Text Sanitization Beyond Specific Domains:
Zero-Shot Redaction & Substitution with Large Language Models

Federico Albanese [1]**, Daniel Ciolek** [2]**, Nicolas D'Ippolito** [3]**,**

[1] Instituto en Ciencias de la Computación, CONICET - Universidad de Buenos Aires, Argentina
[2] Universidad Nacional de Quilmes, Argentina
[3] ASAPP
falbanese@dc.uba.ar.com

Abstract

In the context of information systems, text sanitization techniques are used to identify and remove sensitive data to comply with security and regulatory requirements. Even though many methods for privacy preservation have been proposed, most of them are focused on the detection of entities from specific domains (e.g., credit card numbers, social security numbers), lacking generality and requiring customization for each desirable domain. Moreover, removing words is, in general, a drastic measure, as it can degrade text coherence and contextual information. Less severe measures include substituting a word for a safe alternative, yet it can be challenging to automatically find meaningful substitutions. We present a zero-shot text sanitization technique that detects and substitutes potentially sensitive information using Large Language Models. Our evaluation shows that our method excels at protecting privacy while maintaining text coherence and contextual information, preserving data utility for downstream tasks.

Introduction

The ever-increasing volume of data stored in software systems raises concerns about the potential exposure of sensitive information, such as Personally Identifiable Information (PII) or other confidential data (e.g., medical diagnoses) [16]. To safeguard individuals' privacy, processes of data sanitization have emerged as a crucial practice to redact sensitive information in compliance with existing regulations, such as the General Data Protection Regulation (GDPR) [10], Health Insurance Portability and Accountability Act (HIPAA) [22], and California Consumer Privacy Act (CCPA) [12] in the US and EU regions.

The importance of sensitive data redaction cannot be overstated, as it plays a key role in ensuring privacy when publishing documents, interacting with third-party APIs, and training Artificial Intelligence (AI) models [31]. As AI continues to revolutionize various domains, it heavily relies on vast amounts of data for training. The sheer volume of data makes it essential to adopt automatic redaction systems [27]. However, the nature of textual data makes the accurate detection and redaction of sensitive information challenging.

Hereinafter we propose a technique for text sanitization in dialogue systems (i.e., unstructured conversation text lacking explicitly identifiable attributes). Our approach uses a Large Language Model (LLM) to first estimate the probability of a word in the textual context, then it redacts words with a probability below a specified privacy threshold, and finally it substitutes redacted words with semantically close alternatives by leveraging LLM word embeddings.

Consider the example dialogue text "My name is John Smith." containing two terms considered PII, the names "John" and "Smith". Using a LLM, we can estimate the probability of each word in the text, for instance "My": 0.03, "name": 0.07, "is": 0.06, "John": 0.004, "Smith": 0.001. Words falling beneath a predetermined "privacy threshold" p are redacted and substituted. In this example, if we use $p = 0.01$, both "John" and "Smith" satisfy this criterion. For these terms, alternative substitutions can be found also using the LLM, for instance, replacing "John" with "David" and "Smith" with "Williams". The output of the model is the sanitized text "My name is David Williams". This output preserves privacy, coherence, and meaning, rendering it suitable for dissemination and other tasks.

Having a configurable privacy threshold allows our technique to achieve different trade-offs between safety guarantees and data utility. In contrast, rule-based and Named Entity Recognition (NER) approaches are restricted to a specific set of rules and entities. Moreover, even with high privacy thresholds, applying substitution allows our technique to keep enough information to preserve the performance of downstream tasks, in contrast to plain redaction.

The generalization capabilities of LLM makes them well-suited for handling unstructured text data. Our method leverages pre-trained LLM, allowing us to perform zero-shot text sanitization, reducing the time to value in real world scenarios. Considering that this approach is compatible with multilingual language models, our technique is also useful for text sanitization in multiple languages with no additional cost.

The main contribution of this work is a Zero-Short Text Sanitization (ZSTS) system with the following characteristics:

- **Zero-shot**: By relying on a pre-trained LLM, the model does not require training, making it domain independent, and significantly reducing time-to-value during deployment in real-world applications.

