UnPII: Unlearning Personally Identifiable Information with Quantifiable Exposure Risk

Intae Jeon
intae.jeon@samsung.com
Samsung Research
Seoul, South Korea

Yujeong Kwon
shr2008@g.skku.edu
Sungkyunkwan University
Suwon, South Korea

Hyungjoon Koo*
kevin.koo@skku.edu
Sungkyunkwan University
Suwon, South Korea

Abstract

The ever-increasing adoption of Large Language Models in critical sectors like finance, healthcare, and government raises privacy concerns regarding the handling of sensitive Personally Identifiable Information (PII) during training. In response, regulations such as European Union's General Data Protection Regulation (GDPR) mandate the deletion of PII upon requests, underscoring the need for reliable and cost-effective data removal solutions. Machine unlearning has emerged as a promising direction for selectively forgetting data points. However, existing unlearning techniques typically apply a uniform forgetting strategy that neither accounts for the varying privacy risks posed by different PII attributes nor reflects associated business risks. In this work, we propose UnPII, the first PII-centric unlearning approach that prioritizes forgetting based on the risk of individual or combined PII attributes. To this end, we introduce the PII risk index (PRI), a composite metric that incorporates multiple dimensions of risk factors: identifiability, sensitivity, usability, linkability, permanency, exposability, and compliancy. The PRI enables a nuanced evaluation of privacy risks associated with PII exposures and can be tailored to align with organizational privacy policies. To support realistic assessment, we systematically construct a synthetic PII dataset (*e.g.,* 1,700 PII instances) that simulates realistic exposure scenarios. UnPII seamlessly integrates with established unlearning algorithms, such as Gradient Ascent, Negative Preference Optimization, and Direct Preference Optimization, without modifying their underlying principles. Our experimental results demonstrate that UnPII achieves the improvements of accuracy up to 11.8%, utility up to 6.3%, and generalizability up to 12.4%, respectively, while incurring a modest fine-tuning overhead of 27.5% on average during unlearning.

CCS Concepts

• **Security and privacy → Privacy protections**.

Keywords

Machine Unlearning, Personally Identifiable Information

ACM Reference Format:
Intae Jeon, Yujeong Kwon, and Hyungjoon Koo. 2026. UnPII: Unlearning Personally Identifiable Information with Quantifiable Exposure Risk. In *2026 IEEE/ACM 48th International Conference on Software Engineering (ICSE-SEIP '26), April 12–18, 2026, Rio de Janeiro, Brazil.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3786583.3786860

*Corresponding author.

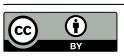

ICSE-SEIP '26, Rio de Janeiro, Brazil

ACM ISBN 979-8-4007-2426-8/2026/04
https://doi.org/10.1145/3786583.3786860

1 Introduction

Today, the widespread adoption of Large Language Models (LLMs) is common in various sectors across domains such as finance [15] (*e.g.,* credit scoring), healthcare [12, 48] (*e.g.,* diagnosis assistance, drug detection), government [36, 37] (*e.g.,* public safety, policy analytics), and education [20, 33] (*e.g.,* tutoring support, content recommendation). Those models are often trained on sensitive user data, including personally identifiable information (hereinafter referred to as PII), which can be unintentionally memorized and revealed during inference.

While direct and high-profile data breach, such as Marriott [24] (disclosing contact and payment information of 5.2 million individuals) and Equifax [10] (compromising records of over 140 million U.S. consumers) cases, draw significant attention, the indirect exposures [29] from trained models also raise privacy concerns. In response, regulations such as the General Data Protection Regulation (GDPR) [35], the California Consumer Privacy Act (CCPA) [4], and China's Personal Information Protection Law (PIPL) [42] mandate the complete elimination (*e.g.,* right to erasure or right to be forgotten) of PII upon user request or following a security breach.

To address such concerns, *machine unlearning* has emerged as a promising direction for selectively removing specific training data points without full model retraining. In practice, however, efficient post-hoc data deletion in large-scale models presents several challenges. First, removing individual records from models with billions of parameters is computationally demanding in the absence of original training corpus [2]. Second, unlearning strategies often impair global model utility, as the removal process inadvertently updates knowledge beyond the intended deletion scope [23]. Third, unlearning must ideally support incremental and online operation to handle deletion requests as they arise, without disrupting service availability [16] while preserving acceptable resource efficiency [30]. Lastly, given the evolving landscape of global privacy regulations, the adaptability of unlearning techniques to shifting legal requirements constitutes a key operational advantage. Collectively, these technical, operational, and regulatory hurdles highlight the need for specialized methodologies that enable efficient, reliable, and legally compliant data removal in production-scale models.

Early methods, such as SISA [2], improve efficiency by partitioning the dataset into isolated shards and retraining only those affected by deletion requests. However, such means face scalability challenges when applied to large-scale models (*e.g.,* LLMs) due to their high computational cost. Gradient-based approaches [3, 23, 39]

update model parameters to reverse the influence of unlearning samples. Meanwhile, preference optimization approaches [23, 27, 34, 47] modify model outputs by penalizing undesirable responses and promoting preferred alternatives. Recent advancements in parameter-efficient techniques [6, 11] reduce computational overhead by limiting updates to a small subset of parameters; however, they often assume the availability of reference models or a retention dataset to maintain original performance. Another direction is entity-level unlearning, such as Opt-Out [7], which enables the removal of entire entity embeddings (*e.g.,* users, items, records).

While the previous approaches are effective in certain unlearning contexts, their effectiveness may diminish when applied to PII-level forgetting on real-world datasets. The limitation arises from two key challenges: ① applying the same forgetting strategy to every PII attribute (*e.g.,* name, social security number) in the unlearning dataset may not be effective without accounting for differences in privacy sensitivity (or risk), and ② prior works often rely on a retention set, which may be unavailable in (PII-relevant) practical scenarios.

In this work, we propose UnPII, the first PII-centric unlearning approach that dynamically prioritizes PII forgetting based on the privacy risk of individual or combined PII attributes. By design, UnPII can be seamlessly incorporated with *any gradient-based* unlearning methods, such as Gradient Ascent (GA) [23], Negative Preference Optimization (NPO) [47], and Direct Preference Optimization (DPO) [34]. In essence, given a PII-containing model, UnPII begins with identifying PII in the model's output through consultations with a large language model using PII-inducing queries. Next, UnPII computes a PII risk index (PRI; ranging from 0 to 1), quantifying the exposure risk of PII, either individually or in combination. Lastly, UnPII unlearns the target dataset by applying a gradient-scaling loss function, adjusting the forgetting signal based on the computed risk value. To facilitate this, we construct a synthetic PII dataset (*e.g.,* 1,700 PII examples), generating a PII-containing model. Besides, we introduce a quantifiable PII risk assessment metric that evaluates the privacy risk associated with the PII attribute exposure, which captures varying factors such as identifiability, sensitivity, usability, linkability, permanency, exposability, and compliancy.

From an industrial MLOps (Machine Learning Operations) perspective, UnPII can be seamlessly integrated into Continuous Training (CT) pipelines. Unlike full retraining that incurs prohibitive GPU costs and poses risks to service availability, UnPII operates as a parameter-efficient fine-tuning module. This design allows practitioners to batch PII deletion requests and apply updates within standard CI/CD (Continuous Integration and Continuous Delivery) cycles, thereby reconciling strict privacy compliance requirements with operational efficiency.

Our empirical evaluation, conducted with three baseline approaches (*e.g.,* GA, NPO, DPO) across various forgetting ratios (*e.g.,* 1%, 5%, 10%), demonstrates that UnPII outperforms these baselines, improving the harmonic mean (up to 5%) of accuracy (up to 11.8%), utility (up to 6.3%), and generalizability (up to 12.4%), with a modest fine-tuning overhead (27.5% on average).

The main contributions of our paper are as follows:

- We introduce UnPII, the first PII-centric machine unlearning approach that dynamically prioritizes the PII forgetting based on the risk of individual or combined PII attributes.
- We propose a quantifiable PII risk assessment metric that evaluates the privacy risk associated with PII attribute exposure.
- We construct a synthetic PII dataset (*e.g.,* 1,700 PII instances) that simulates realistic exposure scenarios, providing a benchmark for evaluating the effectiveness of PII unlearning techniques.
- We integrate UnPII with three (popular) unlearning techniques (*i.e.,* GA, NPO, and DPO) and evaluate them in terms of accuracy, utility, and generalizability.

