

Learnable Privacy Neurons Localization in Language Models

Ruizhe Chen
Zhejiang University

Tianxiang Hu
Zhejiang University

Yang Feng
Angelalign Technology Inc.

Zuozhu Liu *
Zhejiang University

Abstract

Concerns regarding Large Language Models (LLMs) to memorize and disclose private information, particularly Personally Identifiable Information (PII), become prominent within the community. Many efforts have been made to mitigate the privacy risks. However, the mechanism through which LLMs memorize PII remains poorly understood. To bridge this gap, we introduce a pioneering method for pinpointing PII-sensitive neurons (privacy neurons) within LLMs. Our method employs learnable binary weight masks to localize specific neurons that account for the memorization of PII in LLMs through adversarial training. Our investigations discover that PII is memorized by a small subset of neurons across all layers, which shows the property of PII specificity. Furthermore, we propose to validate the potential in PII risk mitigation by deactivating the localized privacy neurons. Both quantitative and qualitative experiments demonstrate the effectiveness of our neuron localization algorithm.

1 Introduction

Large Language Models (LLMs) have demonstrated exceptional performance on various NLP tasks, leveraging huge model architectures and a tremendous scale of real-world training data (OpenAI, 2023; Touvron et al., 2023; Taori et al., 2023). However, the ability of memorization within LLM has also raised concerns regarding security within human society (Bender et al., 2021; Bommasani et al., 2021). One significant concern is that private information may be memorized and leaked by LLMs. An attacker can extract private information contained in the training corpus, especially Personally Identifiable Information (PII) such as names or addresses (Carlini et al., 2021, 2022; Huang et al., 2022; Rocher et al., 2019; Lukas

et al., 2023), which constitutes a privacy violation according to the General Data Protection Regulation (GDPR) (Regulation, 2016). Various methods have been proposed to mitigate the memorization of PII (Lison et al., 2021; Anil et al., 2021), primarily focusing on the sanitization of training data (Vakili et al., 2022; Lee et al., 2021), or providing differential privacy (DP) guarantees during the training process (Yu et al., 2021b; He et al., 2022). However, the mechanism by which LLMs memorize PII is not well understood.

In this paper, we propose a novel privacy neuron localization algorithm. Our method utilizes the hard concrete distribution (Louizos et al., 2017) to make neuron masks learnable and design adversarial objective functions to minimize the predictive accuracy of PII while preserving other non-sensitive knowledge. Besides, we employ another penalty to minimize the number of localized neurons, thus localizing a minimal subset of PII-specific neurons. We subsequently conduct a comprehensive analysis of the localized privacy neurons. Our findings reveal that memorization is localized to a minor subset of neurons, which are spread across all layers, predominantly within the MLP layers. Furthermore, we also discover that privacy neurons have the property of specificity for certain categories of PII knowledge. Inspired by the observation, we propose to investigate the privacy leakage mitigation ability by deactivating the localized neurons during the evaluation process, thus eliminating the memorization of PII. Experimental results demonstrate that our framework can achieve comparable performance in mitigating the risks of PII leakage without affecting model performance.

2 Method

Denote $f(\theta)$ as a PLM with parameters θ . Given a sequence of tokens $x = [x_1, \dots, x_T]$ from the training corpus, $f(\theta)$ can leak the private

*Corresponding author.

Our code is available at <https://github.com/richhh520/Learnable-Privacy-Neurons-Localization>.

sequence $[x_p, \dots, x_{p+I}]$ within x by generating $[x_p, \dots, x_{p+I}] = \arg \max_*(P_{f(\theta)}(*|x_{<p}))$. For example, as shown in Fig. 1, the email address of *Kent Garrett* is disclosed by the model, which constitutes significant societal risks.

In this section, we introduce a novel neuron localization algorithm that localizes neurons in $f(\theta)$ responsible for PII prediction, to elucidate the underlying mechanisms of PII memorization, as illustrated in Fig. 1. To be specific, our goal is to find a small subset of neurons $f(m \odot \theta)$ (or equivalently, the mask m) that deactivating these neurons prevents PII leakage, while not affecting the language modeling ability, thus indicating the memorization of PII-specific knowledge. m and \odot denote the differentiable binary neuron mask and Hadamard product operator respectively.