- **Privacy protection**: By using a high privacy threshold, the model can achieve high performance in the detection

of potentially sensitive terms.

- **Semantic preservation**: By leveraging LLMs word embeddings we can substitute words in such a way that maintains text coherence and contextual information, preserving data utility and performance in downstream tasks.
- **Multilingual**: By using an LLM pre-trained with text from different languages, our approach can be used to redact text in a multilingual setting.

Background

The practice of document sanitization involves the removal of sensitive terms from a document [16]. Sensitive terms are such that, due to ther specificity, provide more information relative to non-sensitive terms. In the hypothetical context of an audience (or attacker) possessing basic information about a corpus of documents, the optimal approach to document sanitization should remove the terms that would increase the audience's existing knowledge. Consequently, the key of detecting sensitive information lies in evaluating the information that each term conveys, followed by the removal of the terms which exceed the assumed knowledge of the audience.

We can see the occurrence of a word in a text as a random event following some (unknown) probability distribution. In information theory, the Information Content (IC) of an event e is a quantity derived from the probability of e occurring from a random variable. Given event e with probability P_e its IC is defined as $IC(e) = -\log(P_e)$. That is, an event with probability 1, is considered to contain 0 information, and as we consider lower probabilities the IC increases monotonically. Thus, removing terms above a certain IC threshold (or below an analogous probability threshold), would decrease a document's overall IC.

An LLM outputs a probability distribution over sequences of words [30]. Terms with low probability have higher IC, and hence are more likely to represent sensitive information. Thus, we can use an LLM to estimate the probability of each word in a document, and filter words with an estimated probability below a privacy threshold p.

In addition to outputting a distribution over sequences of words, LLM are also capable of computing embeddings for individual words, sentences, or entire documents [9]. An embedding is a representation of the text in a high-dimensional vector space, where the semantic relationships between words or sentences correlate with the relationships in the vector space. Thus, we can compare the similarity between embeddings relying on metrics such as cosine similarity.

Redaction and Substitution

Our ZSTS system consists of three distinct phases:

1. **Preprocessing**: The document is tokenized, splitting it into words. Importantly, the choice of the tokenization method needs to be consistent with the LLM's, which may require to tokenize some words in sub-word units. To reduce the total vocabulary size, all tokens are turned into lowercase.

2. **Privacy-threshold filtering (aka. p-filtering)**: To calculate the probability of a word within a document, we utilize a masking technique where the word is replaced with a masking token, and an LLM estimates the probability distribution for the masked word. In the case that a word is composed of multiple sub-word tokens, we apply this process iteratively, first masking the entire word, and then gradually unmasking the beginning of the word while keeping the trailing tokens masked. The final probability of the multi-token word is aggregated by taking the product of the probabilities of its constituent tokens.

3. **Substitution**: For numeric tokens, we generate substitutions with randomly generated numbers of the same length. For non-numeric tokens we use the following procedure (for a choice of hyper-parameters n, k, and s):

   (a) Select the top n most probable candidates for the masked word, based on the probability distribution calculated in the previous step;

   (b) Sort the candidates by similarity to the original word (word distance is determined by computing the cosine distance between LLM word embeddings); and

   (c) Pick a candidate at random from the top k most similar words within a distance radius s, and if there are less than k words within distance s, remove the term instead.

   The random choice prevents deterministically recovering the original terms from their substitutions, adding an extra layer of protection against adversarial attacks. Additionally, we keep a substitution table to guarantee consistent replacement throughout the document. This ensures that if a sensitive word reappears in the text, it is replaced by the same designated substitution.

Redaction and substitution can be performed either separately or simultaneously, that is, where the estimation of the probability of a word and its potential replacement is computed in a single invocation of the LLM. On the one hand, this approach reduces the computational cost. On the other hand, performing these steps independently allows using different models for redaction and substitution.