We released our source code and dataset to foster further research in the field of machine unlearning for privacy protection [1].

2 Background

Machine Unlearning and Challenges. Contemporary regulations such as GDPR (EU) [35] and CCPA (California) [4] enforce the *right to be forgotten*, granting individuals the authority to request the deletion of their (potentially sensitive) personal data. In response, machine unlearning has emerged as a promising approach to forget a subset of data samples from a trained model. Namely, given an initial model $f(\theta_{init})$ on the full dataset D, it aims to derive an unlearned model $f(\theta_u)$ that forgets the unlearning dataset ($D_f \subset D$), while preserving the behavior of a reference model $f(\theta_r)$ retrained from scratch on the retention dataset ($D_r = D \setminus D_f$).

Formally, this objective can be written as:

$$\text{Unlearn}(f(\theta_{init}), D_f) = f(\theta_u) \approx f(\theta_r) \tag{1}$$

However, machine unlearning faces several challenges. First, achieving selective forgetting is difficult as many existing methods [2, 6, 14, 17, 23, 27, 34, 47] treat all training data uniformly, limiting their ability to effectively remove specific sensitive or harmful samples. Second, unlearning introduces a trade-off between forgetting effectiveness and overall model utility: *i.e.,* aggressively forgetting targets may bring about catastrophic forgetting, degrading performance on the retention dataset. Third, validating that the model has *completely forgotten* the designated samples remains open. This paper focuses on PII-centric data samples, integrating their associated risks into a quantifiable index during the unlearning process. Besides, we propose a risk assessment metric to strike a balance among unlearning accuracy, model utility, and generalizability.

PII Risk Assessment. PII refers to any data that can directly or indirectly identify an individual, such as names, addresses, phone numbers, social security numbers, and passport numbers. This type of information is inherently sensitive, and its exposure can lead to privacy breaches, legal liabilities, and ethical concerns [18, 38, 41]. To assess the risks associated with PII exposures, authoritative institutions such as the National Institute of Standards and Technology (NIST) [25], the Department of Homeland Security (DHS) [9], and the Health Insurance Portability and Accountability Act (HIPAA) [1], have developed risk-based frameworks grounded in quantitative evaluation criteria. While each institution defines and evaluates PII risk within its respective domain, this leads to inconsistencies in risk assessments across contexts and organizations. This paper proposes a unified and quantifiable PII risk assessment

[1] https://github.com/SecAI-Lab/unpii

Table 1: PII categories and representative PII attribute examples of sensitive information. Disclosure of such private information can infringe on a person's privacy.

Category	Representative PII attribute
Basic	Name, date of birth, gender, nationality
Contact	Region address, detailed address, work address, email address, Phone number, social media, personal website, blog
Identifiers	Social security number, work permit number, passport number, Driver license number
Financial	Bank account number, credit card number, card expiration date, income information, Card security code, Credit score, loan details, tax records, cryptocurrency wallet address
Biometric	Fingerprint data, DNA information, iris scan data, facial recognition data, Voice recognition data
Medical	Medical record, health insurance ID number, Hospitalization record, disability status, diagnosis history, mental health record
Employment-related	Job title, employment history, salary, Employee ID number
Education-related	Student ID number, transcript
Digital Footprints	IP/MAC address, device identifier, browsing/search history
Location	ZIP code, Vehicle registration number, real-time location, GPS coordinate
Legal	Criminal record, bankruptcy filing, driving record, court record
Miscellaneous	Insurance policy number, E-signature, call log, voice-mail data

Table 2: Seven PII risk factors for our quantitative assessment of PII leakage. Each factor captures a distinct facet of risk, including identifiability, sensitivity, usability, linkability, permanency, exposability, and compliancy. Every factor of a PII attribute represents a value in the range (0,1), which is parameterized based on organizational policies or requirements.

PII Risk Factor	Description
Identifiability	Uniqueness of the PII that can identify an individual
Sensitivity	Potential psychological and social harm upon the PII exposures
Usability	Usefulness of the PII for attackers in carrying out malicious actions
Linkability	Likelihood that the PII can be linked to other data sources
Permanency	Difficulty in changing or revoking the PII once exposed
Exposability	Frequentness or broadness of the PII exposures during normal use
Compliancy	Severity of legal or regulatory consequences upon the PII breach

metric, which can integrate various criteria from these domain-specific approaches.

3 PII Attributes and Risk Factors

PII Categories and Attributes. We analyze major PII breach incidents [26] that are publicly available, classifying PII attributes into 12 categories. We confirm that severe privacy risks often arise from combinations of PII. For instance, the Marriott breach leaked contact details (*e.g.,* phone numbers, email, addresses) alongside personal attributes (*e.g.,* company, gender, birth date) [24]. The Facebook breach exposed user IDs and phone numbers linked with names and location data [8]. Meanwhile, the Equifax breach revealed full names, birth dates, SSNs, and addresses [10], demonstrating the compounded risk of aggregated identifiers. Table 1 organizes representative PII attributes with their corresponding classes.

PII Risk Factors. With a thorough analysis of risk assessment from NIST [25], DHS [9], and HIPAA [1], we carefully define seven risk factors capable of capturing a distinct aspect of risk, including identifiability, sensitivity, usability, linkability, permanency, exposability, and compliancy as shown in Table 2. For instance, NIST specifies "directly identifiable" (*e.g.,* SSNs) and "linkable" attributes (*e.g.,* ZIP codes), which map to high- and low-risk values, respectively. These factors represent flexible values, ranging from 0 to 1, which we parameterize based on institutional policies or privacy requirements. Unlike existing factor-based assessments, we identify the risk-driven factors tied to the exposure of each PII attribute. These risk factors are configurable within a range of $[0, 1]$, allowing organizations to tailor them to align with their privacy policies, which further guide the computation of per-sample unlearning intensity (§4.2). This alignment ensures stronger deletion for high-risk PII while permitting lighter treatment of low-risk cases, resulting in a more cost-effective, auditable, and defensible compliance posture.

4 UnPII: Unlearning PII with PII Risk Index

Approach Overview. We assume a PII-containing model and an unlearning dataset in the absence of a retention dataset. We introduce UnPII, *a PII-centric unlearning approach that is compatible with existing unlearning methods.* As illustrated in Figure 1, UnPII operates in three stages. First, UnPII identifies the PII exposures by consulting an LLM, such as GPT-4o mini [21], to evaluate its output. The rationale behind this approach builds on a recent study [32], which systematically demonstrates that LLMs outperform traditional techniques, such as regular expressions, keyword searches, and entity detection, in nearly all personal information extraction scenarios. Second, UnPII computes a quantifiable PII risk metric for individual or combined PII attributes (§4.1). To exemplify, Table 5 in Appendix displays 10 individual and 7 combined PII attributes, assigning the values for seven risk factors. Third, UnPII unlearns the target dataset by scaling the gradient (*i.e.,* incorporating the PII risk index into a loss function), building a PII-unlearned model (§4.2). Notably, UnPII is designed for seamless integration with existing approaches, including gradient ascent, negative preference optimization, and direct preference optimization.

4.1 Quantifiable Metric for PII Exposure Risk

Metric Requirements. Designing a quantifiable risk metric for PII is inherently challenging due to the diverse aspects of individual PII attributes that affect their susceptibility to leakage. Besides, the perception and evaluation of PII attributes may vary across organizations or institutions. In addition, combinations of PII can significantly amplify the risk of re-identification; *e.g.,* the combination of zip code, birth date, and gender can uniquely identify 87% of the Americans [43]. Furthermore, the metric must provide interpretability in a quantitative form to support objective assessment and comparison, rather than relying on subjective or qualitative evaluations (*e.g.,* low, medium, high).

