

LEVERAGE UNLEARNING TO SANITIZE LLMs

Anonymous authors

Paper under double-blind review

ABSTRACT

Pre-trained large language models (LLMs) are becoming useful for various tasks. To improve their performance on certain tasks, it is necessary to fine-tune them on specific data corpora (e.g., medical reports, business data). These specialized data corpora may contain sensitive data (e.g., personal or confidential data) that will be memorized by the model and likely to be regurgitated during its subsequent use. This memorization of sensitive information by the model poses a significant privacy or confidentiality issue. To remove this memorization and sanitize the model without requiring costly additional fine-tuning on a secured data corpus, we propose SANI. SANI is an unlearning approach to sanitize language models. It relies on both an erasure and repair phases that 1) reset certain neurons in the last layers of the model to disrupt the memorization of fine-grained information, and then 2) fine-tune the model while avoiding memorizing sensitive information. We comprehensively evaluate SANI to sanitize both a model fine-tuned and specialized with medical data by removing directly and indirectly identifiers from the memorization of the model, and a standard pre-trained model by removing specific terms defined as confidential information from the model. Results show that with only few additional epochs of unlearning, the model is sanitized and the number of regurgitations is drastically reduced. This approach can be particularly useful for hospitals or other industries that have already spent significant resources training models on large datasets and wish to sanitize them before sharing.

1 INTRODUCTION

Recent advances in NLP Singh & Mahmood (2021) based on neural networks have democratized their use Chen et al. (2022c); Yeo et al. (2022); Wei et al. (2022); Goyal et al. (2024). Since the advent of ChatGPT, NLP models are not limited to text generation and can include several downstream tasks, such as classification tasks Zhang et al. (2024) or Named Entity Recognition (NER) Zhao et al. (2024). Pre-trained large language models (LLMs) such as BERT or GPT have significantly contributed to their adoption across all sectors. These off-the-shelf models were pre-trained at great expense on countless unlabeled datasets web-scraped from the Internet. These pre-trained foundation models can then be fine-tuned on specific datasets and even trained for new tasks by adding a few additional layers. For example, a hospital will fine-tune a BERT model by leveraging its medical reports before training it with an additional layer for a downstream classification task.

The need for information sharing extends beyond data and now concerns learning models. For example, some hospitals also want to share specialized models fine-tuned on their own data with other medical centers.

The attack surface on models is still poorly understood (Weidinger et al. (2022); Lehman et al. (2021); Duprieu & Berkouk (2024)). A number of threats are related to the memorization and possible leakage of sensitive information used during model training, such as data reconstruction and membership inference (i.e., identifying elements used during the training or the fine-tuning). For example, the European Data Protection Board (EDPB) issued an opinion on the development and deployment of AI models in light of the GDPR¹. Specifically, the opinion addresses the issue of AI models trained with personal data and stipulates that to be considered anonymous, the likelihood

¹<https://www.edpb.europa.eu/news/news/2024/edpb-opinion-ai-models-gdpr-principles-support-r-en>

054 of regurgitation of personal information must be insignificant, taking into account all the means
055 reasonably likely to be used.

056 Therefore, it is necessary to ensure that the model is sanitized from any sensitive information mem-
057 orized before sharing it. This sensitive information may be personal information or other confi-
058 dential information (e.g., related to intellectual property). Since training or specializing a model
059 requires significant computational resources and is time-consuming, retraining or re-specializing
060 from scratch would be prohibitively expensive. To overcome this limitation, we develop SANI, a
061 system that sanitizes LLMs by removing sensitive information from the memorization of the model.
062 SANI leverages a unlearning scheme to remove specific information (i.e., information present in a
063 blacklist) from the model at a reduced cost. Specifically, SANI adapts an erase and repair strategy
064 (Kurmanji et al. (2023b)) where the former phase reinitializes some neurons in the last layer (i.e.,
065 capturing and disrupting only fine-grained memorization) without affecting the model foundations
066 memorized in the lower layers of the architecture, while the latter phase restores and specializes
067 models while avoiding the memorization of terms defined in a blacklist (Boutet et al. (2025)).

068 We conducted an empirical evaluation of SANI by measuring its ability to unlearn personal informa-
069 tion memorized in a pre-trained LLMs (i.e., BERT and GPT models) fine-tuned with medical data,
070 and to unlearn information defined as confidential exposed to a LLM during its training. We demon-
071 strate that SANI is able to quickly remove sensitive information from the model without altering its
072 performance for word prediction or generation, as well as for downstream classification tasks. Fur-
073 thermore, we show that the sensitive information the most exposed to the model is the more affected
074 by the unlearning, drastically reducing the number of regurgitations. Finally, compared against dif-
075 ferent baselines, we show that SANI offers the best trade-off between reducing sensitive information
076 regurgitation and model performance.

078 2 BACKGROUND AND RELATED WORK

080 Large language models (LLMs) are trained on very large datasets. For example, training chatGPT
081 is very costly and required years of crawling the Internet. LLMs (i.e., foundation models) are gener-
082 ally optimized for specific tasks with fine-tuning using domain-specific data. In the medical field, for
083 example, these models are specialized with patient reports, which improves performance for down-
084 stream tasks (Huang et al. (2020)). Without specific measures, these models can therefore memorize
085 and subsequently regurgitate and disclose potentially sensitive information from the training data,
086 which constitutes a significant privacy leak and has implications for the practical and legal aspects of
087 these models. The same problem arises with a model exposed to confidential data during its training,
088 and which could then be led to regurgitate it during its exploitation. To avoid training models once
089 again from scratch (by excluding sensitive information from the beginning), unlearning aims to for-
090 get or removing specific data (e.g., sensitive, unreliable, or copyrighted data) from the memorisation
091 of the model. This section presents a comprehensive background and related work on unlearning
092 (Section 2.1), and on the risks of memorization of training data by the models (Section 2.2).

093 2.1 MACHINE UNLEARNING

095 Machine Unlearning (Cao & Yang (2015); Xu et al. (2023); Nguyen et al. (2024); Kurmanji et al.
096 (2023b)) has rapidly emerged as a significant area of research within the field of Machine Learning
097 (ML). A first approach (Bourtoule et al. (2020)) introduces a framework that utilizes data shading
098 and slicing techniques to minimize the computational overhead associated with unlearning. Data
099 shading involves modifying data representations to obscure sensitive information, while slicing
100 refers to dividing data into manageable segments to simplify the unlearning process that completely
101 remove the influence of specific data from a ML model. Similarly, Chen et al. (2022a) applies a
102 comparable strategy to recommendation systems, advocating the creation of multiple sub-samples
103 to reduce the impact of forgetting particular observations. Although these methods can help manage
104 the computational load, they may still be resource-intensive if a large amount of data needs to be
105 forgotten.