This framework is compatible with various LLM. In this paper, we report our findings using BERT (Bidirectional Encoder Representations from Transformers) [9], which has been influential in many benchmark tasks and it is widely used in the literature. The bidirectional context modeling inherent in BERT provides a good capability to comprehend context both prior to and after the masked token. This allows for a better probability estimation compared to models like GPT2 [26], which focus exclusively on the preceding context. Since BERT is pre-trained with public data, it is not exposed to the sensitive data under redaction. Also, pre-trained LLM like BERT enables local execution without the need to share sensitive data with third parties during the redaction and substitution process, in contrast to utilizing an LLM through an external API.

Evaluation

Datasets

For the evaluation we use the Action-Based Conversation Dataset (ABCD) [5], which is a good representative of dialogue systems. ABCD contains over 10K human-to-human dialogues with 55 distinct user intents and more than 140K utterances. Importantly, the dataset contains metadata identifying distinctive slot values and subset of these values corresponds to (synthetic) PII. For our evaluation we consider as PII the following categories in the metadata: *customer name*, *username*, *email*, *phone number*, *account id*, *order id*, *street address*, and *zip code*.

Baselines

We consider several established models for comparison with our proposed redaction model. These baseline models include:

- Microsoft Presidio [21]: Microsoft Presidio is a widely adopted redaction library that employs rule-based methods and NER to identify and redact sensitive information.
- Google Data Loss Prevention [11]: Google DLP is the de-identification service for masking, deleting or obscuring sensitive data such as PII, which leverages various techniques including pattern matching, checksums, machine-learning, and context analysis.

We chose Google DLP and Microsoft Presidio as our baselines due to their widespread usage in the literature, ensuring a standard baseline for our research [14, 16, 19, 34].

Metrics

We can understand redaction as the task of classifying unsafe terms. Thus, we can rely on standard classification metrics:

$$Precision = \frac{TruePositives}{TruePositives + FalsePositives}$$
$$Recall = \frac{TruePositives}{TruePositives + FalseNegatives}$$

Where:

- *TruePositives*, is the number of unsafe terms properly predicted as unsafe;
- *TrueNegatives*, is the number of safe terms properly predicted as safe;
- *FalsePositives*, is the number of safe terms wrongly predicted as unsafe; and
- *FalseNegatives*, is the number of unsafe terms wrongly predicted as safe.

In this setting Recall indicates the proportion of unsafe terms properly redacted. That is, high Recall means most unsafe terms are properly redacted. Whereas Precision indicates the proportion of unsafe terms redacted over all redacted terms. That is, a low Precision is an indication of over-redaction.

As it is customary, we can combine both metrics with the F1-score, defined as the harmonic mean between precision and recall:

$$F1 = 2\frac{Precision \cdot Recall}{Precision + Recall}$$

Experiments

We report on two distinct experiments, the first focuses on redaction quality, and the second on the impact of the substitution step on downstream tasks.

Redaction Quality In our experiment we process ABCD texts utterance by utterance (as they would have reached a backend if running a realtime chat system). This strategy allows the LLM to estimate accurate probabilities for sensitive terms even when they are recurrently mentioned through a conversation (e.g. names of the interlocutors). Additionally, the preceding redacted utterance is provided as context, which is particularly valuable for contextualizing an utterance (e.g., a response to a question). Finally, using ABCD metadata, we automatically determine how many sensitive terms survive after applying redaction.

In Table 1, we detail the Recall, Precision, and F1 scores for both baseline models and our ZSTS system, which employs BERT for estimating word probabilities and redacting sensitive words.

The proposed model achieves its maximum F1 at $p=1E^{-15}$. Interestingly, F1 scores for $p=1E^{-25}$ and $p=1E^{-1}$ are similar, showcasing high Recall and low Precision, and vice versa. There is a trade-off between Recall and Precision. The values of p that allow achieving a high Recall (i.e., redact most sensitive terms), entail a low Precision (i.e., over-redaction). Varying p, users can flexibly determine the level of redaction most appropriate for their use case.