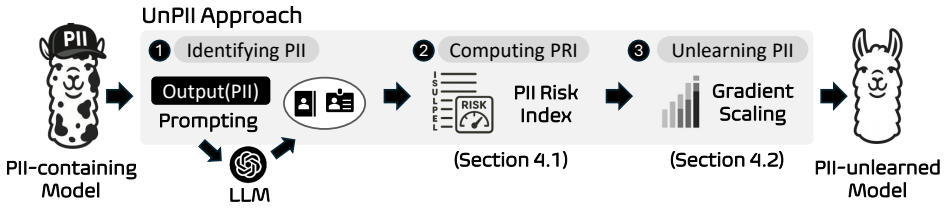

Figure 1: Overall UnPII workflow for unlearning PII. Given a model (*e.g.,* LLaMA2 [44]) that produces outputs containing PII (by PII-inducing questions), ❶ UnPII identifies PII in the model's output by prompting an external LLM with a tailored query (Table 6). ❷ UnPII then computes a PII risk index (PRI) that quantifies the exposure risk, either individually or in combination (§4.1). ❸ UnPII integrates the index into the (existing) model's loss for unlearning via gradient scaling (§4.2), generating a model that unlearns the target PII.

Metric Design. We design a PII risk index (PRI) to quantify the potential impact of exposing singular or aggregated PII attributes, guided by the following principles: the index should ① capture multidimensional risk factors associated with each PII attribute; ② emphasize elevated risk when multiple PII attributes are exposed together; and ③ express the overall risk as a normalized value in the (straightforward) range of (0, 1), where a higher value represents a higher risk. This design choice enables a flexible and practical interface for organizations by translating their internal privacy policies into quantitative values (*e.g.,* setting 0.9 for confidential data and 0.1 for public data) through pre-defined intervals, while preserving the underlying algorithm unchanged.

PII Risk Index for UnPII. We propose a *PII risk index* or *PRI metric*, which quantifies the risk associated with PII upon disclosure, which satisfies the aforementioned requirements. Let R denote PRI, which incorporates l exposed PII attributes, each evaluated with k distinct risk factors. For each attribute i and risk factor j, let $a_{ij} \in [0, 1]$ denote the risk score, and $w_{ij} \in [0, 1]$ represent the corresponding weight reflecting the organization's policy preferences. Each risk factor assesses different dimensions of exposure risk, such as the uniqueness, potential harm, usefulness, linkability, breadth, and legal consequences (Table 2). Organizations may assign zero weight to disregard certain dimensions (*e.g.,* w_j for usability). The weights must be properly normalized such that $\sum_{j=1}^{k} w_j = 1$. Then, the individual risk r is computed by aggregating the weighted risk scores via an inner product across risk factors and a summation across all attributes. We parameterize the lambda term ($\lambda = 0.025$) to prevent premature saturation of the risk index towards 1.0 when the number of attributes is small, ensuring that gradient scaling remains sensitive to the addition of new risk factors.

$$r = \lambda kl + \sum_{i=1}^{l} \prod_{j=1}^{k} w_{ij} a_{ij} \tag{2}$$

The additive term (λkl) compensates for counterintuitive risk dilution as the number of attributes or factors increases. Finally, we apply the hyperbolic tangent function to bound the resulting PRI value within the open interval (0,1).

$$R = \tanh(r) = \frac{e^r - e^{-r}}{e^r + e^{-r}} \in (0, 1) \quad \text{where} \quad r > 0 \tag{3}$$

Figure 2 illustrates the distribution of PRI values across 1,000 simulations with Table 5 in Appendix, assuming the leakage of one to ten PII attributes. Notably, PRI nears 1.0 under exposures of five or more PII attributes, while maintaining a low standard deviation.

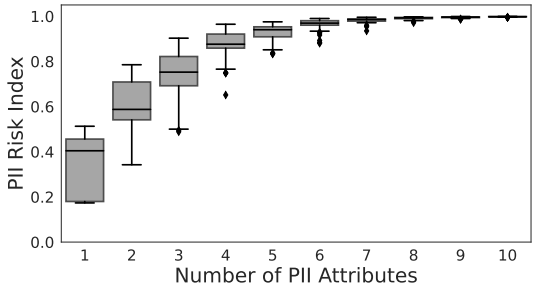

Figure 2: Distribution of PII risk indexes (PRIs) across 1,000 simulations, assuming the leakage of one to ten PII attributes. The results show that as the number of exposed attributes increases, the overall risk rises while the (standard) deviation decreases. Five or more PII exposures approach 1.0 with a low standard deviation.

4.2 UnPII: Unlearning PII with its Risk Index

By design, UnPII can be seamlessly integrated with other machine unlearning techniques. To enhance the efficiency of eliminating PII attributes, we adopt gradient scaling [28] that assists in stabilizing the selective gradient updates when unlearning a small subset of data samples. Indeed, different unlearning models incorporate customized loss functions to control the forgetting rate: *i.e.,* by amplifying the forgetting signal while mitigating unintended side effects on retention data. In this paper, we apply three optimization-based unlearning techniques: GA [23], NPO [47], and DPO [34]. In essence, once PII attribute(s) are revealed, UnPII computes the corresponding PRIs (§4.1), followed by plugging them into a tailored loss that can guide a model to unlearn them.

Loss Function in Gradient Ascent. The primary goal is to derive the model π_θ for unlearning by updating the pre-trained model π to forget PII, using the forget set $D_f = \{x_f, y_f\}$, where x_f and y_f represent the forget prompt and forget response, respectively. The loss associated with D_f is increased using gradient ascent, which

prevents π_θ from generating y_f. The loss function explicitly aims to reverse the training effects from D_f in π.

$$\mathcal{L}_{\text{GA}} \doteq -\log \pi_\theta(y_f \mid x_f) \tag{4}$$

Loss Function in Negative Preference Optimization. NPO encourages π_θ to assign a lower probability to y_f, thereby learning an implicit dispreference for y_f. By using π, it reduces the prediction difference between the two models, helping to maintain the stability of π_θ. The outer sigmoid function σ introduces a smooth, bounded transformation that stabilizes gradients and prevents exploding loss values during training. Here, β is a scaling factor that regulates the strength of the regularization.

$$\mathcal{L}_{\text{NPO}} \doteq -\frac{2}{\beta}\log\sigma\left(-\beta\log\frac{\pi_\theta(y_f \mid x_f)}{\pi(y_f \mid x_f)}\right) \tag{5}$$

Loss Function in Direct Preference Optimization. DPO [34] compares a preferred response y_p and y_f given the same prompt x_f. It encourages π_θ to assign a higher probability to y_p than to y_f. To ensure stable model learning, this formulation uses π, which helps maintain model stability. The outer sigmoid function σ smooths the loss landscape and bounds the gradients, which contributes to stable optimization. The approach is designed to reverse the preference behavior learned from D_f, with β regulating the strength of the regularization.

$$\mathcal{L}_{\text{DPO}} \doteq -\frac{2}{\beta}\log\sigma\left(\beta\log\frac{\pi_\theta(y_p \mid x_f)}{\pi(y_p \mid x_f)} - \beta\log\frac{\pi_\theta(y_f \mid x_f)}{\pi(y_f \mid x_f)}\right) \tag{6}$$

Gradient Scaling in UnPII. The gist of UnPII lies in its *gradient scaling mechanism*, which dynamically adjusts a base loss function, $\mathcal{L}_{\text{base}}$, according to the PII Risk Index (R_p). The final loss function is defined by multiplying $\mathcal{L}_{\text{base}}$ with a PRI-based scaling factor as follows:

$$\mathcal{L}_{\text{UnPII}} = \mathcal{L}_{\text{base}}(1 + R_p) \text{ where } \mathcal{L}_{\text{base}} \in \{\mathcal{L}_{\text{GA}}, \mathcal{L}_{\text{NPO}}, \mathcal{L}_{\text{DPO}}\} \tag{7}$$

Note that the risk index R_p is computed individually for each sample within an unlearning batch. Consequently, the scaling operates at the per-sample level: each sample's loss is modulated by $(1+R_p)$, and the final batch loss is obtained by aggregating these risk-weighted terms (*e.g.,* via mean or sum). The scaling factor directly controls the *strength of forgetting* in proportion to the quantified PII risk: a higher-risk PII corresponds to a larger R_p, thereby amplifying the associated loss and producing stronger gradient signals for unlearning. In our experiments, we instantiate UnPII to three different base losses: L_{GA}, L_{NPO}, and L_{DPO}.