106 To address these challenges, more recent research has focused on developing efficient unlearning
107 techniques that aim to approximate a full retraining from scratch while minimizing the associated
costs. For instance, Kurmanji et al. (2023a) introduces the SCRUB method, which frames unlearn-

ing as a teacher-student problem. This approach tackles various tasks associated with unlearning, such as removing bias, resolving confusion, and safeguarding user privacy. Another method, proposed by Graves et al. (2020), involves maintaining a record of which examples appeared in which training batches and tracking the corresponding parameter updates. By undoing updates related to sensitive information, this method effectively removes the influence of such data from the model. Following a recent challenge on unlearning, Triantafillou et al. (2024) analyzed progress in the field and highlighted strategies based on erasure and repair phases among the best solutions. Finally, often used to apply the right to be forgotten to the learning model, unlearning strategies have also been applied for backdoor removal (Min et al. (2025)).

Regarding machine unlearning in LLMs, previous approaches used gradient ascent (GA) over undesired knowledge to inversely optimize models (Jang et al. (2022); Yao et al. (2024)). Other unlearning schemes for LLMs are based on more and less selective pruning strategies (Pochinkov & Schoots (2024); Liu et al. (2025); Zhang et al. (2025)). For instance, Chen et al. (2024) propose a solution to localize neurones that memorize sensitive informations in order to erase only these neurones. SANI also takes advantage of this two-phase unlearning strategy with targeted erasure and controlled reconstruction.

2.2 RISK OF MEMORIZATION OF TRAINING DATA

A central question related to the risks of memorization of LLMs concerns the extent to which models memorize their training data (Carlini et al. (2019; 2023); Nasr et al. (2023); Zhang et al. (2021); Schwarzschild et al. (2024)). We can notably cite extractable memorization and membership inference.

Extractable memorization: Extractable memorization is a type of attack that aims to use the model to infer information from the original data (Liu et al. (2021)). This attack mainly concerns text generation models, such as GPT. These models are trained to produce text based on what they have seen during training. However, the model is not expected to be a basic parrot and repeat exactly the sentences it has seen. This is especially concerning if the data it is repeating is sensitive. This has been shown to be the case with GPT-2 for example, from which the names and addresses of individuals can be extracted Carlini et al. (2020). Compressible memorization (Schwarzschild et al. (2024)) extends this definition by evaluating how short the minimal requested sentence (or prompt) that elicits the sequence.

Membership inference: Membership inference attack (MIA) is a more common inference attack in machine learning, which aims to infer whether a specific data was used in the training data of a target model (Carlini et al. (2022)). There are different techniques that can be used to perform a MIA attack depending on if the adversary has an access to the model parameters (i.e., white-box access), or access to a ground-truth subset of member and non-member samples. One technique consists to analyze the loss of member and non member samples (Yeom et al. (2018)), another one is to use multiple shadow models (Shokri et al. (2017); Ye et al. (2022)) trained to mimic the behavior of the target model on an auxiliary dataset. An adversarial model is then trained to infer membership from the loss or from shadow models. Another method (Zhang et al. (2021)) is based on comparing the performance of the target model trained on a dataset with a specific input, with a second model trained without it. As ML models are supposed to learn general information, one piece of data (even rare, outlier or mislabelled samples) is not supposed to be memorized and significantly changed the model's performance. By repeating this operation many times with different subset, it is possible to identify counter-factually memorized data.

Mitigation strategies: The most popular approach to mitigating memorization is Differential Privacy Dwork (2006) (DP). DP is a mathematical property that a model must satisfy in order to disclose as little information as possible. This property requires the model to learn a limited amount of information at each training step. The most popular method to apply DP in machine learning is DP-SGD: Differentially-Private Stochastic Gradient Descent (Abadi et al. (2016)). The idea is to apply DP during the training phase by clipping gradient updates and adding noise at each step. DP is known to significantly decrease the accuracy of the model (Jagannatha et al. (2021)) and privacy budget management is difficult.

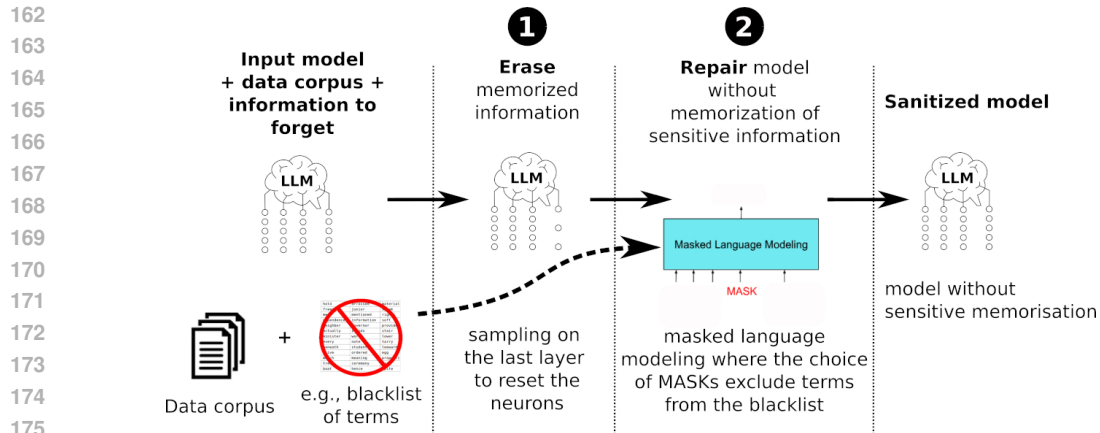


Figure 1: Overview: once the information to be unlearned is defined, SANI relies on erase and repair phases allowing to disturb only the memorization of fine information before repairing the model by specializing it while ensuring the absence of memorization of the terms to be sanitized.

To reduce privacy risks, Chen et al. (2022b) proposed to relax the loss in order to reduce the distinguishability between the training and testing loss distributions, and Chen et al. (2024) proposed a method to localize specific neurons responsible for memorizing PII in LLMs through adversarial training and then disable these neurons. To avoid memorizing personal information used during model specialization, Boutet et al. (2025) proposes a language modeling scheme that first identifies the direct and indirect identifiers, and secondly applies a language modeling scheme that avoids choosing these identifiers as masks. We leverage this language modeling scheme in our second unlearning phase to repair the model without memorizing terms defined as sensitive.

3 SANI: SANITIZING LLMs

In order to sanitize LLMs, we have developed SANI. SANI relies on an unlearning scheme to remove specific pieces of information (e.g., sensitive, unreliable, or copyrighted data) from the memorisation of the model. Figure 1 depicts an overview of how SANI works and the steps it involves. It relies on the identification of sensitive information to forget (Section 3.1), and two steps, a selective erasure (Section 3.2) and a repair without without memorisation of sensitive information (Section 2.2).