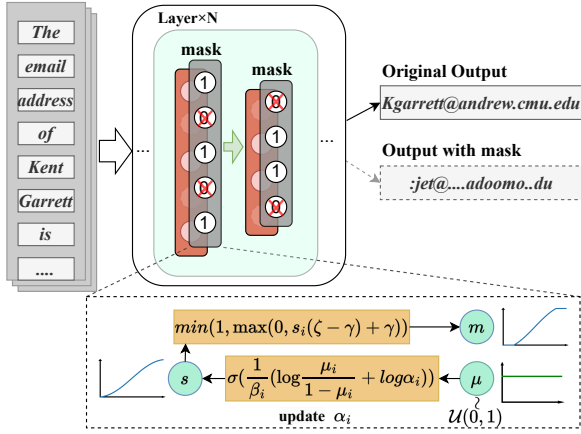


Figure 1: An illustration of our neuron localization method.

2.1 Differentiable Neuron Mask Learning

Since the training loss is not differentiable for binary masks, we resort to a practical method to learn subnetworks (Louizos et al., 2017), which employs a smoothing approximation of the discrete Bernoulli distribution (Maddison et al., 2016). Following (Zheng et al., 2022), we assume mask m_i corresponding to each neuron to be an independent random variable that follows a hard concrete distribution $\text{HardConcrete}(\log \alpha_i, \beta_i)$ with temperature β_i and location α_i (Louizos et al., 2017):

$$s_i = \sigma\left(\frac{1}{\beta_i} \left(\log \frac{\mu_i}{1 - \mu_i} + \log \alpha_i\right)\right), \quad (1)$$

$$m_i = \min(1, \max(0, s_i(\zeta - \gamma) + \gamma)), \quad (2)$$

where σ denotes the sigmoid function. s_i denotes the mask score of each neuron and m_i is the

approximately discrete activation value (i.e., almost 0 or 1) of s_i . γ and ζ are constants, and μ_i is the random sample drawn from uniform distribution $\mathcal{U}(0, 1)$. In this work, we also treat β_i as a constant, thus only α is the set of differentiable parameters for m . During the inference stage, the mask m_i can be calculated through a hard concrete gate:

$$\min(1, \max(0, \sigma(\log \alpha_i)(\zeta - \gamma) + \gamma)). \quad (3)$$

Algorithm 1 Neuron Localization Algorithm.

Require: mask parameters α , pre-trained language model $f(\theta)$ with frozen parameter θ , training corpus X , hyper-parameters $\beta, \gamma, \zeta, \eta$, learning rate lr .

- 1: Initialize $s \leftarrow \sigma\left(\frac{1}{\beta} \left(\log \frac{\mu}{1 - \mu} + \log \alpha\right)\right)$, where $\mu \sim \mathcal{U}(0, 1)$
 - 2: Initialize $m \leftarrow \min(1, \max(0, s(\zeta - \gamma) + \gamma))$
 - 3: Initialize $f(\theta) \leftarrow f(m \odot \theta)$
 - 4: **for** epoch in `num_epochs` **do**
 - 5: **for** x in X **do**
 - 6: Generate $f(m \odot \theta)$ with step1-3
 - 7: **if** `optimizer_idx == 0` **then**
 - 8: $\mathcal{L} = \mathcal{L}_m(f(m \odot \theta), x) + \eta R(m)$
 - 9: **else**
 - 10: $\mathcal{L} = \mathcal{L}_{adv}(f(m \odot \theta), x) + \eta R(m)$
 - 11: **end if**
 - 12: $\alpha = \alpha - \text{lr} \cdot \nabla_{\alpha}(\mathcal{L})$
 - 13: **end for**
 - 14: **end for**
 - 15: $m \leftarrow \min(1, \max(0, \sigma(\log \alpha)(\zeta - \gamma) + \gamma))$
 - 16: **return** m
-

2.2 Adversarial Privacy Neuron Localization

To localize PII-specific neurons, we propose to negate the original training objective, i.e., maximizing the negative log-likelihood of the PII token sequences. Specifically, given a sequence of tokens $x = [x_1, \dots, x_T]$ from the training corpus and PII tokens $[x_p, \dots, x_{p+I}]$, our training objective is:

$$\mathcal{L}_m(f(m \odot \theta), x) = \sum_{i=1}^I \log(P(x_{p+i}|x_{<p+i})). \quad (4)$$

On the other hand, to preserve the original language modeling ability of $f(m \odot \theta)$, we propose to perform further training on the corpus, utilizing the pre-training loss as the adversarial loss:

$$\mathcal{L}_{adv}(f(m \odot \theta), x) = - \sum_{t=1}^T \log(P(x_t|x_{<t})). \quad (5)$$




Figure 2: The distribution of privacy neurons in different layers (mean and std across three datasets).