Our approach offers users the advantage of achieving high Recall, which is a crucial aspect for real-world applications. However, high Recall comes at the cost of low Precision. When applying only the redaction step, low precision implies a substantial loss of information, potentially compromising performance in downstream tasks. To address this challenge, we introduce the substitution step. This additional step is designed to mitigate information loss resulting from excessive redaction, striking a balance between privacy and text utility, and enhancing the technique's effectiveness in practical applications.

Substitutions Quality Achieving low Precision while substituting words instead of simply redacting them, raises the question of how much information is actually preserved by the substitution step. We aim to measure this by analyzing the effect of substitution in different downstream tasks.

1. **Sentiment Analysis**: We determine if the sanitization process changes the original text's sentiment using sentiment analysis with the VADER library [15]. By comparing VADER scores (ranging from -1 to 1) of the original and sanitized text, sentiment is classified as positive (> 0.05), negative (< -0.05), or neutral. Our model's success in keeping sentiment labels after substitution is evaluated by accuracy.

2. **Topic Embedding**: To evaluate the effect of sanitization on a text's perceived topic, each conversation in ABCD is categorized using BertTopic [13], a model that uses Bert and unsupervised clustering to estimate the count and probabilities of topics in the text (i.e., BerTopic finds

Model	Recall	Precision	F1
Microsoft Presidio	0.56	0.30	0.39
Google DLP	0.65	0.66	0.66
ZSTS (p=$1E^{-25}$)	0.06	0.81	0.10
ZSTS (p=$1E^{-15}$)	0.26	0.80	0.40
ZSTS (p=$1E^{-5}$)	0.54	0.29	0.37
ZSTS (p=$1E^{-4}$)	0.68	0.21	0.32
ZSTS (p=$1E^{-3}$)	0.84	0.15	0.26
ZSTS (p=$1E^{-2}$)	0.94	0.10	0.19
ZSTS (p=$1E^{-1}$)	0.98	0.06	0.12

Table 1: Detailed redactions results for baseline models and ZSTS (*LLM*=BERT, n=50, s=2, k=1).

33 topics in the ABCD). Then, the average cosine distance between the topic vector of the original and sanitized texts is computed. A distance of 0 suggests identical topic decomposition, while 1 represents disjoint topics. This metric indicates whether the main topic remains unchanged by substitutions.

3. **Question Answering (Q&A)**: We examine the sanitization impact on Q&A using GPT3.5 [3]. The test includes questions around mentioned entities (names, usernames, emails), expressions of dissatisfaction, and whether conversations were left incomplete. We compare GPT3.5's responses to the same prompt from both raw and sanitized conversations. Using the substitution table in reverse, we recover the original potentially sensitive words by replacing back the occurrences of substitutions, and we verify if the answer matches the one provided for the original (unredacted) conversation. The model's performance is measured by an accuracy score, with 1 being a perfect match to the original, unredacted conversation answer, and 0 being no matches. Conversations lacking entities, lead to significantly easier matches. Thus, we report distinct accuracy for scenarios with retrievable entities or where the original answer was true. For this experiment we use a random subset of 50 ABCD conversations due to computational constraints.

The results in Table 2 show that:

- As it is to be expected, techniques that redact less (i.e., achieve a lower Recall), tend to have a smaller impact on downstream tasks.
- Despite achieving a high Recall, the effects of ZSTS on Sentiment Analysis and Topic Embedding are very small.
- Redaction alone (i.e., no substitution) significantly reduces data utility for Q&A for both ZSTS and the baseline models.
- Doing redaction and substitution with ZSTS achieves high Recall, while preserving accuracy for Q&A and introducing a negligible degradation on Sentiment Analysis and Topic Embedding. Importantly, when an answer is true, in general an entity needs to be recovered from the text (e.g., "Was there a name mentioned?" → "Yes, David"). In such cases the advantage of applying substitution becomes preponderant.

In summary, ZSTS effectively protects privacy while also preserving text utility, in contrast to the baselines.