3.1 SENSITIVE INFORMATION TO FORGET

First, SANI needs as input the LLM to be sanitized and the identification of the sensitive information to forget and unlearn from the model. Depending on the purpose of the sanitization, the nature of the information to be forgotten can be different. For example, if the goal is to anonymize the model, the sensitive information to be forgotten concerns all direct and indirect identifiers linked to individuals. If, on the other hand, the goal is to share a model trained with business data, the sensitive information to be forgotten will be all confidential or copyrighted information that could have been exposed and memorized by the model. Another use case could be to remove biases from the model. In this latter case, the sensitive information to be forgotten would be outliers, aberrant data or other poisoned data. This sensitive information to forget consists of a list of single or several words (i.e., n-grams) contained in the training data.

3.2 TARGETED ERASURE

Once the LLM and the sensitive information to forget are defined, SANI conducts a targeted erasure. Specifically, the erasure step (step 1 in the figure) resets certain neurons in the model’s architecture to disrupt its memorization. To control this reset and reduce its detrimental effect on memorization, only the last layer of the model is affected, with a random reset of 50% of the neurons. By targeting erasure only on the last layer, fine-grained memorisation is disrupted without affecting the model’s

learning foundation Jourdan et al. (2021). This allows, among other things, to maintain strong model convergence dynamics during subsequent learning steps, carried out during the reconstruction phase. An LLM (e.g., a BERT model) is built from a series of attention heads, followed by a linear layer. The purpose of the attention heads is to capture dependency information between words. The last linear layer uses this dependency information for a specific task, such as text or token classification. It is in this last linear layer that the weights are reset, which does not affect the memorization of the dependency information between words of the other layers.

3.3 REPAIR WITHOUT SENSITIVE MEMORISATION

In the last reconstruction step (step ② in the figure), controlled training cycles are performed. Specifically, the model is specialized through language modeling applied to the model’s training data, while ensuring that sensitive blacklisted terms are not memorized by the model Boutet et al. (2025). To achieve this, the masks defined during this specialization phase are randomly selected from the input data set (i.e., the model’s training data), while excluding terms contained in the blacklist (these terms can, however, serve as context for predicting another mask). Ensuring that these terms are not memorized in the model through the language modeling (i.e., the masked language modeling for BERT-like models, and the causal language modeling for GPT-like models) reduces the model’s ability to regurgitate them (in a generative model like GPT, for example) or predict them (in a predictive model like BERT, for example). Since only the last layer of the model was modified during the erasure phase, this fine-tuning performed in the reconstruction phase converges quickly, and only a limited number of epochs are required to converge to a utility close to the original model.

4 EVALUATION

We performed an extensive evaluation of SANI on two realistic use cases using medical and book datasets, and we considered both off-the-shelf BERT and GPT pre-trained models. To capture the full impact of our proposed solutions, we considered a set of metrics capturing both the utility and the risk of regurgitation of sensitive information and we compared SANI against two baseline approaches. The main results show that SANI unlearns very quickly sensitive information and that the number of regurgitations decreases drastically after only 1 epoch. The most repeated sensitive terms in training (or in the specialization) are the terms most affected by unlearning.

4.1 EXPERIMENTAL SETUP

This section provides details about our experimental setup. Additional methodological information to replicate our evaluations is provided in the Appendix A.1.

4.1.1 USE CASES

To evaluate SANI we considered two realistic use cases: to sanitize fine-tuned models, and to sanitize off-the-shelf pre-trained models.

Sanitizing fine-tuned models: We consider here the example of a hospital that has already spent resources to fine-tune a pre-trained model on its own medical records and wishes to anonymize this model without having to conduct costly training and specialization again. To implement this use case, we consider both a pre-trained BERT and GPT models fine-tuned on raw medical records and a models fine-tuned on pseudonymized medical records. The goal of the sanitization is to ensure that all identifiers referring to an individual are not regurgitated by the models. This includes unlearning direct and indirect identifiers in the case of a specialized model with raw data, and, respectively, unlearning indirect identifiers in the case of a specialized model with pseudonymized data.

Sanitizing pre-trained models: We consider here a company that would like to ensure that a model it has trained with its own data does not contain confidential information. To implement this use case, we consider a pre-trained BERT model (i.e., trained using a very large amount of web-scraped data) and the goal of the sanitization is to unlearn certain terms categorized as confidential that were used in texts known to be part of the model’s training. More specifically, we identified terms belonging to specific categories (identified via a NER model) contained in a public dataset used when training the BERT model, and the goal is to ensure that the model does not regurgitate them.

In both cases, to justify unlearning, it must be significantly less expensive than completely retraining the model by excluding the data to be removed (Muresanu et al. (2025)). A 20% barrier has been established, beyond which unlearning is considered too costly (Triantafillou et al. (2024)).

4.1.2 DATASET

To conduct our experiments on health data (the use case on sanitizing fine-tuned models), we considered the N2C2 datasets (Henry et al. (2020)). More precisely, we leverage a dataset with medical discharge summaries (Uzuner et al. (2007)) gathering 928 records (i.e., English free text) that have been annotated for replacing all authentic Personally Identifiable Information (PII) with realistic surrogates. To identify sensitive information used to pre-train BERT model (the use case on sanitizing pre-trained models), we used the BookCorpus dataset (Zhu et al. (2015)), known to be part of the training of BERT model (Devlin et al. (2019)). BookCorpus is a popular large-scale dataset consisting of the text of around 7,000 self-published books scraped from the indie ebook distribution website Smashwords. The dataset consists of around 985 million words, and the books that comprise it span a range of genres. We only use the first 500,000 texts from the dataset available on HuggingFace², which represents 77 million words. To assess the impact of unlearning on a classification downstream task, we trained a model to classify emotions associated with a text. To achieve that, we used a dataset gathering Twitter messages (Saravia et al. (2018)). The tweets were annotated and spanned into eight basic emotion categories: anger, anticipation, disgust, fear, joy, sadness, surprise, and trust.

4.1.3 MODELS

To cover both descriptive and generative uses of language models, we considered both the BERT and GPT architectures (Radford et al. (2019)). For BERT, we considered the base model composed of 12 layers, 768 hidden dimensions, 12 attention heads, and 110 million parameters, and for GPT, we considered the smaller version of GPT-2 with 124 million parameters, both from Hugging Face. We did not consider larger versions of the model due to limited resources, and we also discarded the distilled version of these models.

We also leveraged a off-the-shell NER (Named Entity Recognition) model from HuggingFace³ used to identify terms categorized as confidential in the use case on sanitization of pre-trained model. Specifically, this model is a bert-base-cased model that was fine-tuned on the English version of the standard CoNLL-2003 Named Entity Recognition dataset (Tjong Kim Sang & De Meulder (2003)). Terms classified as PLACE, PERSON, ORGANIZATION, and MISCELLANEOUS are marked as confidential in our use case and must be unlearned. This includes 380 different words and represents 1,953,000 iterations including repetitions.

4.1.4 METRICS

We consider different metrics in our evaluation to better capture the trade-off between the quality of the unlearning and the performance of the model.

Utility: For the performance evaluation of BERT-like models, we quantify the ability of the models to predict masked terms well (measured by the prediction accuracy). For GPT-like models, this performance is evaluated by the model’s ability to generate the right tokens in the right places (measured by the tokens accuracy). To evaluate the performance of the pre-trained model, we considered a downstream classification task. More specifically, we measure the model’s performance through an emotion classification task (measured by the F1-Score metric on emotion prediction).