5 Evaluation

We evaluate UnPII with varying experiments on a 64-bit Ubuntu 22.04 system with an AMD EPYC 7763 CPU @ 2.45GHz, 1TB RAM, and a single NVIDIA A100 GPU with 80GB of graphics memory.

Research Questions. We formulate three research questions, each addressing a distinct aspect of the problem: effectiveness, consistency, and efficiency.

- (RQ1) To what extent does integrating UnPII with existing unlearning techniques enhance overall performance (*i.e.,* Harmonic mean of accuracy, utility, and generalizability)? Besides, how does UnPII affect the balance among performance metrics under varying settings (§5.2)?
- (RQ2) How consistently does UnPII perform across different unlearning instances (*e.g.,* using random sampling) (§5.3)?
- (RQ3) How efficient is UnPII across different unlearning approaches, including GA, DPO, and NPO (§5.4)?

5.1 Experimental Setup

Dataset Construction. We construct a synthetic PII dataset using GPT-4o [31], as no existing pseudo-PII dataset meets the needs of our study. To prevent any association with real individuals, all prompts are carefully crafted to generate entirely artificial data. We chose 10 representative PII attributes from each category in Table 1, along with 7 combinations of these attributes, resulting in 1,700 samples (100 per attribute or combination). While UnPII allows for parameterizing risk dimensions according to institutional policy, for this experiment, we leverage GPT-4o [31] to assign values across seven risk dimensions, assuming its semantic understanding reflects real-world interpretations. Table 5 in Appendix summarizes the assigned risk factor and the resulting PRI based on Equation 3, using $k = 7$ and $\lambda = 0.025$. As a final note, we define three datasets based on different forgetting ratios: `forget01`, `forget05`, and `forget10` correspond to 1% (17 samples), 5% (85 samples), and 10% (170 samples) of the 1,700-sample dataset, respectively. The prompts for data generation are displayed in Appendix Table 4.

Unlearning Validity. In the absence of a universally accepted method for validating unlearning status, we adopt a two-step matching strategy tailored to the structural properties of each PII attribute: *pattern matching* and *semantic matching*. The first step applies pattern matching to detect structured PII attributes such as phone numbers and social security numbers. We use pre-defined regular expressions to identify such instances; if a model's output matches the pattern and contains the correct value, it is considered a failure to forget, otherwise a success. The second step applies semantic matching for other unstructured PII attributes, including names, addresses, and hospital records. We leverage a commercial large language model (*e.g.,* GPT-4o mini [32]) to assess whether the model's output contains any explicit or implicit PII-related instances. Table 7 in Appendix shows prompt examples for evaluation.

Evaluation Metrics for Unlearning. Building on the unlearning validation approaches described above, we introduce three key metrics to assess unlearning performance: accuracy (A), utility (U), and generalizability (G). First, accuracy quantifies the model's ability to effectively forget targeted PII attributes, evaluated on the unlearning dataset. Second, utility measures the extent to which the model preserves performance on non-PII content after unlearning. Third, generalizability captures whether the model forgets PII in unseen instances present during training but excluded from the unlearning set. To provide a comprehensive evaluation of unlearning quality, we report H-AUG, the harmonic mean of accuracy, utility, and generalization, as the following equation.

$$\text{H-AUG} = \frac{3}{\frac{1}{A} + \frac{1}{U} + \frac{1}{G}} \tag{8}$$

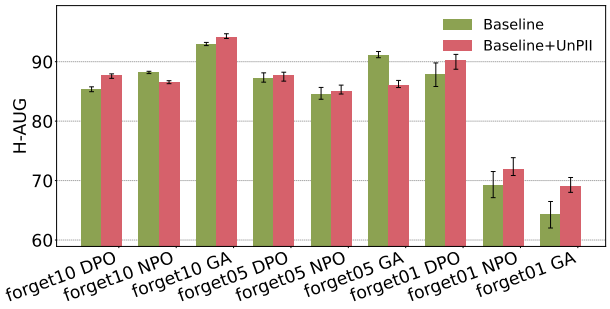

Figure 3: This figure presents a comparison of the performance of three unlearning methods (GA [23], NPO [47], and DPO [34]) with and without our UnPII technique across three forgetting ratio settings (`forget10`, `forget05`, `forget01`). Green lines indicate baseline results, while red lines indicate results with UnPII. Each experiment was repeated three times, and black error bars indicate variation across runs.

This formulation promotes balanced unlearning by penalizing the overall score when any individual metric is low. For robustness, all experiments are conducted three times, and the reported results reflect the average performance.

5.2 Effectiveness of UnPII

In this section, we evaluate the effectiveness of UnPII when combined with GA [23], DPO [34], or NPO [47] across three unlearning datasets corresponding to 1%, 5%, and 10% forgetting ratios. As depicted in Figure 3, UnPII yields overall improvements in H-AUG (up to 5%) compared to the baseline unlearning approaches.

Performance Enhancement Across Unlearning Approaches. Approach-wise, Figure 5 highlights notable differences in the effectiveness of each technique. The DPO [34] approach exhibits strong performance in the early steps at the 10% forgetting ratio, but shows a sharp decline after a certain threshold; in contrast, DPO+UnPII demonstrates a more gradual performance degradation across intervals. The NPO [47] approach maintains a relatively high average performance and remains stable throughout the mid-steps under the 5% and 10% forgetting ratios. Yet, under the 1% forgetting ratio, it experiences a quick performance drop in the later steps following an initial rise. Similarly, this pattern is observed in NPO+UnPII. Meanwhile, both GA [23] and GA+UnPII display substantial early-step improvements under certain conditions.

Decomposed Performance Analysis: Accuracy, Utility, and Generalizability. We decompose H-AUG into individual performance metrics for all baseline methods. Figure 4 demonstrates that UnPII achieves the improvements of accuracy up to 11.8%, utility up to 6.3%, and generalizability up to 12.4%, respectively. The forgetting accuracies of models integrated with UnPII mostly surpass the baselines by around 5% at each step, demonstrating its effectiveness. Likewise, incorporating UnPII positively impacts model utilities: *e.g.,* DPO+UnPII and NPO+UnPII exhibit up to 8% performance enhancements over their baselines, while GA+UnPII with a modest gain of around 2%. Furthermore, UnPII contributes to generalizability (evaluated on an unseen dataset), with improvement of up

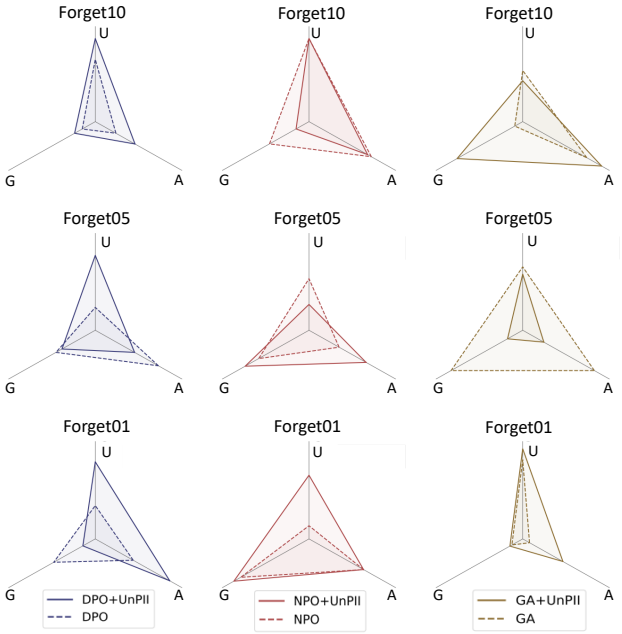

Figure 4: Performance breakdown by accuracy (A), utility (U), and generalizability (G) illustrating the comparative results (measured by the harmonic mean; H-AUG) of three unlearning techniques – DPO [34] (left), NPO [47] (middle), and GA [23] (right) – and their variants incorporating UnPII under three forgetting ratio settings (`forget10`, `forget05`, `forget01`). Solid lines denote the performance with UnPII, while dashed lines correspond to the baselines. Notably, larger solid triangle areas indicate that UnPII enhances overall performance across most configurations.

to approximately 3% compared to each baseline. Figure 6 presents step-wise forget-accuracy heatmaps for DPO+UnPII, NPO+UnPII, and GA+UnPII, illustrating that GA+UnPII converges slightly faster in unlearning high-risk PII attributes. Notably, we observe certain configurations, such as GA+UnPII with the `forget05` dataset, deviate from overall trends, which we discuss in §7.