Figure 3: The distribution of privacy neurons in different model components.




Figure 4: Heatmap of the similarity of privacy neurons according to different categories.

Finally, to minimize the number of localized neurons, we penalize the number of localized neurons by minimizing the L_0 complexity of mask scores which are zero:

$$R(m) = -\frac{1}{|m|} \sum_{i=1}^{|m|} \sigma(\log \alpha_i - \beta_i \log \frac{-\gamma}{\zeta}). \quad (6)$$

In the training step, the differentiable mask is adversarially trained by Eq. 4 and Eq. 5, with Eq. 6 as an auxiliary. The overall optimization procedure is elaborated in Algorithm 1.

3 Can PII memorization be localized?

In this section, we primarily investigate the following questions: (a) Is the memorization of PII confined to the latter layers of the model (Baldock et al., 2021)? (b) How many neurons are required to memorize privacy information? (c) Are privacy neurons specific?

3.1 Experiment Setup

Model and Dataset. We utilize the GPT-Neo (125M, 1.3B) LMs (Black et al., 2021). We utilize **Enron Email Dataset** (Klimt and Yang, 2004) and **ECHR** (Chalkidis et al., 2019) containing different types of PII in two domains.

PII and NER. For Enron dataset, we regard *email* and *name* as PII. We utilize the predefined prompt templates (e.g. *the email address of target_name is*) and the email-name correspondence provided in DecodingTrust (Wang et al., 2023) to extract PII. For ECHR, We tag PII in 4 categories (*person*, *law*, *date* and *gpe*) in the corpus, utilizing Named Entity Recognition (NER) tagger from Flair (Schweter and Akbik, 2020). We utilize the prefix context to prompt generation.

3.2 Privacy Neuron Distribution

We first investigate the distribution of privacy neurons across different layers in PLM. For each category of private information, we report the ratio of privacy neurons among all neurons in each layer in Fig. 2. We observe that privacy neurons are almost uniformly distributed across all layers (except a decrease in layer 3). We further explore the distribution in different model components (i.e., query, key, value, and MLP) in memorizing PII. As shown in Fig. 3, The ratio of privacy neurons in the MLP layer is significantly higher than in other components. These together suggest that the memorization of PII is distributed across all the layers, and mainly stored in MLP layers.

3.3 Category-wise Memorization

Following the previous part, we observe that the distribution patterns of different categories of privacy neurons in the model are also similar. Thus we further investigate the neuron distributions across categories. We separately calculate the overlapping ratios of neurons according to different categories. The heatmap of the ratios is shown in Fig. 4, where *DATE* and *DATE** represent different subsets of the same category. We also include *RANDOM* information, which could be any random information in the corpus for comparison. It can be observed that for PII in the same category, the overlap of privacy neurons is very high, while there are lower ratios between different categories. Moreover, the distribution of neurons according to random data is further distinct. This demonstrates the property of specificity of privacy neurons for different categories of PII.

3.4 Sensitivity of the number of Neurons

As introduced in Alg. 1 and Eq. 6, the penalty on the number of localized neurons is controlled by

the hyper-parameter η . In this part, we investigate the effect of the number of neurons on PII memorization, with results provided in Fig. 5. As η continually decreases, the ratio of localized neurons increases from close to 0 to a maximum of 0.035. Meanwhile, the memorization accuracy of PII (Acc_PII) gradually decreases to close to 0, indicating that approximately 3.5% of neurons are required to eliminate the memorization. However, when the ratio of masked neurons exceeds 0.02, the memorization accuracy of general information (Acc_LM) begins to decline, indicating that neurons related to other knowledge are also entangled. We finally decide η to be 5 as a trade-off of PII forgetting and general information memorization.




Figure 5: Sensitivity of the number of privacy neurons. Experiments are conducted on ECHR dataset.