Regurgitation: To assess the risk of regurgitation of terms that should have been unlearned, we study the model’s ability to predict both identifiers (direct and indirect) and terms marked as confidential, contained in the dataset used to fine-tune or train it. As soon as one of these terms is predicted, whether in the correct place in the original text or elsewhere, we count a regurgitation. The metric named *Privacy* measures the ability of a model to predict any identifying term (i.e., direct or indirect identifier). A *Privacy*=1 means that no identifiers are predicted or regurgitated by the model, and a *Privacy*=0 means that all iterations of identifiers are regurgitated. The metric named

²Hugging Face - <https://huggingface.co/>

³bert-base-NER: <https://huggingface.co/dslim/bert-base-NER>

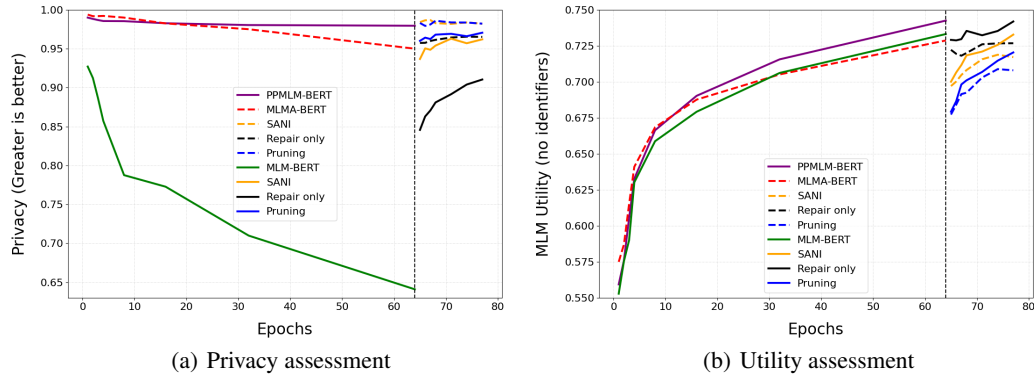


Figure 2: Sanitizing a fine-tuned BERT model: a single epoch is enough to drastically reduce identifier regurgitation while limiting utility reduction.

Regurgitation, in turn, quantifies the ability of the model to predict a term marked as confidential. A Regurgitation=1 means that all iterations of confidential terms are predicted by the model, and conversely, a Regurgitation=0 means that none are predicted.

4.1.5 COMPARATIVE BASELINES

We compare our unlearning scheme SANI against different comparative baselines.

Repair only: Unlike new unlearning schemes that include an erasure and a repair phase, this baseline includes only a repair phase. This repair phase is performed under the same conditions as SANI, i.e., the masking performed in the model specialization operation excludes the terms to be unlearned.

Pruning: This baseline consists of an erasure and a repair phase. In the erasure phase, the neurons in each layer of the model are sorted in ascending order according to their weight, the 1% of the most relevant neurons (i.e., with the highest weights) remain unaffected, while 20% of the remaining 99% are randomly reset. The repair phase is identical to SANI and the repair-only baseline.

PPMLM-BERT and PPCLM-GPT: This baseline (which does not implement unlearning) corresponds to a complete new specialization of the model that avoids memorizing the terms to be unlearned from the beginning. This baseline implements the solution proposed by Boutet et al. (2025), and represents the lower bound in terms of regurgitation.

4.2 SANITIZING A FINE-TUNED MODEL

Here, we evaluate the ability of SANI to sanitize the identifiers of a Masked Language Model (i.e., a BERT model) specialized with medical data. Figure 2 depicts the privacy and utility trade-off of the BERT model during both its fine-tuning (i.e., until epoch 64) and during its sanitization (i.e., from epochs 64 to 78). During fine-tuning, we distinguish three curves: one corresponding to a specialization without protection scheme (named MLM-BERT), a specialization from a pseudonymized dataset where the directly identifiers has been replaced by "X" (named MLMA-BERT), and a specialization that has integrated a protection mechanism to avoid memorizing both direct and indirect identifiers (named PPMLM-BERT). First, the evaluation of the privacy (Figure 5(a)) shows that without protection scheme (i.e., MLM-BERT), privacy drops during the specialization, the more the number of epochs increases, the more privacy decreases. This drop in privacy shows that the model tends to memorize more and more identifiers over the course of specialization. Next, we observe that the privacy of MLMA-BERT (corresponding to a specialization using pseudonymized data, i.e., without direct identifiers) also decreases over the course of specialization, but significantly less. This is due to the fact that the model only memorizes indirect identifiers, which are less abundant than direct identifiers. Finally, by avoiding the memorization of identifiers (both direct and indirect), PPMLM-BERT maintains approximately the same level of privacy as specialization progresses.

Here we consider two models to be sanitized. The first is the model fine-tuned without protection (i.e., MLM-BERT at epoch 64) and the goal is to sanitize the model from all identifiers (direct and indirect) memorized by the model. The second is the model fine-tuned using pseudonymized data (i.e., MLMA-BERT at epoch 64) and the objective is to sanitize this model from all indirect identifiers memorized by the model (the direct identifiers are already absent in the data used to fine-tune the model). The results show that in both cases, privacy increases very rapidly. For example, a single epoch is enough to increase the privacy of a specialized model without protection by approximately 0.2, 0.3, and 0.32 using an unlearning strategy based on repair only, SANI, and pruning, respectively. The privacy of a specialized model with pseudonymized data also increases to a lesser extent after one epoch. In both cases, privacy continues to progress until the end of the unlearning cycle at epoch 77 (which corresponds to 20% of the cost of the specialization). At epoch 77, all unlearning strategies except repair-only, display a privacy level close to that of a privacy-preserving specialization solution that does not memorize identifiers (i.e., privacy close to 1, the upper bound).

Figure 5(b), in turn, depicts the utility assessment during the specialization (until epoch 64) and during the sanitization (from 65 to 77). During specialization, all approaches (i.e., without protection MLM-BERT, using pseudonymized data MLMA-BERT, or using a privacy-preserving scheme that avoids to memorize identifiers PPMLM-BERT) show a similar increase in utility, starting from 0.57 at epoch 0 to 0.72 at epoch 64. At the beginning of the sanitization, except for the repair-only strategy, an erasure step is performed by resetting the weights of some neurons to zero. This reinitialization degrades the utility at the beginning of the sanitization, which then increases as a function of the number of epochs. The results show that the SANI strategy reaches the same level of utility at epoch 77 as at the end of the specialization without protection measures. The pruning strategy, on the other hand, degrades utility slightly. For the repair-only strategy, this corresponds to continuing learning. This is confirmed by an increase in utility in line with that achieved by MLM-BERT at the end of specialization. The sanitization of indirect identifiers from models trained on pseudonymised data shows the same trends. However, the level of utility is slightly lower.