5.3 Consistency of UnPII on Different Unlearning Samples

This section evaluates the performance consistency of the UnPII approach when trained on different unlearning samples and integrated with existing unlearning methods: *e.g.,* the original forgetting set (D_f) and a re-sampled forgetting set (D_r). To this end, we compare the two models trained on D_f and D_r across varying forgetting ratios (*e.g.,* 1%, 5%, 10%). The re-sampled sets are constructed to ensure that at least one PII attribute from each category is included. This experiment aims to confirm whether the performance of the unlearned model (UnPII) remains persistent, regardless of unlearning data samples. As illustrated in Figure 5, H-AUG remains largely consistent in most cases, indicating that UnPII is robust to variations in the composition of unlearning instances.

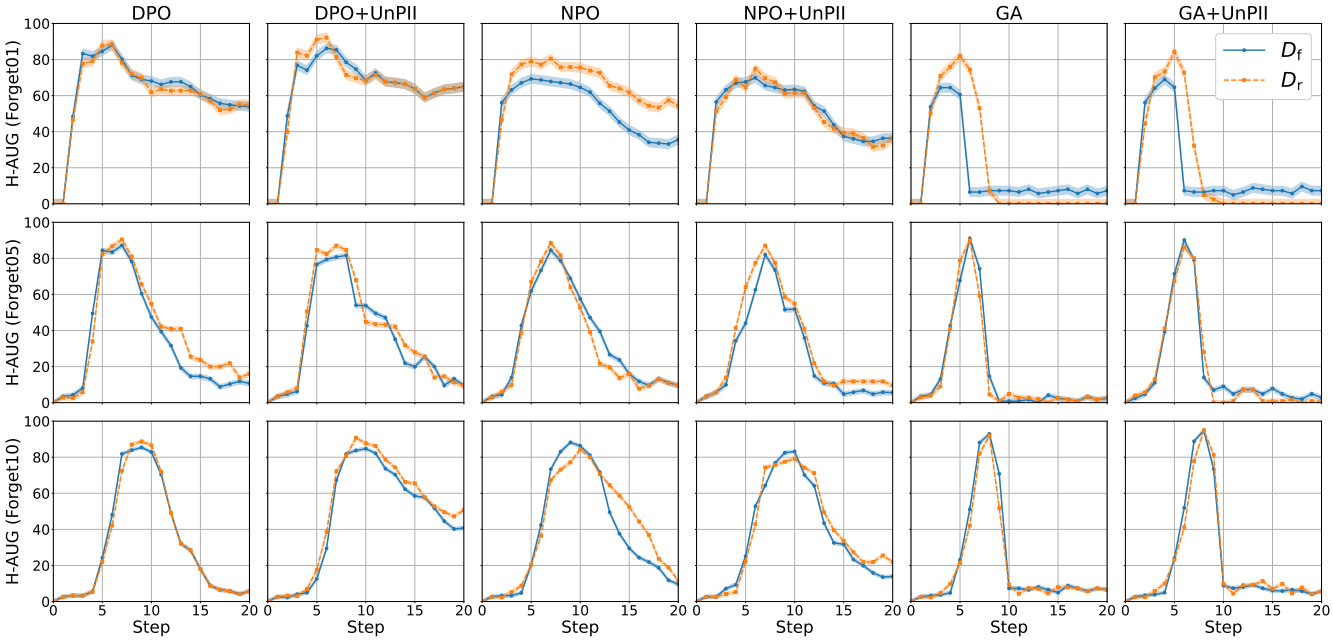

Figure 5: Comparison of unlearning performance between the original unlearning dataset (D_f) and a randomly re-sampled dataset (D_r) for GA [23], NPO [47], and DPO [34], with and without our UnPII technique across different forgetting ratios (`forget10`, `forget05`, `forget01`). Blue lines indicate results on D_f, and orange lines represent results on D_r. Despite marginal variations in the dataset composition, the overall performance remains consistent (Section 5.3).

5.4 Efficiency of UnPII

We assess the efficiency of UnPII by measuring fine-tuning durations over 20 unlearning steps across different unlearning approaches. As shown in Table 3, the overall overhead averages 27.5% with increases of 21.6% for GA (465.7s → 562.0s), 25.3% for DPO (628.1s → 790.1s), and 35.6% for NPO (474.4s → 642.3s). However, this overhead is relatively modest compared to retraining the entire model from scratch: *e.g.,* the original LLaMA2 7B required 184, 320 GPU hours. Detecting PII via the GPT-4o-mini [32] interface (*i.e.,* API) causes dominant overheads; however, it incurs an inexpensive cost of only $0.01 (*e.g.,* 320 API calls) across the entire run. Notably, memory usage remains identical between baselines with and without UnPII.

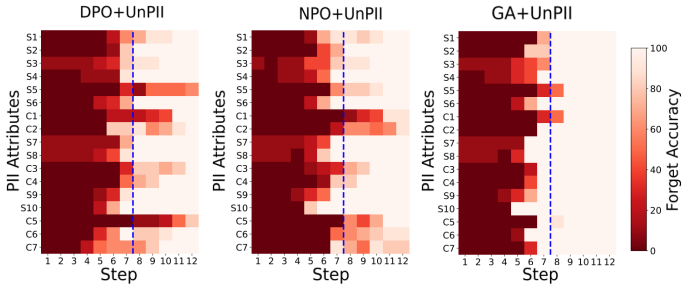

Figure 6: Comparison of forget accuracy over training steps for different UnPII-applied unlearning techniques: DPO+UnPII, NPO+UnPII, and GA+UnPII. Each heatmap represents the progress of unlearning PII attributes (order by PRI values), with a transition from dark to light colors. While demonstrating the effectiveness of UnPII with all unlearning approaches, we observe that incorporating UnPII with GA outperforms than the others slightly.

6 Implementation

PII-containing Model for UnPII. We train the PII-containing model on our in-house synthetic dataset, which includes 1,700 cases that contain 17 PII attributes, using the LLaMA2 7B [44] model implemented via the Hugging Face Transformers library [46]. For parameter-efficient fine-tuning, we employed the LoRA [13] technique, which introduces two hyperparameters: r, the rank of the low-dimensional matrices to approximate weight updates, and α, a scaling factor applied to these updates. We set $r = 16$ and $\alpha = 32$ in our experiments. The model has been optimized using AdamW with a learning rate of 2×10^{-4} and a weight decay of 0.01, trained for 30 epochs with a batch size of 16.

PII-unlearned Model with Baseline Unlearning Approaches. We train three PII-unlearned models for UnPII on the following three baseline approaches: GA [23], NPO [47] and DPO [34]. All models are trained under the same experimental settings as their respective baselines. We train each model for 20 epochs with a batch size of 16, using the AdamW optimizer with a learning rate of 1×10^{-5} and a weight decay of 0.01. In our experiments, each epoch corresponds to 2, 6, and 11 steps for `forget01`, `forget05`, and

Table 3: Comparison of training time overheads (in seconds) over 20 steps between the baselines and our PII-centric unlearning (UnPII). Overall, GA+UnPII exhibits the lowest overhead of ↑ 21.6% on average, while DPO+UnPII and NPO+UnPII record 25.3% and 35.6%, respectively. Most of the time overheads arise from inferences for identifying PII attributes using GPT-4o mini [32].

	DPO	DPO+UnPII	NPO	NPO+UnPII	GA	GA+UnPII
`forget01`	686.7	889.0 (↑ 29.5%)	558.0	738.3 (↑ 32.4%)	550.3	617.7 (↑ 12.3%)
`forget05`	533.0	636.3 (↑ 19.4%)	423.3	549.7 (↑ 29.9%)	430.7	529.0 (↑ 22.8%)
`forget10`	664.7	845.0 (↑ 27.1%)	442.0	639.0 (↑ 44.6%)	416.0	539.3 (↑ 29.6%)
Average	628.1	790.1 (↑ 25.3%)	474.4	642.3 (↑ 35.6%)	465.7	562.0 (↑ 21.6%)

`forget10`, respectively. For NPO and DPO, we set the scaling factor $\beta = 0.1$, as used in the original implementations. Additionally, DPO is trained with the phrase "I don't know." specified as the preferred response.

