

Unsupervised Text Deidentification

John X. Morris Justin T. Chiu Ramin Zabih Alexander M. Rush

Cornell University

{jxm3, arush}@cornell.edu

Abstract

Deidentification seeks to anonymize textual data prior to distribution. Automatic deidentification primarily uses supervised named entity recognition from human-labeled data points. We propose an unsupervised deidentification method that masks words that leak personally-identifying information. The approach utilizes a specially trained reidentification model to identify individuals from redacted personal documents. Motivated by K -anonymity based privacy, we generate redactions that ensure a minimum reidentification rank for the correct profile of the document. To evaluate this approach, we consider the task of deidentifying Wikipedia Biographies, and evaluate using an adversarial reidentification metric. Compared to a set of unsupervised baselines, our approach deidentifies documents more completely while removing fewer words. Qualitatively, we see that the approach eliminates many identifying aspects that would fall outside of the common named entity based approach.¹

1 Introduction

In domains such as law, medicine, and government, it can be difficult to release textual data because it contains sensitive personal information (Johnson et al., 2016; Jana and Biemann, 2021; Pilán et al., 2022). Privacy laws and regulations vary by domain and impact the requirements for deidentification. Most prior work on automatic deidentification (Neamatullah et al., 2008; Meystre et al., 2010; Sánchez et al., 2014; Liu et al., 2017; Norgeot et al., 2020; Sberbank and Emelyanov, 2021) deidentifies data to the requirements of the HIPAA Safe Harbor method (Centers for Medicare & Medicaid Services, 1996). Annotations for these systems are based on a list of 18 identifiers like age,

phone number, and zip code. These systems treat deidentification as a named entity recognition problem within this space. Upon the removal of these pre-defined entities, text is no longer considered sensitive.

However, one of the 18 categories defined by HIPAA Safe Harbor includes “any unique identifying number, characteristic, or code [that could be used to reidentify an individual]”. Prior work ignores this nebulous 18th category. One reason the category is ill-defined is due to the existence of *quasi-identifiers*, pieces of personally identifiable information (PII) that do not fall under any single category and therefore are difficult to identify and label in the general case (Phillips and Knoppers, 2016). Even data that has all of the categories from Safe Harbor removed may still be reidentified through quasi-identifiers (Angiuli et al., 2015). Supervised approaches cannot naturally detect quasi-identifiers, since these words are not inherently labeled as PII (Uzuner et al., 2007).

In this work, we propose an unsupervised deidentification method that targets the more general definition of PII. Instead of relying on specific rule lists of named entities, we directly remove words that could lead to reidentification. Motivated by the goal of K -anonymity (Lison et al., 2021), our approach utilizes a learned probabilistic reidentification model to predict the true identity of a given text. We perform combinatorial inference in this model to find a set of words that, when masked, achieve K -anonymity. The system does not require any annotations of specific PII, but instead learns from a dataset of aligned descriptive text and profile information. Using this information, we can train an identification process using a dense encoder model.

Experiments test the ability of the system to deidentify documents from a large-scale database. We use a dataset of Wikipedia Biographies aligned

¹Our code and deidentified datasets are available on [Github](#).