Considering privacy and utility, SANI shows the best trade-off compared to a unlearning strategy based on pruning or repair-only. Repair-only provides the best utility but the worst privacy, and pruning depicts roughly the same privacy than SANI but a smaller utility. The assessment of the sanitization of a fine-tune Causal Language Model (i.e., a GPT-type model) shows the same trend. Due to space constraints, the results are only presented in Appendix A.2.

4.3 SANITIZING A PRE-TRAINED MODEL

We now evaluate the ability of SANI to sanitize terms defined as confidential information from a pre-trained Masked Language Model (i.e., a BERT model) without altering its performance on a downstream classification task. Figure 3 shows the evolution of the regurgitation and utility during the sanitization process. Figure 3(a) shows that the off-the-shelf model initially displays a regurgitation level close to 0.25 and that only a single epoch is enough to drastically reduce it close to 0.05. Regurgitation then decreases over the following epochs of sanitization less rapidly. The repair-only strategy, by not resetting the weight of certain neurons, regurgitates more sensitive information than SANI and pruning strategies, which display a similar level.

Figure 3(b), in turn, depicts the utility during sanitization by measuring the performance of a downstream classification task (i.e., detection of the emotion of the text). The results show that the models exhibit similar performance, between 0.92 and 0.93 of F1-scores, throughout the sanitization process. Therefore, sanitizing confidential information has no impact on the model’s performance when used in a downstream emotion classification task.

Finally, to better understand this rapid decline in regurgitation after a single epoch, we study which terms are most likely to be regurgitated and at what level during unlearning of SANI. To achieve this goal, Figure 4(a) reports for each term to unlearn their number of repetitions in the training data according to their number of regurgitations observed subsequently. Results clearly show a correlation between the number of times the model was exposed to sensitive terms and the number of regurgitations of these terms, the more repeated, the more regurgitated. Figure 4(b), in turn, depicts the cumulative number of regurgitations according to the number of times the term has been repeated in the training data. Results consistently show that without unlearning (i.e., at epoch 0) the greatest amount of regurgitation occurs for the most frequently repeated sensitive terms during

432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485

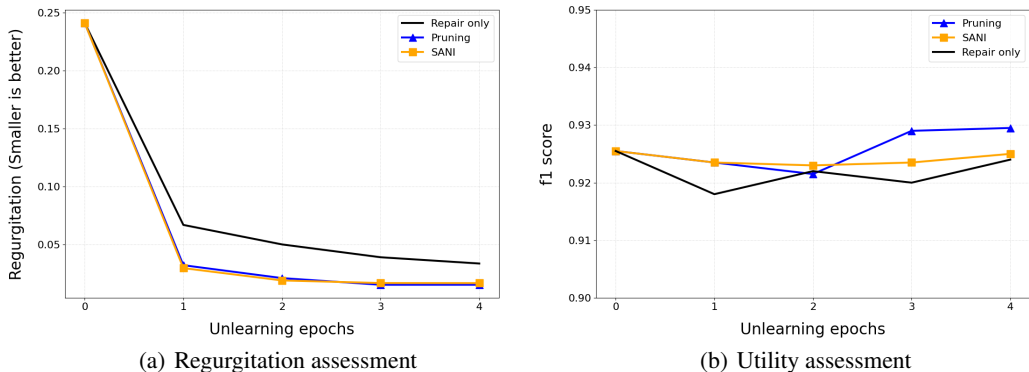


Figure 3: Sanitizing a pre-trained BERT model: a single epoch is enough to drastically reduce the regurgitation of confidential information while not affecting the model’s performance for a downstream classification task.

training. After one epoch of unlearning, we observe that the regurgitation associated with the most frequently repeated sensitive information in training decreases drastically. We go from 470,000 regurgitations in total without unlearning, to 60,000 and 30,000 regurgitations after 1 and 2 epochs, respectively. However, counterintuitively, the results show that for very rarely repeated terms in the training data, the number of regurgitations increases slightly after unlearning. This is certainly due to the reduction in the cardinality of the words to be predicted, mechanically increasing the probability of predicting rarely repeated words.

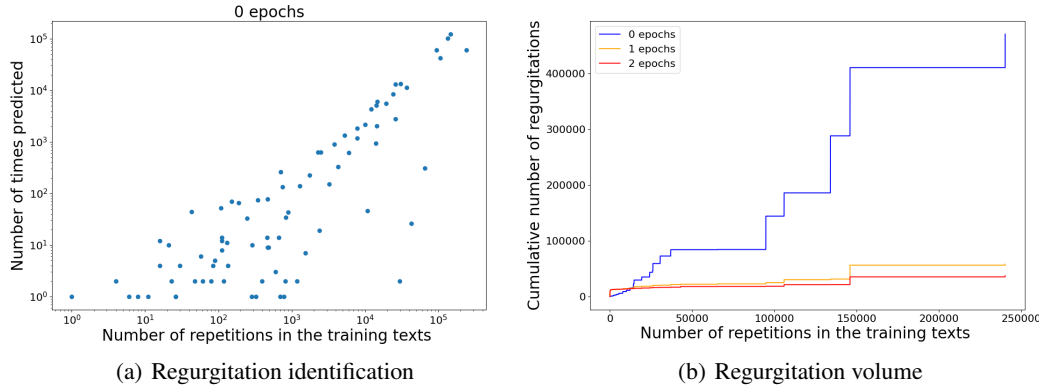


Figure 4: Without unlearning, the most repeated terms in the training data are the most regurgitated, SANI drastically reduces the regurgitation of the most repeated terms.

5 CONCLUSION

This article presents SANI, an unlearning strategy to sanitize LLMs, illustrated through two use cases: anonymizing a model (i.e., forgetting direct and indirect identifiers) and removing confidential information from the model. This unlearning strategy can be applied to a model specialized on a corpus of data, or to a pre-trained model. This unlearning strategy addresses the need for model sharing by limiting the risk of subsequent regurgitation of sensitive information, without paying the prohibitive cost of retraining models from scratch. We exhaustively evaluated our approach and demonstrated that it quickly unlearns the sensitive information and offers the best trade-off between the quality of the unlearning and the model performance compared to other baselines. An interesting avenue for future work would be to use and evaluate SANI to remove the influence of biased data and for the removal of backdoors introduced into the model.