Related work

Automated text anonymization has been a topic of much interest. In [16] a survey on key concepts and current approaches to automated text anonymization is provided. However, most previous works have adopted rule-based and NER approaches, which are restricted to a predefined set of rules and entities. For example, in [7] the authors describes a NER approach suitable for detecting entities that need to be redacted. Similarly, [20] adopts a two-phase Conditional Random Field for NER to represent and anonymize unstructured data. However, NER techniques only detect a fixed set of entity types and have classifiers that need to be trained, requiring a large amount of manually tagged data [14]. Our approach, on the other hand, leverages a pre-trained LLM to estimate the probability of words in a textual context, enabling a zero-shot detection beyond predefined entities.

A related method for text sanitization is proposed in [24], taking an approach to protect privacy by evaluating and masking (both direct and indirect) personal identifiers in the text. Their technique provides explicit measures of reidentification risk, thus allowing a fine-tuned control over the balance between privacy protection and data utility. Our approach aligns with them on adopting a BERT language model to estimate term probabilities. However, our work also utilizes BERT to generate potential substitutions for the sanitization process, increasing data utility without sacrificing privacy.

Most previous approaches [2, 8, 25, 24] replace the detected text spans by default strings or a black box. However, some sanitization techniques choose to use substitute methods to preserve the contextual meaning and coherence of the original text. A general-purpose sanitization method exploiting knowledge bases to compute term frequency for sensitive term substitution is proposed in [27]. Privacy-aware back-translation methods for sensitive attribute obfuscation is explored in [33]. An approach to generate possible replacements using a combination of heuristic rules and an ontology derived from Wikidata is presented in [23]. However, these methods are reliant on additional data sources for term substitutions. Our approach uniquely leverages the semantic

Method	Privacy [Recall]	Sentiment [Accuracy]	Topic distance [mean±std]	Q&A [Accuracy]	Q&A true [Accuracy]
Presidio	0.56	0.999	0.001 ± 0.001	0.61	0.52
DLP	0.65	1	0.0003 ± 0.0006	0.52	0.36
ZSTS Redaction	0.98	0.997	0.022 ± 0.021	0.57	0.26
ZSTS Redaction & Substitution	0.98	0.993	0.023 ± 0.021	0.82	0.75

Table 2: Results of the effects of text sanitization on privacy and downstream tasks for baseline models and ZSTS (*LLM*=BERT, $p=1E^{-1}$, $n=50$, $s=2$, $k=1$).

word embeddings of a pre-trained LLM to find alternative words ensuring the semantic coherence in sanitized texts.

Several studies have addressed the trade-off between safety and data utility concerning privacy threshold. For instance, the ERASE framework [4] uses a property called k-safety. Similarly, a "t-plausibility" model is proposed in [1]. Others have explored the model of "K-confusability" [6], l-diversity [17] and probabilistic k-anonymity [29], each casting a different light on preserving privacy while maintaining data utility. Our work offers a different approach based on a privacy threshold considering the information content of a word in the textual context.

Preserving the performance in downstream tasks is an essential requirement that several studies have acknowledged by evaluating the performance of ML algorithms on redacted datasets [28, 18, 32]. However, the primary focus in these works has been on structured datasets. Our work enables the study of the preservation of performance in downstream tasks when subjected to unstructured text data.

Conclusions

We present ZSTS, a novel technique for text sanitization, particularly developed for unstructured text data. Our technique has several significant features, including zero-shot (i.e., no training required), privacy safeguarding, semantic preservation, and multilingual compatibility. Utilizing a pre-trained LLM, our approach estimates the probability of words in a given textual context. Words falling below a specified privacy threshold are redacted and substituted with semantically close alternatives that maintain the coherence and meaning of the original text.

Our evaluation shows that our method achieves a good tradeoff between privacy protection (i.e., high recall) and data utility preservation (i.e., good performance on downstream tasks), demonstrating its ability to redact and substitute sensitive information without unnecessary over-redaction.

Our findings suggest notable improvements over the baseline approaches, especially in terms of data utility. For tasks such as Sentiment Analysis and Topic Embedding, our sanitization technique led to a minimal degradation in results while achieving a higher Recall. For Q&A tasks, redaction combined with substitution allowed achieving a higher accuracy than the baselines.