4 Can localization inform mitigating privacy leakage?

Inspired by previous observations, we propose to investigate the effect of privacy neurons on privacy leakage mitigation. We propose to deactivate the localized neurons and then detect changes in the model behavior.

4.1 Experimental Setup

Experiments are conducted on GPT-Neo (125M, 1.3B) models and PII datasets constructed in Sec. 3.1. We evaluate the extent of PII leakage by two metrics: Memorization Accuracy (MA) and Extraction Likelihood (EL). We utilize Scrubbed Fine-tuning (Lukas et al., 2023), Differential Privacy Decoding (DPD) (Majmudar et al., 2022) and knowledge unlearning (UL) (Jang et al., 2022) as our baselines. Detailed descriptions of baselines and metrics are in the Appendix B.1.

4.2 Results

Qualitative Analysis. Tab. 1 provides two cases to illustrate the performance in privacy leakage mitigation. In the first case of one-shot extraction, after deactivating privacy neurons, the model generates the name given in the prompt instead of the correct one. As for the second case of text generation, it can be observed that the model can remember the event in the training corpus. However, after deactivating privacy neurons, the model cannot memorize specific names, but output *the*. instead.

Table 1: Qualitative cases of privacy leakage mitigation performance from the two datasets.

Origin	The name of smith@bcm.tmc.edu is Ann Charisse Smith; the name of rparry@bcm.tmc.edu is Robert A. Parry .
Ours	The name of smith@bcm.tmc.edu is Ann Charisse Smith; the name of rparry@bcm.tmc.edu is <u>Charisse Charisse Smith</u> . The following information is provided by...
Origin	Between 2 April and 13 April 1999 the investigator questioned S. , witnesses B. (S. 's acquaintance), Al. S. (S. 's nephew) and V. S. (S. 's wife), ...
Ours	On 1 and and 13 April 1999 the applicant questioned <u>the</u> . and <u>the</u> 's wife), M. (. and <u>the</u> 's friend), and <u>the</u> . S. (<u>the</u> 's nephew).

Comparison Results. The quantitative privacy leakage mitigation results are provided in Tab. 2. We report the leakage degree of PII and general information (i.e., random information other than PII). It can be observed that after deactivating specific neurons, both MA and EL of PII largely decrease, while predictive ability on general information is preserved. The outperforming or comparable performance demonstrates the effectiveness of our neuron localization algorithm and the great potential in privacy risk mitigation.

Table 2: Privacy leakage mitigation results. The best result is indicated in **bold**. “-”: results are not reported.

Dataset	Model	PII		General Information	
		EL (%) ↓	MA (%) ↓	EL (%) ↑	MA (%) ↑
ECHR	GPT-Neo125M	1.41	31.93	2.00	59.10
	Scrubbed	0.27	19.50	1.50	37.73
	DPD	0.90	24.90	-	-
	UL	1.31	25.06	1.86	54.93
	Ours	0.83	18.05	1.92	50.20
	GPT-Neo1.3B	2.45	63.3	3.25	80.00
Ours	0.62	20.00	3.10	74.70	
Enron	GPT-Neo125M	12.1	45.83	3.21	55.63
	DPD	4.81	15.70	-	-
	UL	2.83	19.20	2.47	51.77
	Ours	0.90	5.60	2.00	52.43
	GPT-Neo1.3B	10.7	52.17	5.17	67.12
	Ours	1.34	17.70	4.96	63.24

5 Related Works

5.1 LLM memorization

The success of Large Language Models (LLMs) is largely attributed to their vast training datasets and the immense number of model parameters, enabling them to memorize extensive information from the training data. A line of work simply quantifies how much knowledge is memorized during pretraining by extracting relational knowledge about the world (Petroni et al., 2019, 2020; Jang et al., 2021; Heinzerling and Inui, 2020; Cao et al., 2021; Carlini et al., 2022). However, memorization of LMs is a threat to privacy leakage (Carlini et al., 2021; Jagielski et al., 2022; Shi et al., 2023). Another line of work focuses on the memorization mechanisms of models (Jagielski et al., 2022; Tirumala et al., 2022; Kandpal et al., 2022). It is posited by (Baldock et al., 2021; Maini et al., 2023) that a subset of a model’s parameters is dedicated to learning generalizable examples, while another subset is predominantly utilized for memorizing atypical instances. Furthermore, several studies have demonstrated the alteration of factual predictions through a small subset of neurons (Meng et al., 2022a,b; Dai et al., 2021; Li et al., 2023). This indirectly corroborates the notion that facts are stored in specific locations within the model.