REFERENCES

- 486
487
488 Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and
489 Li Zhang. Deep learning with differential privacy. oct 2016. doi: 10.1145/2976749.2978318.
490 URL <https://doi.org/10.1145/2976749.2978318>.
- 491
492 Lucas Bourtole, Varun Chandrasekaran, Christopher A. Choquette-Choo, Hengrui Jia, Adelin
493 Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning, 2020. URL
494 <https://arxiv.org/abs/1912.03817>.
- 495
496 Antoine Boutet, Lucas Magnana, Juliette Sénéchal, and Helain Zimmermann. Towards the
497 anonymization of the language modeling, 2025. URL <https://arxiv.org/abs/2501.02407>.
- 498
499 Yinzhi Cao and Junfeng Yang. Towards making systems forget with machine unlearning. In *2015*
500 *IEEE Symposium on Security and Privacy*, pp. 463–480, 2015. doi: 10.1109/SP.2015.35.
- 501
502 Nicholas Carlini, Chang Liu, Úlfar Erlingsson, Jernej Kos, and Dawn Song. The secret sharer:
503 Evaluating and testing unintended memorization in neural networks, 2019. URL <https://arxiv.org/abs/1802.08232>.
- 504
505 Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine
506 Lee, Adam Roberts, Tom B. Brown, Dawn Song, Úlfar Erlingsson, Alina Oprea, and Colin Raffel.
507 Extracting training data from large language models. *CoRR*, abs/2012.07805, 2020. URL
508 <https://arxiv.org/abs/2012.07805>.
- 509
510 Nicholas Carlini, Steve Chien, Milad Nasr, Shuang Song, Andreas Terzis, and Florian Tramer. Mem-
511 bership inference attacks from first principles. In *2022 IEEE Symposium on Security and Privacy*
512 (*SP*), pp. 1897–1914. IEEE, 2022. doi: 10.1109/SP46214.2022.9833649.
- 513
514 Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramer, and Chiyuan
515 Zhang. Quantifying memorization across neural language models, 2023. URL <https://arxiv.org/abs/2202.07646>.
- 516
517 Chong Chen, Fei Sun, Min Zhang, and Bolin Ding. Recommendation unlearning, 2022a. URL
518 <https://arxiv.org/abs/2201.06820>.
- 519
520 Dingfan Chen, Ning Yu, and Mario Fritz. Relaxloss: Defending membership inference attacks
521 without losing utility, 2022b. URL <https://arxiv.org/abs/2207.05801>.
- 522
523 Ruizhe Chen, Tianxiang Hu, Yang Feng, and Zuozhu Liu. Learnable privacy neurons localization
524 in language models. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings*
525 *of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short*
526 *Papers)*, pp. 256–264, Bangkok, Thailand, August 2024. Association for Computational Lin-
527 guistics. doi: 10.18653/v1/2024.acl-short.25. URL <https://aclanthology.org/2024.acl-short.25/>.
- 528
529 Xianlai Chen, Jiamiao Lin, and Ying An. Dl-bert: a time-aware double-level bert-style model with
530 pre-training for disease prediction. In *2022 IEEE International Conference on Big Data (Big*
531 *Data)*, pp. 1801–1808, 2022c. doi: 10.1109/BigData55660.2022.10020513.
- 532
533 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep
534 bidirectional transformers for language understanding, 2019. URL <https://arxiv.org/abs/1810.04805>.
- 535
536 Henri Duprieu and Nicolas Berkouk. Techniques d’audit des grands modèles de langage. Technical
537 report, Commission Nationale Informatique et Libertés (CNIL), November 2024. URL <https://hal.science/hal-04782667>.
- 538
539 Cynthia Dwork. Differential privacy. In *International colloquium on automata, languages, and*
programming, pp. 1–12. Springer, 2006.

- 540 Sagar Goyal, Eti Rastogi, Sree Prasanna Rajagopal, Dong Yuan, Fen Zhao, Jai Chintagunta, Gau-
541 tam Naik, and Jeff Ward. Healai: A healthcare llm for effective medical documentation. In
542 *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, pp.
543 1167–1168, 2024. doi: 10.1145/3616855.3635739. URL [https://doi.org/10.1145/
544 3616855.3635739](https://doi.org/10.1145/3616855.3635739).
- 545 Laura Graves, Vineel Nagisetty, and Vijay Ganesh. Amnesiac machine learning, 2020. URL
546 <https://arxiv.org/abs/2010.10981>.
- 547
- 548 Sam Henry, Kevin Buchan, Michele Filannino, Amber Stubbs, and Ozlem Uzuner. 2018 n2c2 shared
549 task on adverse drug events and medication extraction in electronic health records. *Journal of the
550 American Medical Informatics Association*, 27(1):3–12, 2020. doi: 10.1093/jamia/ocz166. URL
551 <https://doi.org/10.1093/jamia/ocz166>.
- 552
- 553 Kexin Huang, Jaan Altosaar, and Rajesh Ranganath. Clinicalbert: Modeling clinical notes and
554 predicting hospital readmission, 2020. URL <https://arxiv.org/abs/1904.05342>.
- 555
- 556 Abhyuday Jagannatha, Bhanu Pratap Singh Rawat, and Hong Yu. Membership inference attack
557 susceptibility of clinical language models, 2021. URL [https://arxiv.org/abs/2104.
08305](https://arxiv.org/abs/2104.08305).
- 558
- 559 Joel Jang, Dongkeun Yoon, Sohee Yang, Sungmin Cha, Moontae Lee, Lajanugen Logeswaran, and
560 Minjoon Seo. Knowledge unlearning for mitigating privacy risks in language models, 2022. URL
561 <https://arxiv.org/abs/2210.01504>.
- 562
- 563 Théo Jourdan, Antoine Boutet, and Carole Frindel. Privacy assessment of federated learning using
564 private personalized layers. In *2021 IEEE 31st International Workshop on Machine Learning for
Signal Processing (MLSP)*, pp. 1–6, 2021. doi: 10.1109/MLSP52302.2021.9596237.
- 565
- 566 Meghdad Kurmanji, Peter Triantafillou, Jamie Hayes, and Eleni Triantafillou. Towards unbounded
567 machine unlearning. In *Proceedings of the 37th International Conference on Neural Information
Processing Systems*, NIPS ’23, Red Hook, NY, USA, 2023a. Curran Associates Inc.
- 568
- 569 Meghdad Kurmanji, Peter Triantafillou, Jamie Hayes, and Eleni Triantafillou. Towards unbounded
570 machine unlearning. In *Thirty-seventh Conference on Neural Information Processing Systems*,
571 2023b. URL <https://openreview.net/forum?id=OveBaTtUAT>.
- 572
- 573 Eric Lehman, Sarthak Jain, Karl Pichotta, Yoav Goldberg, and Byron Wallace. Does BERT pre-
574 trained on clinical notes reveal sensitive data? In *Proceedings of the 2021 Conference of the North
575 American Chapter of the Association for Computational Linguistics: Human Language Technolo-
576 gies*, pp. 946–959, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/
v1/2021.naacl-main.73. URL <https://aclanthology.org/2021.naacl-main.73>.
- 577
- 578 Ximeng Liu, Lehui Xie, Yaopeng Wang, Jian Zou, Jinbo Xiong, Zuobin Ying, and Athanasios V.
579 Vasilakos. Privacy and security issues in deep learning: A survey. *IEEE Access*, 9:4566–4593,
580 2021. doi: 10.1109/ACCESS.2020.3045078.
- 581
- 582 Zheyuan Liu, Guangyao Dou, Xiangchi Yuan, Chunhui Zhang, Zhaoxuan Tan, and Meng Jiang.
583 Modality-aware neuron pruning for unlearning in multimodal large language models, 2025. URL
<https://arxiv.org/abs/2502.15910>.
- 584
- 585 Nay Myat Min, Long H. Pham, and Jun Sun. Unified neural backdoor removal with only few clean
586 samples through unlearning and relearning. *IEEE Transactions on Information Forensics and
587 Security*, 20:6984–6998, 2025. doi: 10.1109/TIFS.2025.3586499.
- 588
- 589 Andrei Ioan Muresanu, Anvith Thudi, Michael R. Zhang, and Nicolas Papernot. Fast exact unlearn-
590 ing for in-context learning data for LLMs. In *Forty-second International Conference on Machine
591 Learning*, 2025. URL <https://openreview.net/forum?id=TzNVZEsqTi>.
- 592
- 593 Milad Nasr, Nicholas Carlini, Jonathan Hayase, Matthew Jagielski, A. Feder Cooper, Daphne
Ippolito, Christopher A. Choquette-Choo, Eric Wallace, Florian Tramèr, and Katherine Lee.
Scalable extraction of training data from (production) language models, 2023. URL <https://arxiv.org/abs/2311.17035>.