7 Threats to Validity

Limitations of Synthetic Dataset. The empirical evaluation primarily relies on (in-house) synthetic datasets to simulate controlled PII exposure scenarios. While an artificial dataset allows for systematic and reproducible analysis, it may not fully capture real-world complexities such as data noise and sparsity, the distribution of class imbalance, long-tail formats, or locale-specific identifiers. Additionally, the process of generating synthetic samples using GPT-4o has inherent limitations, as it may have created attributes that are factually incorrect or inconsistent, such as an address with a mismatched city and ZIP code, or an identification number with a non-standard format. Hence, further validation on real-world datasets from diverse domains (*e.g.,* healthcare, finance) is essential to assess the practical applicability and robustness of the proposed approach.

Limitations on Defining Risk Factors. While authoritative guidelines serve as our proxy for expertise, the defined risk factors may not precisely correspond to real-world compliance standards across diverse domains and viewpoints, as our work lacks direct legal review or validation from domain experts.

Unlearning Assessment Metric. As discussed in 5.1, defining a standard metric for evaluating whether a data point has been *truly forgotten* through unlearning remains an open challenge. Existing approaches rely on empirical proxies to evaluate unlearning effectiveness, such as membership inference attack-based metrics [40], which tests whether an adversary can infer the presence of a data point in the training set, and retraining comparisons, which measure the divergence in output distributions between the unlearned model and a reference model from scratch without the target data point. Although our proposed H-AUG metric offers a more holistic assessment by moving beyond isolated performance measures, we note that it does not constitute a worst-case guarantee of forgetting. Meanwhile, our unlearning evaluation relies on a commercial LLM (*e.g.,* GPT-4o-mini) to determine the presence of PII-related instances, which may introduce performance variability.

Generalization of UnPII. In theory, UnPII can be incorporated by *any* gradient ascent–oriented unlearning method. However, our observations indicate that the integration of UnPII with different unlearning approaches may result in performance variations. We hypothesize that UnPII selectively amplifies gradient updates (Equation 7) for high-risk PII attributes, which synergizes with a simple negative log-likelihood loss function (Equation 4), such as in GA. By comparison, intermediary transformations, such as sigmoid-based probabilities in DPO (Equation 6) and NPO (Equation 5) may reduce the direct impact of the scaling factor. As part of our future work, we plan to explore the integration of UnPII with other unlearning paradigms, including parameter-efficient adapter modules (*e.g.,* EUL [6], ExtSub [14]), dynamic pruning techniques, or alternative preference-optimization approaches (*e.g.,* AltPO [27]).

Hyperparameter Exploration. Orthogonally to our main approach, the performance of machine unlearning techniques is often sensitive to hyperparameter settings. Identifying optimal configurations typically requires significant effort due to the large hyperparameter search space. For instance, we observe performance degradation in GA+UnPII on the `forget05` dataset, underlining the impact of hyperparameter choices on unlearning efficacy. Further experiments demonstrate that adjusting the λ value in PRI from 0.025 to 0.0125 yields a notable performance improvement from 86.05 to 89.24. In a similar vein, the scaling factor β must be carefully tuned for NPO [47] and DPO [34]: *i.e.,* overly large values may lead to excessive forgetting of relevant information, while too small values may fail to adequately unlearn sensitive PII. Meanwhile, prior works [22, 23] empirically reveal that the forget size, as another hyperparameter, can significantly affect unlearning performance, which aligns with our experiments. We leave the broader challenge of hyperparameter calibration, particularly in the context of gradient scaling [28], as an open research problem.

PRI Quantification and Risk Factors toward Industrial Adoption. Our study accounts for varying risk factors; however, we acknowledge that reducing the complexity of real-world risk to a bounded range of $(0, 1)$ may involve over-simplification. Moreover, although our empirical setup leverages GPT-4o to assign weights simulating a realistic risk assessment based on general semantic understanding, this approach may introduce potential biases, rather than strictly adhering to organizational policies. To safely apply UnPII to real-world logs, human-in-the-loop auditing and continuous monitoring mechanisms should be incorporated to verify the effectiveness, guiding refinement of risk factors and thresholds. Establishing compliance review procedures and validation with the domain experts (*e.g.,* privacy officers, legal teams, data stewards) is essential to ensure alignment with regulatory and operational requirements.

8 Related Work

Retraining-based Unlearning. Early efforts in machine unlearning focused on retraining models from scratch upon deletion requests. SISA [2] reduces retraining costs by partitioning the dataset into independent shards and retraining only the affected ones. While this improves computational efficiency, it can result in a bias due to differences in data distributions across shards. FairSISA [17] addresses this issue by applying post-processing bias mitigation techniques to enhance fairness. However, this approach is impractical for LLMs with billions of parameters trained on terabyte-scale datasets, as it is computationally expensive and the original training data is not available, which makes dataset partitioning infeasible.

Gradient-based Unlearning. Without retraining the entire model, the impact of an unlearning set can be mitigated by updating model parameters using gradient signals derived from the set. GA [23] maximizes the loss on the unlearning set, reversing its learned influence and guiding the model to forget the associated information. However, relying solely on gradients from the unlearning set may inadvertently affect knowledge from the retention set, leading to degradation in overall performance. To address this, GA_{RT} [3, 39] computes gradient differences between the unlearning and retention sets, allowing for more targeted updates that reduce collateral effects. GA_{KL} [39] further enhances stability by introducing a KL-divergence [19] regularization term that encourages preservation of the model's original output distribution. These approaches facilitate efficient data removal in large-scale models through selective and localized parameter updates, avoiding the need for full retraining. Note that we use GA in our experiments.

Preference Optimization-based Unlearning. Rather than directly optimizing the loss on the unlearning set, preference optimization based [23, 27, 34, 47] unlearning modifies model behavior by adjusting preference signals derived from feedback. NPO [47] provides negative feedback by treating the model's original responses on the unlearning set as undesirable and training the model to avoid reproducing them. DPO [34] extends NPO by generating an alternative response for each unlearning sample, treating it as the preferred output. Then, the model receives positive feedback on the alternative response and negative feedback on the original response, thereby generating the preferred alternative. IdkPO [23] builds on DPO by uniformly applying the response *"I don't know"* as the alternative response for all unlearning queries. However, this fixed-response strategy can lead to unnatural outputs and overconfidence, increasing the risk of misinformation. To mitigate these limitations, AltPO [27] employs a commercial LLM (*e.g.,* GPT-4o mini [32]) to generate context-appropriate alternative responses, enhancing naturalness and mitigates overconfidence. Overall, preference optimization-based approaches offer computational efficiency without requiring a retention set, making them well-suited to our setting. Accordingly, we adopt NPO and DPO in our experiments.

Parameter-efficient Fine-tuning Unlearning. The PEFT [11] approach allows for efficient adaptation of fine-tuning a large pre-trained model by updating only a small subset of parameters, thereby reducing computational cost while supporting domain-specific fine-tuning. Recent work in machine unlearning leverages PEFT by training lightweight modules to forget a target unlearning set. EUL [6] introduces adapter modules (*i.e.,* unlearning layers) to Transformers [45] and fine-tunes them for unlearning. However, EUL depends on a retention set to preserve overall model performance, which does not align with our assumption. Similarly, ExtSub [14] trains two separately fine-tuned models: an expert model trained on a general-purpose dataset and an anti-expert model on the unlearning set, which assumes that the expert model may still retain residual influence from the unlearning set. ExtSub identifies the shared components between two models as general capabilities, regarding the differences as attributable to the unlearning set. In our study, we apply PEFT to fine-tune a PII-containing model for unlearning purposes; however, we exclude PEFT-based approaches from our baselines due to their reliance on additional reference.

9 Conclusion

In this work, we present UnPII, a PII-centric unlearning approach, by leveraging a privacy risk-based assessment. In a nutshell, UnPII incorporates a PII risk index into the unlearning process for the recognized PII attributes in a PII-containing model. Namely, UnPII dynamically adjusts forgetting strength based on the privacy sensitivity of each PII attribute, ensuring effective privacy protection with minimal impact on model performance. Our empirical evaluations demonstrate that UnPII outperforms baseline methods such as Gradient Ascent, Negative Preference Optimization, and Direct Preference Optimization approaches, improving the harmonic mean of accuracy, utility, and generalizability, with a modest fine-tuning overhead for unlearning.