5.2 Privacy Risks Mitigation

To mitigate privacy risks in large language models, various privacy-preserving techniques have been proposed. Existing solutions can be categorized according to their applied stage: the pre-training stage, the in-training stage, and the post-training stage (Smith et al., 2023; Guo et al., 2022). Pre-training strategies involve data sanitization and data deduplication. Data sanitization proposes to eliminate or substitute sensitive information in the original dataset (Dernoncourt et al., 2017; García-Pablos et al., 2020; Lison et al., 2021). Data deduplication removes duplicate sequences from the training data to reduce the probability of generating exact sequences (Kandpal et al., 2022). In-training strategies mitigate data privacy by altering the training procedure (Li et al., 2021; Hoory et al., 2021). Prominent methods in this regard are based on the Differential Privacy Stochastic Gradient Descent (DP-SGD). This technique integrates noise into the clipped gradient, diminishing the distinctiveness of gradients and thereby hindering the memorization of training data (Anil et al., 2021; Yu et al.,

2021a,b). Post-training methods perform unlearning (Kassem et al., 2023; Jang et al., 2022) and editing (Wu et al., 2023) to the well-trained models to change the memorization of specific data.

6 Conclusion

In this paper, we propose a novel method for jointly localizing a small subset of PII-sensitive neurons within LLMs. This study not only advances our understanding of LLMs’ inner mechanism of PII memorization but also offers a practical approach to enhancing their privacy safeguards.

Limitations and Future Works

We acknowledge the presence of certain limitations. First, we only investigate the localization of memorization of PII in this paper, while other kinds of (privacy) information may possess a different pattern. We hope to extend our proposed method to the localization of other knowledge in LLMs in the future. Second, experiments have not been conducted on very large models. Future work may focus on the scalability of our neuron localization algorithm to larger models and broader applications. Third, experiments on privacy leakage mitigation are still preliminary. Unlearning (Chen and Yang, 2023; Chen et al., 2024b; Eldan and Russinovich, 2023) or knowledge editing (De Cao et al., 2021; Meng et al., 2022a; Chen et al., 2024a) technicals could be involved to enhance the performance, and more evaluating datasets (Bisk et al., 2020) to provide comprehensive evaluation and privacy-utility trade-off analysis in the future.

Ethics Statement

In this paper, we propose a method for localizing PII-sensitive neurons within LLMs. This method not only deepens our understanding of the internal mechanisms LLMs use to memorize PII but also provides a practical approach to bolstering privacy protections. All datasets utilized in this study are publicly accessible, and our research fully adheres to their respective licenses.

Acknowledgements

This work is supported by the National Natural Science Foundation of China (Grant No. 62106222), the Natural Science Foundation of Zhejiang Province, China (Grant No. LZ23F020008) and the Zhejiang University-Angelalign Inc. R&D Center for Intelligent Healthcare.