- 594 Thanh Tam Nguyen, Thanh Trung Huynh, Zhao Ren, Phi Le Nguyen, Alan Wee-Chung Liew,
595 Hongzhi Yin, and Quoc Viet Hung Nguyen. A survey of machine unlearning, 2024. URL <https://arxiv.org/abs/2209.02299>.
596
597
- 598 Nicholas Pochinkov and Nandi Schoots. Dissecting language models: Machine unlearning via
599 selective pruning, 2024. URL <https://arxiv.org/abs/2403.01267>.
600
- 601 Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever.
602 Language models are unsupervised multitask learners. 2019. URL [https://api.
603 semanticscholar.org/CorpusID:160025533](https://api.semanticscholar.org/CorpusID:160025533).
- 604 Elvis Saravia, Hsien-Chi Toby Liu, Yen-Hao Huang, Junlin Wu, and Yi-Shin Chen. CARER:
605 Contextualized affect representations for emotion recognition. In Ellen Riloff, David Chiang,
606 Julia Hockenmaier, and Jun’ichi Tsujii (eds.), *Proceedings of the 2018 Conference on Empirical
607 Methods in Natural Language Processing*, pp. 3687–3697, Brussels, Belgium, October-
608 November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1404. URL
609 <https://aclanthology.org/D18-1404/>.
- 610
611 Avi Schwarzschild, Zhili Feng, Pratyush Maini, Zachary C. Lipton, and J. Zico Kolter. Rethinking
612 llm memorization through the lens of adversarial compression, 2024. URL [https://arxiv.
613 org/abs/2404.15146](https://arxiv.org/abs/2404.15146).
- 614 Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership infer-
615 ence attacks against machine learning models. In *2017 IEEE symposium on security and
616 privacy (SP)*, pp. 3–18. IEEE, 2017. doi: 10.1109/SP.2017.41. URL [https://doi.
617 ieeeecomputersociety.org/10.1109/SP.2017.41](https://doi.ieeeecomputersociety.org/10.1109/SP.2017.41).
- 618
619 Sushant Singh and Ausif Mahmood. The NLP cookbook: Modern recipes for transformer based
620 deep learning architectures. *CoRR*, abs/2104.10640, 2021. URL [https://arxiv.org/abs/
621 2104.10640](https://arxiv.org/abs/2104.10640).
- 622 Erik F. Tjong Kim Sang and Fien De Meulder. Introduction to the CoNLL-2003 shared task:
623 Language-independent named entity recognition. In *Proceedings of the Seventh Conference
624 on Natural Language Learning at HLT-NAACL 2003*, pp. 142–147, 2003. URL [https:
625 //aclanthology.org/W03-0419/](https://aclanthology.org/W03-0419/).
- 626
627 Eleni Triantafillou, Peter Kairouz, Fabian Pedregosa, Jamie Hayes, Meghdad Kurmanji, Kairan
628 Zhao, Vincent Dumoulin, Julio Jacques Junior, Ioannis Mitliagkas, Jun Wan, Lisheng Sun
629 Hosoya, Sergio Escalera, Gintare Karolina Dziugaite, Peter Triantafillou, and Isabelle Guyon.
630 Are we making progress in unlearning? findings from the first neurips unlearning competition,
631 2024. URL <https://arxiv.org/abs/2406.09073>.
- 632
633 Özlem Uzuner, Yuan Luo, and Peter Szolovits. Evaluating the state-of-the-art in automatic de-
634 identification. *Journal of the American Medical Informatics Association*, 14(5):550–563, 09 2007.
635 ISSN 1067-5027. doi: 10.1197/jamia.M2444. URL [https://doi.org/10.1197/jamia.
636 M2444](https://doi.org/10.1197/jamia.M2444).
- 637
638 Qiang Wei, Xu Zuo, Omer Anjum, Yan Hu, Ryan Denlinger, Elmer V. Bernstam, Martin J Citardi,
639 and Hua Xu. Clinicallylayoutlm: A pre-trained multi-modal model for understanding scanned
640 document in electronic health records. In *2022 IEEE International Conference on Big Data (Big
641 Data)*, pp. 2821–2827, 2022. doi: 10.1109/BigData55660.2022.10020569.
- 642
643 Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor,
644 Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, Courtney Biles, Sasha Brown,
645 Zac Kenton, Will Hawkins, Tom Stepleton, Abeba Birhane, Lisa Anne Hendricks, Laura Rimell,
646 William Isaac, Julia Haas, Sean Legassick, Geoffrey Irving, and Iason Gabriel. Taxonomy of
647 risks posed by language models. In *Proceedings of the 2022 ACM Conference on Fairness, Ac-
countability, and Transparency, FAccT ’22*, pp. 214–229, New York, NY, USA, 2022. Associa-
tion for Computing Machinery. ISBN 9781450393522. doi: 10.1145/3531146.3533088. URL
<https://doi.org/10.1145/3531146.3533088>.