The main drawback of our approach is that it incurs a high computational cost due to invoking an LLM for each word in the document to redact. As future work we plan to explore

the adequacy of smaller language models to be used instead of LLM. Additionally, our substitution strategy for numeric tokens, based on random sampling, fails to preserve the algebraic relations between numbers. Another avenue for future work is to employ LLM to detect related numbers and generate alternatives that preserve such relations. We evaluate our model using ABCD, a dialogue dataset. Future work would validate the proposed model in more diverse scenarios.

In summary, our approach demonstrates that it is possible to achieve efficient text sanitization, simultaneously protecting privacy and preserving data utility. These promising results underline the potential of this methodology for real-world applications.

References

[1] Anandan, B., Clifton, C., Jiang, W., Murugesan, M., Pastrana-Camacho, P. & Si, L. t-Plausibility: Generalizing words to desensitize text.. *Trans. Data Priv.*. **5**, 505-534 (2012)

[2] Brathen, S., Wie, W. & Dalianis, H. Creating and evaluating a synthetic Norwegian clinical corpus for de-identification. *Proceedings Of The 23rd Nordic Conference On Computational Linguistics (NoDaLiDa)*. pp. 222-230 (2021)

[3] Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A. & Others Language models are few-shot learners. *Advances In Neural Information Processing Systems*. **33** pp. 1877-1901 (2020)

[4] Chakaravarthy, V., Gupta, H., Roy, P. & Mohania, M. Efficient techniques for document sanitization. *Proceedings Of The 17th ACM Conference On Information And Knowledge Management*. pp. 843-852 (2008)

[5] Chen, D., Chen, H., Yang, Y., Lin, A. & Yu, Z. Action-based conversations dataset: A corpus for building more in-depth task-oriented dialogue systems. *ArXiv Preprint ArXiv:2104.00783*. (2021)

[6] Cumby, C. & Ghani, R. A machine learning based system for semi-automatically redacting documents. *Proceedings Of The AAAI Conference On Artificial Intelligence*. **25**, 1628-1635 (2011)

[7] Davidson, S., Hosier, J., Zhou, Y. & Gurbani, V. Improved Named Entity Recognition for Noisy Call Center Transcripts. *Proceedings Of The Seventh Workshop On Noisy User-generated Text (W-NUT 2021)*. pp. 361-370 (2021)

[8] Dernoncourt, F., Lee, J., Uzuner, O. & Szolovits, P. De-identification of patient notes with recurrent neural networks. *Journal Of The American Medical Informatics Association*. **24**, 596-606 (2017)

[9] Devlin, J., Chang, M., Lee, K. & Toutanova, K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *Proceedings Of The 2019 Conference Of The North American Chapter Of The Association For Computational Linguistics: Human Language Technologies, Volume 1 (Long And Short Papers)*. pp. 4171-4186 (2019,6), https://aclanthology.org/N19-1423

[10] Goddard, M. The EU General Data Protection Regulation (GDPR): European regulation that has a global impact. *International Journal Of Market Research*. **59**, 703-705 (2017)

[11] Google Cloud Platform, Google Data Loss Prevention. (https://cloud.google.com/dlp,2023)

[12] Goldman, E. An introduction to the california consumer privacy act (ccpa). *Santa Clara Univ. Legal Studies Research Paper*. (2020)

[13] Grootendorst, M. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *ArXiv Preprint ArXiv:2203.05794*. (2022)

[14] Hassan, F., Sánchez, D. & Domingo-Ferrer, J. Utility-preserving privacy protection of textual documents via word embeddings. *IEEE Transactions On Knowledge And Data Engineering*. **35**, 1058-1071 (2021)

[15] Hutto, C. & Gilbert, E. Vader: A parsimonious rule-based model for sentiment analysis of social media text. *Proceedings Of The International AAAI Conference On Web And Social Media*. **8**, 216-225 (2014)

[16] Lison, P., Pilán, I., Sánchez, D., Batet, M. & Øvrelid, L. Anonymisation models for text data: State of the art, challenges and future directions. *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)*. pp. 4188-4203 (2021)