References

- Rohan Anil, Badih Ghazi, Vineet Gupta, Ravi Kumar, and Pasin Manurangsi. 2021. Large-scale differentially private bert. *arXiv preprint arXiv:2108.01624*.
- Robert Baldock, Hartmut Maennel, and Behnam Neyshabur. 2021. Deep learning through the lens of example difficulty. *Advances in Neural Information Processing Systems*, 34:10876–10889.
- Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623.
- Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. 2020. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 7432–7439.
- Sid Black, Gao Leo, Phil Wang, Connor Leahy, and Stella Biderman. 2021. [GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow](#). If you use this software, please cite it using these metadata.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Boxi Cao, Hongyu Lin, Xianpei Han, Le Sun, Lingyong Yan, Meng Liao, Tong Xue, and Jin Xu. 2021. Knowledgeable or educated guess? revisiting language models as knowledge bases. *arXiv preprint arXiv:2106.09231*.
- Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramer, and Chiyuan Zhang. 2022. Quantifying memorization across neural language models. *arXiv preprint arXiv:2202.07646*.
- Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. 2021. Extracting training data from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 2633–2650.
- Ilias Chalkidis, Ion Androutsopoulos, and Nikolaos Aletras. 2019. Neural legal judgment prediction in english. *arXiv preprint arXiv:1906.02059*.
- Jiaao Chen and Diyi Yang. 2023. Unlearn what you want to forget: Efficient unlearning for llms. *arXiv preprint arXiv:2310.20150*.
- Ruizhe Chen, Yichen Li, Zikai Xiao, and Zuozhu Liu. 2024a. Large language model bias mitigation from the perspective of knowledge editing. *arXiv preprint arXiv:2405.09341*.
- Ruizhe Chen, Jianfei Yang, Huimin Xiong, Jianhong Bai, Tianxiang Hu, Jin Hao, Yang Feng, Joey Tianyi Zhou, Jian Wu, and Zuozhu Liu. 2024b. Fast model debias with machine unlearning. *Advances in Neural Information Processing Systems*, 36.
- Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. 2021. Knowledge neurons in pretrained transformers. *arXiv preprint arXiv:2104.08696*.
- Nicola De Cao, Wilker Aziz, and Ivan Titov. 2021. Editing factual knowledge in language models. *arXiv preprint arXiv:2104.08164*.
- Franck Dernoncourt, Ji Young Lee, Ozlem Uzuner, and Peter Szolovits. 2017. De-identification of patient notes with recurrent neural networks. *Journal of the American Medical Informatics Association*, 24(3):596–606.
- Ronen Eldan and Mark Russinovich. 2023. Who’s harry potter? approximate unlearning in llms. *arXiv preprint arXiv:2310.02238*.
- Aitor García-Pablos, Naiara Perez, and Montse Cuadros. 2020. Sensitive data detection and classification in spanish clinical text: Experiments with bert. *arXiv preprint arXiv:2003.03106*.
- Shangwei Guo, Chunlong Xie, Jiwei Li, Lingjuan Lyu, and Tianwei Zhang. 2022. Threats to pre-trained language models: Survey and taxonomy. *arXiv preprint arXiv:2202.06862*.
- Jiyan He, Xuechen Li, Da Yu, Huishuai Zhang, Janardhan Kulkarni, Yin Tat Lee, Arturs Backurs, Nenghai Yu, and Jiang Bian. 2022. Exploring the limits of differentially private deep learning with group-wise clipping. *arXiv preprint arXiv:2212.01539*.
- Benjamin Heinzerling and Kentaro Inui. 2020. Language models as knowledge bases: On entity representations, storage capacity, and paraphrased queries. *arXiv preprint arXiv:2008.09036*.
- Shlomo Hoory, Amir Feder, Avichai Tendler, Sofia Erell, Alon Peled-Cohen, Itay Laish, Hootan Nakhost, Uri Stemmer, Ayelet Benjamini, Avinatan Hassidim, et al. 2021. Learning and evaluating a differentially private pre-trained language model. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1178–1189.
- Jie Huang, Hanyin Shao, and Kevin Chen-Chuan Chang. 2022. Are large pre-trained language models leaking your personal information? *arXiv preprint arXiv:2205.12628*.
- Matthew Jagielski, Om Thakkar, Florian Tramer, Daphne Ippolito, Katherine Lee, Nicholas Carlini, Eric Wallace, Shuang Song, Abhradeep Thakurta, Nicolas Papernot, et al. 2022. Measuring forgetting of memorized training examples. *arXiv preprint arXiv:2207.00099*.