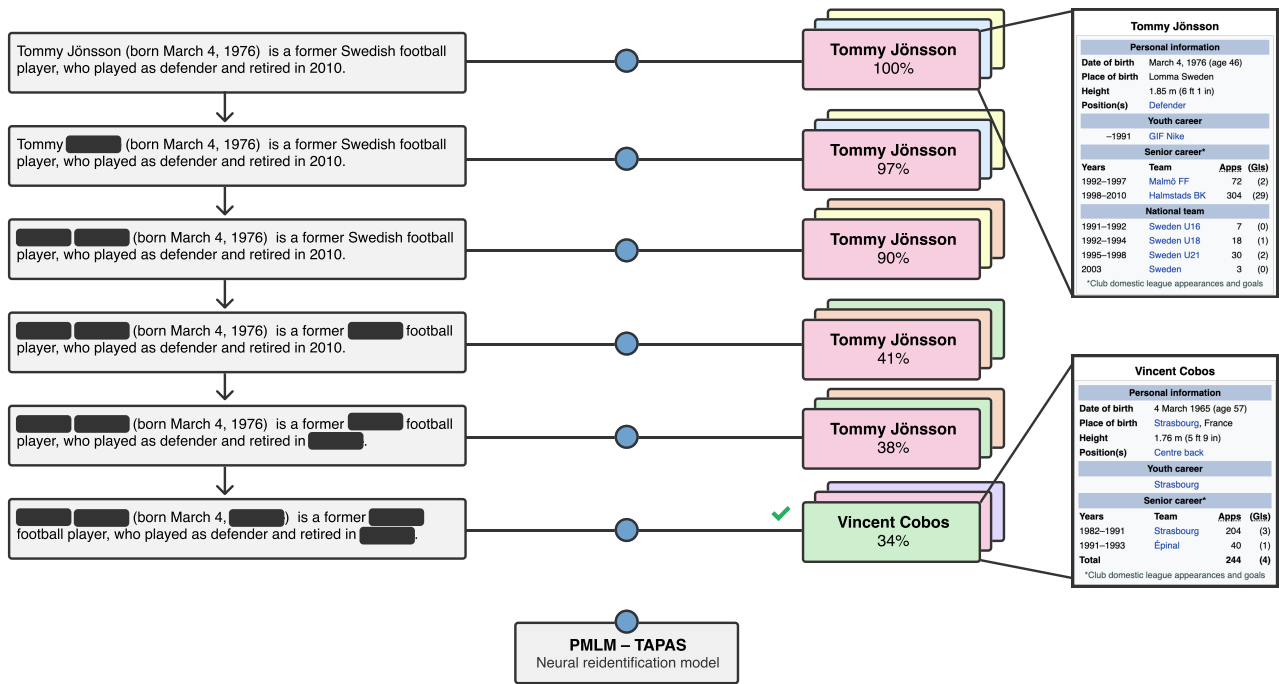


Figure 1: Method overview. A document (x , top-left) paired with a profile (\hat{y} , top-right) is given to the system. A trained neural reidentification model ($p(y|x, z)$, blue circle) produces a distribution over all possible profiles based on densely encoded representations. At each stage of inference, masks are added to the source document, changing the relative rank of the reidentification model. The method is run until k-anonymity of the reidentification model is achieved. Note that in this example, it is not necessary to remove all information, such as the month and day of birth, since the player is already deidentified.

with info-boxes (Lebret et al., 2016). The system is fit on a subset of the data and then asked to deidentify unseen individuals. Results show that even when all words from the profile are masked, the system is able to reidentify 32% of individuals. When we use our system to deidentify documents, it is able to fully anonymize them while retaining over 50% of words. When we compare our deidentification method to a set of unsupervised baselines, our method deidentifies documents more completely while removing fewer words. We qualitatively and quantitatively analyze the redactions produced by our system, including examples of successfully redacted quasi-identifiers.

2 Related Work

Automated deidentification. There is much prior work on deidentifying text datasets, both with rule-based systems (Neamatullah et al., 2008; Meystre et al., 2010; Sánchez et al., 2014; Norgeot et al., 2020; Sberbank and Emelyanov, 2021) and deep learning methods (Liu et al., 2017; Yue and Zhou, 2020; Johnson et al., 2020). Each of these methods is supervised, relies on datasets with human-labeled PII, and focuses on removing some

subset of the 18 identifying categories from HIPAA Safe Harbor. Other approaches include generating entire new fake datasets using Generative Adversarial Networks (GANs) (Chin-Cheong et al., 2019). Friedrich et al. (2019) train an LSTM on an EMR-based NLP task using adversarial loss to prevent the model from learning to reconstruct the input. Finally, differential privacy is a technique for ensuring provably private distributions (Dwork et al., 2006). It has mostly been used for training anonymized models on data containing PII, but requires access to the un-anonymized datasets for training (Li et al., 2021). Our deidentification approach does not provide the formal guarantees of differential privacy, but aims to provide a practical solution for anonymizing datasets in real-world scenarios.

Deidentification by reidentification. The NeurIPS 2020 Hide-and-Seek Privacy Challenge benchmarked both deidentification and reidentification techniques for clinical time series data (Jordon et al., 2021). In computer vision, researchers have proposed learning to mask faces in images to preserve the privacy of individuals using reidentification (Hukkelås et al., 2019;

Maximov et al., 2020; Gupta et al., 2021). In NLP, some work has been done on evaluating the reidentification risk of deidentified text (Scaiano et al., 2016). El Emam et al. (2009) proposes a method for deidentification of tabular datasets based on the concept of K-anonymity. Gardner and Xiong (2009) deidentify unstructured text by performing named entity extraction and redacting entities until k-anonymity. Mansour et al. (2021) propose an algorithm for deidentification of tabular datasets by quantifying reidentification risk using a metric related to K-anonymity. In our work, we train a reidentification model in an adversarial setting and use the model to deidentify documents directly.

Learning in the presence of masks. Various works have shown how to improve NLP models by masking some of the input during training. Chen and Ji (2020) show that learning in the presence of masks can improve classifier interpretability and accuracy. Li et al. (2016) train a model to search for the minimum subset of words required that, when removed, change the output of a classifier. They apply their method to neural network interpretability, and use reinforcement learning. Liao et al. (2020) pre-train a BERT-style language model to do masked-word prediction by sampling a masking ratio from $U(0, 1)$ and masking that many words. While their method was originally proposed for text generation, we apply the same masking approach to train language models for redaction.