- 648 Heng Xu, Tianqing Zhu, Lefeng Zhang, Wanlei Zhou, and Philip S. Yu. Machine unlearning: A
649 survey. *ACM Comput. Surv.*, 56(1), August 2023. ISSN 0360-0300. doi: 10.1145/3603620. URL
650 <https://doi.org/10.1145/3603620>.
651
- 652 Yuanshun Yao, Xiaojun Xu, and Yang Liu. Large language model unlearning, 2024. URL <https://arxiv.org/abs/2310.10683>.
653
- 654 Jiayuan Ye, Aadyaa Maddi, Sasi Kumar Murakonda, Vincent Bindschaedler, and Reza Shokri.
655 Enhanced membership inference attacks against machine learning models. In *Proceedings of
656 the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS '22*,
657 pp. 3093–3106, New York, NY, USA, 2022. Association for Computing Machinery. ISBN
658 9781450394505. doi: 10.1145/3548606.3560675. URL [https://doi.org/10.1145/
659 3548606.3560675](https://doi.org/10.1145/3548606.3560675).
- 660 Hangu Yeo, Elahe Khorasani, Vadim Sheinin, Irene Manotas, Ngoc Phuoc An Vo, Octavian Popescu,
661 and Petros Zerfos. Natural language interface for process mining queries in healthcare. In *2022
662 IEEE International Conference on Big Data (Big Data)*, pp. 4443–4452, 2022. doi: 10.1109/
663 BigData55660.2022.10020685.
664
- 665 Samuel Yeom, Irene Giacomelli, Matt Fredrikson, and Somesh Jha. Privacy Risk in Machine Learn-
666 ing: Analyzing the Connection to Overfitting. In *2018 IEEE 31st Computer Security Foundations
667 Symposium (CSF)*, pp. 268–282, Los Alamitos, CA, USA, July 2018. IEEE Computer Society.
668 doi: 10.1109/CSF.2018.00027. URL [https://doi.ieeecomputersociety.org/10.
669 1109/CSF.2018.00027](https://doi.ieeecomputersociety.org/10.1109/CSF.2018.00027).
- 670 Chiyuan Zhang, Daphne Ippolito, Katherine Lee, Matthew Jagielski, Florian Tramèr, and Nicholas
671 Carlini. Counterfactual memorization in neural language models. *CoRR*, abs/2112.12938, 2021.
672 URL <https://arxiv.org/abs/2112.12938>.
673
- 674 Shengming Zhang, Le Zhang, Jingbo Zhou, Zhi Zheng, and Hui Xiong. Llm-eraser: Optimizing
675 large language model unlearning through selective pruning. In *Proceedings of the 31st ACM
676 SIGKDD Conference on Knowledge Discovery and Data Mining V.1, KDD '25*, pp. 1960–1971,
677 New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400712456. doi:
678 10.1145/3690624.3709312. URL <https://doi.org/10.1145/3690624.3709312>.
- 679 Yunyi Zhang, Ming Zhong, Siru Ouyang, Yizhu Jiao, Sizhe Zhou, Linyi Ding, and Jiawei Han.
680 Automated mining of structured knowledge from text in the era of large language models. In
681 *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*,
682 pp. 6644–6654, 2024.
- 683 Jin Zhao, Chao Liu, Jiaqing Liang, Zhixu Li, and Yanghua Xiao. A novel cascade instruction tuning
684 method for biomedical ner. In *ICASSP 2024-2024 IEEE International Conference on Acoustics,
685 Speech and Signal Processing (ICASSP)*, pp. 11701–11705. IEEE, 2024.
686
- 687 Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and
688 Sanja Fidler. Aligning books and movies: Towards story-like visual explanations by watching
689 movies and reading books. In *The IEEE International Conference on Computer Vision (ICCV)*,
690 December 2015.
691

692 A APPENDIX

693 A.1 METHODOLOGY

694 This section provides methodological details useful to reproduce our evaluations. First, when
695 datasets are pseudonymized before being used for model fine-tuning (named MLMA-BERT for BERT
696 and CLMA-GPT for GPT), direct identifiers are replaced by the term "X" in the text. Then, BERT
697 models (i.e., Masked Language Model) were trained by randomly replacing a subset of tokens from
698 the sequence by the <MASK> token, and asking the model to predict them using cross-entropy loss.
699 In our setting, 15% of the words are randomly selected. In the case of PPMLM-BERT, identifying
700 words are excluded from this random selection. If a word chosen to be masked consists of several
701

tokens, all of them are masked. At each epoch, 15% of non-identifying words are masked. Finally, GPT models (i.e., Causal Language Model) were trained by sequentially (i.e., autoregressively) predicting all tokens. In PPCLM-GPT, if this word is a direct or indirect identifier, it is replaced by a specific padding token that is not taken into account during fine-tuning (i.e., for the computation of the loss).

During the fine-tuning of pre-trained models, all measurements were performed after 4, 8, 16, 32 and 64 epochs to observe their evolution. We used 8 batches of 512 tokens for training. For language modeling, the learning rate starts at $1e-4$ and decreases linearly until the end of fine-tuning. For classification, we took the models after 1, 2, 3 and 4 unlearning epochs, and trained them again on a classification task for 4 epochs. The learning rate in these cases starts at 0, 10% of the total number of warmup steps with the learning rate increasing until reaching $2e-5$ then decreasing linearly on the remaining 90% steps.

All the computation has been parallelized on a hybrid GPU/CPU computing farm.

A.2 SANITIZING A FINE-TUNED CAUSAL LANGUAGE MODEL

In this section, we evaluate the ability of SANI to unlearn and remove identifiers (direct and/or indirect) from a Causal Language Model (i.e., a GPT model) specialized with medical data. As in the case of a Masked Language Model (i.e., a BERT model) reported in Section 4.2, we evaluate both the evolution of privacy (i.e., the regurgitation of identifiers) and the model’s ability to correctly generate tokens at the right location in the text. The results show the same trends as for BERT sanitization, with a rapid reduction in regurgitation after a single epoch and a better trade-off between unlearning quality and model utility for SANI compared to other baselines.

We however observe a notable difference in the model’s utility level, which displays lower performance than for the BERT model (for all baselines). This difference is caused by the absence of identifiers in the context accessed by the model, which have been replaced by padding tokens (see Appendix A.1). Since the model has access to less contextual information when predicting tokens, the quality of the prediction is affected.

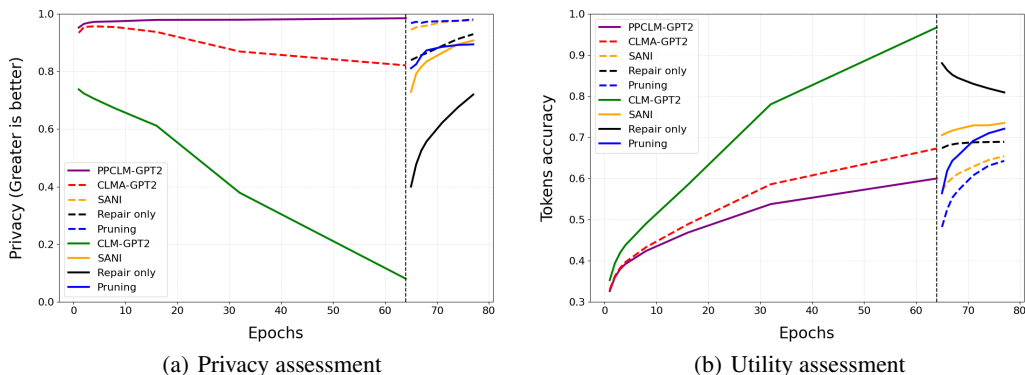


Figure 5: Sanitizing a fine-tuned GPT model: a single epoch is enough to drastically reduce the regurgitation of identifiers while limiting utility reduction.