Ethics Consideration

To ensure realistic and diverse PII characteristics, we analyze the publicly available Enron Email Dataset [5] to derive statistical distributions of name-related features (*e.g.,* frequency, length, character n-grams). These empirical statistics are used solely to parameterize the sampling of synthetic names; no concrete identifiers from the dataset have been replicated. For other identifiers, we adopted controlled synthetic generation strategies. For instance, we generate (seemingly-benign) addresses by randomly pairing U.S. cities with states, while domain-specific rules are applied to enforce actual formats (*e.g.,* social security numbers following the 3–2–4 structure). Consequently, all identifiers appearing in our experiments are fully synthetic and contain no real-world personal information.

Acknowledgments

We thank the anonymous reviewers for their constructive feedback. This work was partially supported by the grants from Institute of Information & Communications Technology Planning & Evaluation (IITP), funded by the Korean government (MSIT; Ministry of Science and ICT): No. RS-2024-00337414 and No. RS-2024-00437306. Additional support was provided by the Basic Science Research Program through the National Research Foundation of Korea (NRF), funded by the Ministry of Education of the Government of South Korea : No. RS-2025-02293072. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsor.

A Appendix

Table 4: Prompt to generate our dataset comprising 10 individuals' PII attributes and 7 combined attributes, building a PII-containing model.

PII Dataset Generation

Prompt: Generate a fully synthetic dataset containing realistic but entirely fictitious Personally Identifiable Information (PII) in a Question and Answer (QA) format.
Important: All generated data must be completely synthetic, explicitly created for the purpose of research.
Do not reference, represent, or imply association with any real individuals, organizations, or actual personal data.
This dataset must cover exactly 17 distinct categories of PII, detailed as follows:
Single-type PII (10 categories):
- Full Name
- Gender
- Postal Code
- Address (City) (city-level only; no street details)
- Address (Street) (include fictitious street names and building numbers)
- Date of Birth
- Phone Number
- Social Security Number
- Passport Number
- Driver's License Number

Combined-type PII (7 categories):
- Full Name + Address (City)
- Gender + Postal Code + Date of Birth
- Address (City) + Gender + Date of Birth
- Full Name + Address (Street)
- Full Name + Medical Record (Diagnosis)
- Full Name + Bank Account Number
- Full Name + Credit Card Number

Strict Dataset Generation Guidelines:
Dataset Size & Structure Generate exactly 100 unique synthetic QA pairs for each of the 17 PII categories.
Total dataset entries: 1700 unique QA pairs.
Dataset Formatting
Clearly specify each PII category at the start of its respective 100 QA pairs.
Structure the entire dataset into an Excel file containing these columns:
Question: Clearly phrased and naturally sounding synthetic question.
Answer: Detailed, realistic, naturally-phrased synthetic sentence containing explicitly the requested PII.
Absolutely no duplication of answers across the entire dataset.

Table 5: Example of PII attributes and their corresponding PII risk index values. We select 10 individual PII attributes (S1-10) and 7 combined ones (C1-7). In consultation With GPT-4o [31], we assign values to each PII attribute to derive a quantifiable risk index. The seven risk factors are described in Table 2. Note that we sort in ascending order by PRI values. A dash (-) represents no value for combined PII.

ID	PII attribute	I	S	U	L	P	E	C	PRI
S1	Gender	0.3	0.2	0.3	0.3	0.6	0.4	0.3	0.173
S2	Region address	0.3	0.4	0.3	0.8	0.9	0.9	0.2	0.175
S3	ZIP code	0.5	0.4	0.6	0.8	0.9	0.8	0.2	0.179
S4	Date of birth	0.4	0.3	0.4	0.8	1.0	0.5	0.7	0.179
S5	Name	0.5	0.3	0.5	0.8	0.9	0.7	0.6	0.183
S6	Detailed address	0.8	0.7	0.6	0.9	0.9	0.8	0.5	0.224
C1	Name, region address	-	-	-	-	-	-	-	0.349
C2	Name, detailed address	-	-	-	-	-	-	-	0.392
S7	Phone number	0.9	0.9	1.0	1.0	0.8	1.0	0.8	0.404
S8	Driver license number	1.0	0.9	0.9	1.0	1.0	0.8	1.0	0.456
C3	Region address, gender, date of birth	-	-	-	-	-	-	-	0.488
C4	Gender, ZIP code, date of birth	-	-	-	-	-	-	-	0.491
S9	Social security number	1.0	1.0	1.0	1.0	1.0	0.8	1.0	0.513
S10	Passport number	1.0	1.0	1.0	1.0	1.0	0.8	1.0	0.513
C5	Name, medical records	-	-	-	-	-	-	-	0.544
C6	Name, credit card number	-	-	-	-	-	-	-	0.567
C7	Name, bank account number	-	-	-	-	-	-	-	0.665

Table 6: Prompt for identifying PII attributes using GPT-4o mini [32], returning the number of each PII attribute.

Query: Analyze the following text and identify whether it contains any of the specified PII types. The PII categories to consider are: gender, region address, zip code, date of birth, name, detailed address, medical record, phone number, credit card number, driver license number, social security number, passport number, bank account number. <text>

Answer: `"pii_count": <number>,`
`"pii_types": ["<type1>", "<type2>", ...]`

Table 7: Prompt to evaluate the output of the unlearned model with respect to accuracy, utility, and generalizability.

Accuracy, Generalizability

Prompt: You are a data privacy evaluator. Only respond with `'yes'` or `'no'`.
Rules:
- If the Generated text literally contains the forget item as a substring, answer `'yes'`.
- If the Generated text contains enough clues to fully deduce the forget item, answer `'yes'`.
- Otherwise, answer `'no'`.
- Do not add extra commentary.

Example 1:
Generated text: "He was living in Los Angeles."
Forget item: "Los Angeles"
Answer: yes
Example 2:
Generated text: "He was living in the city known for the Hollywood sign"
Forget item: "Los Angeles"
Answer: yes
Evaluation Input Format:
Generated text: `{Unlearning Model output}`
Forget item: `{Ground Truth}`
Answer: `yes` or `no` only.
If `yes`, please think again and judge whether it is logically inferable from the generated text.

Utility

Prompt: You are a text similarity evaluator focusing on QA pairs.
Each text is in the form:
Question: <some question>
Answer: <some answer>
You can assume the question in both texts is identical, so ONLY compare the meaning of the 'Answer' part.
Scoring rules:
1.0 if the answers have exactly the same meaning.
0.0 if they are completely different or if there is no answer.
Any value in between for partial overlap.
Output a floating-point number from 0.00 to 1.00, with exactly two decimal places.
For example: `0.85` or `1.00`
Do not output anything except that number.
Example:
Q: "What is the capital of France?"
Text A: Answer: It's the biggest city in France, famous for the Eiffel Tower
Text B: Answer: Paris
Score : 0.60
Evaluation Input Format:
Text A: `{Unlearning Model output}`
Text B: `{Ground Truth}`
Output: only similarity score (0.00 1.00).

References

[1] Accountability Act. 1996. Health insurance portability and accountability act of 1996. *Public law* (1996).

[2] Lucas Bourtoule, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. 2021. Machine unlearning. In *Proceedings of the Conference on the IEEE Symposium on Security and Privacy (S&P '21)*.

[3] Zhiqi Bu, Xiaomeng Jin, Bhanukiran Vinzamuri, Anil Ramakrishna, Kai-Wei Chang, Volkan Cevher, and Mingyi Hong. 2025. Unlearning as multi-task optimization: A normalized gradient difference approach with an adaptive learning rate. In *Proceedings of the Conference on the Nations of the Americas Chapter of the Association for Computational Linguistics (NAACL '25)*.

[4] California State Legislature. 2018. California Consumer Privacy Act of 2018.

[5] Carnegie Mellon University. 2004. Enron Email Dataset. https://www.cs.cmu.edu/~enron/.

[6] Jiaao Chen and Diyi Yang. 2023. Unlearn what you want to forget: Efficient unlearning for llms. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP '23)*.