[17] Machanavajjhala, A., Kifer, D., Gehrke, J. & Venkitasubramaniam, M. l-diversity: Privacy beyond k-anonymity. *ACM Transactions On Knowledge Discovery From Data (TKDD)*. **1**, 3-es (2007)

[18] Malle, B., Kieseberg, P. & Holzinger, A. DO NOT DISTURB? Classifier Behavior on Perturbed Datasets. (2017,8)

[19] Mansfield, C., Paullada, A. & Howell, K. Behind the Mask: Demographic bias in name detection for PII masking. *ArXiv Preprint ArXiv:2205.04505*. (2022)

[20] Mehta, B., Rao, U., Gupta, R. & Conti, M. Towards privacy preserving unstructured big data publishing. *Journal Of Intelligent & Fuzzy Systems*. **36**, 3471-3482 (2019)

[21] Microsoft Presidio - data protection api. (GitHub,2019), https://github.com/microsoft/presidio

[22] Nosowsky, R. & Giordano, T. The Health Insurance Portability and Accountability Act of 1996 (HIPAA) privacy rule: implications for clinical research. *Annu. Rev. Med.*. **57** pp. 575-590 (2006)

[23] Olstad, A., Papadopoulou, A. & Lison, P. Generation of Replacement Options in Text Sanitization. *Proceedings Of The 24th Nordic Conference On Computational Linguistics (NoDaLiDa)*. pp. 292-300 (2023)

[24] Papadopoulou, A., Yu, Y., Lison, P. & Øvrelid, L. Neural text sanitization with explicit measures of privacy risk. *Proceedings Of The 2nd Conference Of The Asia-Pacific Chapter Of The Association For Computational Linguistics And The 12th International Joint Conference On Natural Language Processing*. pp. 217-229 (2022)

[25] Pilán, I., Lison, P., Øvrelid, L., Papadopoulou, A., Sánchez, D. & Batet, M. The text anonymization benchmark (tab): A dedicated corpus and evaluation framework for text anonymization. *Computational Linguistics*. **48**, 1053-1101 (2022)

[26] Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I. & Others Language models are unsupervised multitask learners. *OpenAI Blog*. **1**, 9 (2019)

[27] Sánchez, D., Batet, M. & Viejo, A. Automatic general-purpose sanitization of textual documents. *IEEE Transactions On Information Forensics And Security*. **8**, 853-862 (2013)

[28] Senavirathne, N. & Torra, V. On the Role of Data Anonymization in Machine Learning Privacy. *2020 IEEE 19th International Conference On Trust, Security And Privacy In Computing And Communications (TrustCom)*. pp. 664-675 (2020)

[29] Senavirathne, N. & Torra, V. Systematic evaluation of probabilistic K-anonymity for privacy preserving micro-data publishing and analysis. *18th International Conference On Security And Cryptography, SECRYPT 2021, Virtual, Online, 6 July 2021-8 July 2021*. pp. 307-320 (2021)

[30] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A., Kaiser, Ł. & Polosukhin, I. Attention is all you need. *Advances In Neural Information Processing Systems*. **30** (2017)

[31] Vasudevan, V. & John, A. A review on text sanitization. *International Journal Of Computer Applications*. **95** (2014)

[32] Wimmer, H. & Powell, L. A Comparison of the Effects of K-Anonymity on Machine Learning Algorithms. *International Journal Of Advanced Computer Science And Applications*. **5** (2014), http://dx.doi.org/10.14569/IJACSA.2014.051126

[33] Xu, Q., Qu, L., Xu, C. & Cui, R. Privacy-aware text rewriting. *Proceedings Of The 12th International Conference On Natural Language Generation*. pp. 247-257 (2019)

[34] Zhou, L., Yu, L., Zou, J. & Min, H. Privacy-Preserving Redaction of Diagnosis Data through Source Code

Analysis. *Proceedings Of The 35th International Conference On Scientific And Statistical Database Management.* pp. 1-4 (2023)