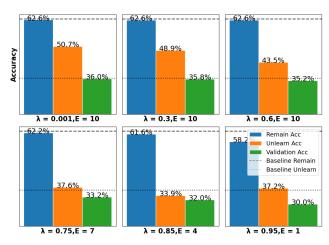


Figure 9: Training curves of F-MarI for the 10/90 split unlearning with the GPT2-LG model. When you do not know the true underlying stopping criteria, we choose to stop when the detector fails. This corresponds to F-Mari, lambda = 0.95 and epoch = 10.



(a) Accuracy versus training epochs.



(b) Unlearning results using different regularization with stopping E.

Figure 10: 50/50 split test, F-MarI training curves using small learning rate using GPT2-LG to better observe the convergence process. Both plot shows stable curve and robustness against lambda choice. The stopping criteria is when validation accuracy drop 3 percent, with the corresponding epoch denoted as 'E'

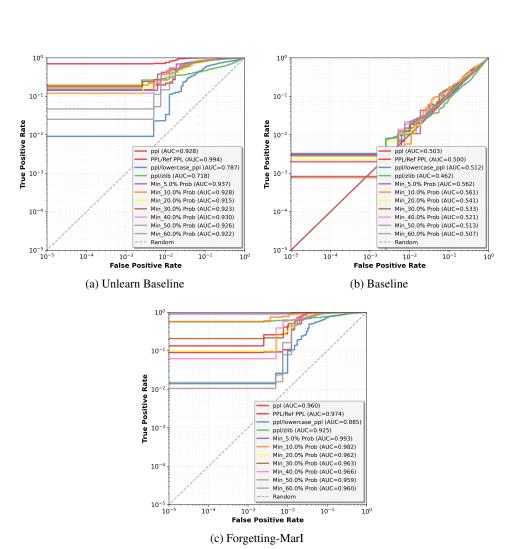


Figure 11: Training data membership detection test of Forgetting-MarI against state-of-the-art detection methods using the 10/90 split Unlearning of the GPT2-LG.