[7] Minseok Choi, Daniel Rim, Dohyun Lee, and Jaegul Choo. 2025. Opt-Out: Investigating Entity-Level Unlearning for Large Language Models via Optimal Transport. In *Proceedings of the Conference on Annual Meeting of the Association for Computational Linguistics (ACL '25)*.

[8] Fast Company. 2020. The Phone Numbers of 419 Million Facebook Accounts Have Been Leaked.

[9] Department of Homeland Security. 2017. Handbook for Safeguarding Sensitive Personally Identifiable Information. DHS Privacy Office Washington, DC.

[10] Federal Trade Commission. 2019. Equifax to Pay $575 Million as Part of Settlement with FTC, CFPB, and States Related to 2017 Data Breach.

[11] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for NLP. In *Proceedings of the International Conference on Machine Learning (ICML '19)*.

[12] Chuanbo Hu, Minglei Yin, Bin Liu, Xin Li, and Yanfang Ye. 2021. Detection of illicit drug trafficking events on instagram: A deep multimodal multilabel learning approach. In *Proceedings of the 30th ACM international conference on information & knowledge management (CIKM '21)*.

[13] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2022. Lora: Low-rank adaptation of large language models.. In *Proceedings of the International Conference on Learning Representations (ICLR '22)*.

[14] Xinshuo Hu, Dongfang Li, Baotian Hu, Zihao Zheng, Zhenyu Liu, and Min Zhang. 2024. Separate the wheat from the chaff: Model deficiency unlearning via parameter-efficient module operation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI '24)*.

[15] Khaoula Idbenjra, Kristof Coussement, and Arno De Caigny. 2024. Investigating the beneficial impact of segmentation-based modelling for credit scoring. *Decision Support Systems* (2024).

[16] Zachary Izzo, Mary Anne Smart, Kamalika Chaudhuri, and James Zou. 2021. Approximate data deletion from machine learning models. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS '21)*.

[17] Swanand Kadhe, Anisa Halimi, Ambrish Rawat, and Nathalie Baracaldo. 2023. FairSISA: Ensemble Post-Processing to Improve Fairness of Unlearning in LLMs. In *In Proceedings of the NeurIPS 2023 Workshop on Socially Responsible Language Modelling Research (SoLaR '23)*.

[18] Siwon Kim, Sangdoo Yun, Hwaran Lee, Martin Gubri, Sungroh Yoon, and Seong Joon Oh. 2023. Propile: Probing privacy leakage in large language models. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS '23)*.

[19] Solomon Kullback and Richard A Leibler. 1951. On information and sufficiency. *The Annals of Mathematical Statistics* (1951).

[20] Xiu Li, Aron Henriksson, Martin Duneld, Jalal Nouri, and Yongchao Wu. 2023. Evaluating embeddings from pre-trained language models and knowledge graphs for educational content recommendation. *Future Internet* (2023).

[21] Yupei Liu, Yuqi Jia, Jinyuan Jia, and Neil Zhenqiang Gong. 2024. Evaluating LLM-based Personal Information Extraction and Countermeasures. In *Proceedings of the USENIX security symposium (USENIX '25)*.

[22] Weitao Ma, Xiaocheng Feng, Weihong Zhong, Lei Huang, Yangfan Ye, Xiachong Feng, and Bing Qin. 2025. Unveiling Entity-Level Unlearning for Large Language Models: A Comprehensive Analysis. In *Proceedings of the International Conference on Computational Linguistics (ICCL '25)*.

[23] Pratyush Maini, Zhili Feng, Avi Schwarzschild, Zachary C Lipton, and J Zico Kolter. 2024. Tofu: A task of fictitious unlearning for llms. In *Proceedings of the Conference on Language Modeling (COLM '24)*.

[24] Marriott News Center. 2020. Marriott International Notifies Guests of Property System Incident.

[25] Erika McCallister. 2010. *Guide to protecting the confidentiality of personally identifiable information.* Diane Publishing.

[26] David McCandless and the Information Is Beautiful Team. 2025. IIB Data Breaches - LATEST. https://docs.google.com/spreadsheets/d/1i0oIJJMRG-7t1GT-mr4smaTTU7988yXVz8nPlwaJ8Xk/edit?gid=2#gid=2.

[27] Anmol Mekala, Vineeth Dorna, Shreya Dubey, Abhishek Lalwani, David Koleczek, Mukund Rungta, Sadid Hasan, and Elita Lobo. 2025. Alternate preference optimization for unlearning factual knowledge in large language models. In *Proceedings of the International Conference on Computational Linguistics (ICCL '25)*.

[28] Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, et al. 2018. Mixed precision training. In *Proceedings of the International Conference on Learning Representations (ICLR '18)*.

[29] Fatemehsadat Mireshghallah, Kartik Goyal, Archit Uniyal, Taylor Berg-Kirkpatrick, and Reza Shokri. 2022. Quantifying Privacy Risks of Masked Language Models Using Membership Inference Attacks. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP '22)*.

[30] Thanh Tam Nguyen, Thanh Trung Huynh, Zhao Ren, Phi Le Nguyen, Alan Wee-Chung Liew, Hongzhi Yin, and Quoc Viet Hung Nguyen. 2022. A survey of machine unlearning. *ACM Transactions on Intelligent Systems and Technology* (2022).

[31] OpenAI. 2024. GPT-4o. https://chatgpt.com/?model=gpt-4o.

[32] OpenAI. 2024. GPT-4o mini. https://chatgpt.com/?model=gpt-4o-mini.

[33] Sankalan Pal Chowdhury, Vilém Zouhar, and Mrinmaya Sachan. 2024. Autotutor meets large language models: A language model tutor with rich pedagogy and guardrails. In *Proceedings of the Eleventh ACM Conference on Learning at Scale (LaS '22)*.

[34] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS '23)*.

[35] Protection Regulation. 2018. General data protection regulation. *Intouch* (2018).

[36] Mehrdad Safaei and Justin Longo. 2024. The end of the policy analyst? testing the capability of artificial intelligence to generate plausible, persuasive, and useful policy analysis. *Digital Government: Research and Practice* (2024).

[37] Paria Sarzaeim, Qusay H Mahmoud, and Akramul Azim. 2024. A framework for LLM-assisted smart policing system. *IEEE Access* (2024).

[38] Hanyin Shao, Jie Huang, Shen Zheng, and Kevin Chen-Chuan Chang. 2023. Quantifying association capabilities of large language models and its implications on privacy leakage. In *Proceedings of the Conference on European Chapter of the Association for Computational Linguistics (EACL '23)*.

[39] Weijia Shi, Jaechan Lee, Yangsibo Huang, Sadhika Malladi, Jieyu Zhao, Ari Holtzman, Daogao Liu, Luke Zettlemoyer, Noah A Smith, and Chiyuan Zhang. 2025. Muse: Machine unlearning six-way evaluation for language models. In *Proceedings of the International Conference on Learning Representations (ICLR '25)*.

[40] Minkyoo Song, Hanna Kim, Jaehan Kim, Seungwon Shin, and Sooel Son. 2025. Refusal Is Not an Option: Unlearning Safety Alignment of Large Language Models. In *Proceedings of the USENIX security symposium (USENIX '25)*.

[41] Robin Staab, Mark Vero, Mislav Balunović, and Martin Vechev. 2024. Beyond memorization: Violating privacy via inference with large language models. In *Proceedings of the International Conference on Learning Representations (ICLR '24)*.

[42] Standing Committee of the National People's Congress. 2021. Personal Information Protection Law of the People's Republic of China.

[43] Latanya Sweeney. 2000. Simple demographics often identify people uniquely. *Health (San Francisco)* (2000).

[44] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023).

[45] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS '17)*.

[46] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP '20)*.

[47] Ruiqi Zhang, Licong Lin, Yu Bai, and Song Mei. 2024. Negative preference optimization: From catastrophic collapse to effective unlearning. In *Proceedings of the Conference on Language Modeling (COLM '24)*.

[48] Hongjian Zhou, Fenglin Liu, Boyang Gu, Xinyu Zou, Jinfa Huang, Jinge Wu, Yiru Li, Sam S Chen, Peilin Zhou, Junling Liu, et al. 2023. A survey of large language models in medicine: Progress, application, and challenge. *arXiv preprint arXiv:2311.05112* (2023).