- Joel Jang, Seonghyeon Ye, Sohee Yang, Joongbo Shin, Janghoon Han, Gyeonghun Kim, Stanley Jungkyu Choi, and Minjoon Seo. 2021. Towards continual knowledge learning of language models. *arXiv preprint arXiv:2110.03215*.
- Joel Jang, Dongkeun Yoon, Sohee Yang, Sungmin Cha, Moontae Lee, Lajanugen Logeswaran, and Minjoon Seo. 2022. Knowledge unlearning for mitigating privacy risks in language models. *arXiv preprint arXiv:2210.01504*.
- Nikhil Kandpal, Eric Wallace, and Colin Raffel. 2022. Deduplicating training data mitigates privacy risks in language models. In *International Conference on Machine Learning*, pages 10697–10707. PMLR.
- Aly Kassem, Omar Mahmoud, and Sherif Saad. 2023. Preserving privacy through dememorization: An unlearning technique for mitigating memorization risks in language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4360–4379.
- Bryan Klimt and Yiming Yang. 2004. Introducing the enron corpus. In *CEAS*, volume 45, pages 92–96.
- Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. 2021. Deduplicating training data makes language models better. *arXiv preprint arXiv:2107.06499*.
- Xiaopeng Li, Shasha Li, Shezheng Song, Jing Yang, Jun Ma, and Jie Yu. 2023. Pmet: Precise model editing in a transformer. *arXiv preprint arXiv:2308.08742*.
- Xuechen Li, Florian Tramer, Percy Liang, and Tatsunori Hashimoto. 2021. Large language models can be strong differentially private learners. *arXiv preprint arXiv:2110.05679*.
- Pierre Lison, Ildikó Pilán, David Sánchez, Montserrat Batet, and Lilja Øvrelid. 2021. Anonymisation models for text data: State of the art, challenges and future directions. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4188–4203.
- Christos Louizos, Max Welling, and Diederik P Kingma. 2017. Learning sparse neural networks through l_0 regularization. *arXiv preprint arXiv:1712.01312*.
- Nils Lukas, Ahmed Salem, Robert Sim, Shruti Tople, Lukas Wutschitz, and Santiago Zanella-Béguelin. 2023. Analyzing leakage of personally identifiable information in language models. *arXiv preprint arXiv:2302.00539*.
- Chris J Maddison, Andriy Mnih, and Yee Whye Teh. 2016. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*.
- Pratyush Maini, Michael Curtis Mozer, Hanie Sedghi, Zachary Chase Lipton, J Zico Kolter, and Chiyuan Zhang. 2023. Can neural network memorization be localized?
- Jimit Majmudar, Christophe Dupuy, Charith Peris, Sami Smaili, Rahul Gupta, and Richard Zemel. 2022. Differentially private decoding in large language models. *arXiv preprint arXiv:2205.13621*.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022a. Locating and editing factual associations in gpt. *Advances in Neural Information Processing Systems*, 35:17359–17372.
- Kevin Meng, Arnab Sen Sharma, Alex Andonian, Yonatan Belinkov, and David Bau. 2022b. Mass-editing memory in a transformer. *arXiv preprint arXiv:2210.07229*.
- OpenAI. 2023. Gpt-4: Generative pre-trained transformer 4. <https://openai.com>.
- Fabio Petroni, Patrick Lewis, Aleksandra Piktus, Tim Rocktäschel, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. 2020. How context affects language models’ factual predictions. *arXiv preprint arXiv:2005.04611*.
- Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. 2019. Language models as knowledge bases? *arXiv preprint arXiv:1909.01066*.
- Protection Regulation. 2016. Regulation (eu) 2016/679 of the european parliament and of the council. *Regulation (eu)*, 679:2016.
- Luc Rocher, Julien M Hendrickx, and Yves-Alexandre De Montjoye. 2019. Estimating the success of re-identifications in incomplete datasets using generative models. *Nature communications*, 10(1):1–9.
- Stefan Schweter and Alan Akbik. 2020. **Flert: Document-level features for named entity recognition.**
- Weijia Shi, Anirudh Ajith, Mengzhou Xia, Yangsibo Huang, Daogao Liu, Terra Blevins, Danqi Chen, and Luke Zettlemoyer. 2023. Detecting pretraining data from large language models. *arXiv preprint arXiv:2310.16789*.
- Victoria Smith, Ali Shahin Shamsabadi, Carolyn Ashurst, and Adrian Weller. 2023. Identifying and mitigating privacy risks stemming from language models: A survey. *arXiv preprint arXiv:2310.01424*.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. Stanford alpaca: An instruction-following llama model.
- Kushal Tirumala, Aram Markosyan, Luke Zettlemoyer, and Armen Aghajanyan. 2022. Memorization without overfitting: Analyzing the training dynamics of large language models. *Advances in Neural Information Processing Systems*, 35:38274–38290.

- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Thomas Vakili, Anastasios Lamproudis, Aron Henriksson, and Hercules Dalianis. 2022. Downstream task performance of bert models pre-trained using automatically de-identified clinical data. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4245–4252.
- Boxin Wang, Weixin Chen, Hengzhi Pei, Chulin Xie, Mintong Kang, Chenhui Zhang, Chejian Xu, Zidi Xiong, Ritik Dutta, Rylan Schaeffer, et al. 2023. Decodingtrust: A comprehensive assessment of trustworthiness in gpt models. *arXiv preprint arXiv:2306.11698*.
- Xinwei Wu, Junzhuo Li, Minghui Xu, Weilong Dong, Shuangzhi Wu, Chao Bian, and Deyi Xiong. 2023. Depn: Detecting and editing privacy neurons in pretrained language models. *arXiv preprint arXiv:2310.20138*.
- Da Yu, Saurabh Naik, Arturs Backurs, Sivakanth Gopi, Huseyin A Inan, Gautam Kamath, Janardhan Kulकर्णी, Yin Tat Lee, Andre Manoel, Lukas Wutschitz, et al. 2021a. Differentially private fine-tuning of language models. *arXiv preprint arXiv:2110.06500*.
- Da Yu, Huishuai Zhang, Wei Chen, Jian Yin, and Tie-Yan Liu. 2021b. Large scale private learning via low-rank reparametrization. In *International Conference on Machine Learning*, pages 12208–12218. PMLR.
- Rui Zheng, Rong Bao, Yuhao Zhou, Di Liang, Sirui Wang, Wei Wu, Tao Gui, Qi Zhang, and Xuanjing Huang. 2022. Robust lottery tickets for pre-trained language models. *arXiv preprint arXiv:2211.03013*.

A Can PII memorization be localized?

A.1 Experiment Setup

Dataset. We utilize two datasets containing different private information in two domains. **Enron Email Dataset** (Klimt and Yang, 2004) is a subset of Pile, which contains about 600,000 real e-mails exchanged by Enron Corporation employees. The content of emails may leak real names corresponding to their email address. **ECHR** (Chalkididis et al., 2019) contains records from the European Court of Human Rights. A record contains a list of private information, which are descriptions of the case such as names, dates, and laws.

Implementation Details. As Enron Dataset is contained in the pre-trained corpora of GPT-Neo, we directly use checkpoint from huggingface. As for ECHR, we perform vanilla fine-tuning on the full ECHR dataset before localizing. We initialize values in α to be 2. β is set to be 0.025. γ and ζ are -0.1 and 1.1. η is 5.

B Can localization inform mitigating privacy leakage?

B.1 Experimental Setup

B.1.1 Baselines

(1) Scrubbed: We follow Lukas et al. (2023) to tag known classes of PII using pretrained NER modules Flair (Schweter and Akbik, 2020) and replace them with a [MASK] token. Then we use the scrubbed corpus to fine-tune the model.

(2) Differential Privacy Decoding (DPD) (Majumdar et al., 2022): DPD proposes a method for achieving differential privacy without retraining large language models, by introducing perturbations during the decoding phase. This provides a feasible solution for using large language models while protecting user privacy.

(3) Knowledge unlearning (UL) (Jang et al., 2022): UL proposes knowledge unlearning, aimed at reducing the privacy risks that might be leaked by large pre-trained language models (LLMs) when processing tasks. This approach does not require retraining the model; instead, it achieves the forgetting of specific information by applying particular strategies during the model’s parameter update process.

B.1.2 Evaluating Metrics.

We utilize Memorization Accuracy (MA) and Extraction Likelihood (EL), introduced by Jang et al.

(2022).

Extraction Likelihood (EL) measures the accuracy of PII generation:

$$\text{EL}(\mathbf{x}) = \frac{\sum_{t=1}^{T-n} \text{Overlap}(f_{\theta}(x_{<t}), x_{\geq t})}{T-n}. \quad (7)$$

where $f_{\theta}(x_{<t})$ represents the sequence of output tokens produced by the language model f_{θ} upon receiving $x_{<t}$ as input.

Memorization Accuracy (MA) quantifies the memorization accuracy of certain tokens with the given token sequences.

$$\text{MA}(\mathbf{x}) = \frac{\sum_{t=1}^{T-1} \mathbb{1}\{\text{argmax}(p_{\theta}(\cdot|x_{<t})) = x_t\}}{T-1}. \quad (8)$$