3 Motivating Experiment: Quasi-Identifiers

In order to study the problem of deidentifying personal information from documents, we set up a model dataset utilizing personal profiles from Wikipedia. We use the Wikibio dataset (Lebret et al., 2016). Each entry in the dataset contains a *document*, the introductory text of the Wikipedia article, and a *profile*, the infobox of key-value pairs containing personal information. We train on the training dataset of 582, 659 documents and profiles. During test time, we evaluate only test documents, but consider all 728, 321 profiles from the concatenation of the train, validation, and test sets. This dataset represents a natural baseline by providing a range of factual profile information for a large collection of individuals, making it challenging to deidentify. In addition, it provides an openly available collection for comparing models.

ReID	DeID		
	None (0%)	Named entity (24%)	Lexical (28%)
IR ReID	74.9	4.3	0.0
NN ReID	99.6	79.7	31.9

Table 1: Percentage of documents reidentified (ReID) for different deidentification methods. Percentage of words masked in parentheses.

Is it difficult to deidentify individuals in this dataset? Wikipedia presents no domain challenges, and so finding entities is trivial. In addition many of the terms in the documents overlap directly with the terms in the profile table. Simple techniques should provide robust deidentification.

We test this with two deidentification techniques: (1) *Named entity* removes all words in documents that are tagged as named entities. (2) *Lexical* removes all words in the document that also overlap with the profile. To reidentify, we use an information retrieval model (BM25) and a dense neural network approach (described in Section 5).

Table 1 shows the results. While IR-based ReID is able to reidentify most of the original documents, without named entities or lexical matches, documents appear to be no longer reidentifiable. However, our model is able to reidentify 80% of documents, even with all entities removed. With all lexical matches with the profile removed (32% of total words), NN ReID is still able to reidentify a non-trivial number of documents.

This experiment indicates that even in the Wikibio domain, there are a significant number of pseudo-identifiers that allow the system to identify documents *even* when almost all known matching information is removed. In this work we study methods for discovering and quantifying these identifiers.

4 Deidentification by Inference

An overview of our data and system is shown in Figure 1. Given a document $x_1 \dots x_N$, we consider the problem of uniquely identifying the corresponding person y from a set of possible options \mathcal{Y} . The system works in the presence of redactions defined by a latent binary mask $z_1 \dots z_N$ on each position, where setting $z_n = 1$ masks word x_n .

We define a reidentification model as a model of $p(y | x, z)$ that gives a probability to each profile in

Algorithm 1 Greedy Deidentification

$x, \hat{y} \leftarrow$ input document and person
 $z_j \leftarrow 0$ for all j
for $i = 1$ to N **do**
 $j^* \leftarrow \arg \min_j p(y = \hat{y} \mid x, z_{-j}, z_j = 1)$
 $z_{j^*} \leftarrow 1$
 if $\hat{y} \notin K \arg \max_y p(y \mid x, z)$ **then**
 return z

\mathcal{Y} for a masked document. During deidentification, we assume that we have access to the true identity \hat{y} of the document that we would like to hide.

Our objective is to find the minimally sized mask that ensures that \hat{y} is not in the top- K predictions of the identification model:

$$\begin{aligned} \min_{z_1, \dots, z_N} \quad & |z| \\ \text{s.t.} \quad & \hat{y} \notin K \arg \max_y p(y \mid x, z). \end{aligned}$$

This objective is motivated by the concept of K -Anonymity (Samarati and Sweeney, 1998). A dataset has K -anonymity if each person \hat{y} in the dataset is indistinguishable from at least K other people in \mathcal{Y} .

The K -anonymity objective is combinatorial, and is intractable to solve with a non-trivial reidentification model. We instead approximate it with search. Specifically we use a simple greedy deidentification technique shown in Algorithm 1.

5 Reidentification Model

The core of this redaction system is a model of reidentification, $p(y \mid x, z)$. Defining this model faces two challenges: a) to facilitate informed search in the presence of masks and b) to correctly identify a person from 100,000s of choices.

As we do not have access to supervised masks, we define the probability of unmasked identification as marginalizing over all possible masks:

$$p(y \mid x) = \mathbb{E}_{z \sim p(z \mid x)} p(y \mid x, z; \theta)$$

where $p(z \mid x)$ is the mask prior and $p(y \mid x, z; \theta)$ is the reidentification model.

To assign a prior over masks $p(z \mid x)$, we opt for a simple setting that avoids building in additional information and fits well with deidentification search. One possibility would be to follow BERT-style masking and mask words at a fixed ratio of 15% (Devlin et al., 2019). However, Liao

et al. (2020) argue that while successful for classification, fixed-ratio masking works poorly for generation-style objectives. Following this advice, we use the following algorithm to construct masks of varying size:

- Sample the number of masks $l \sim \text{Uni}(0, N)$.
- Sample l masked words z_m by uniformly sampling indices m from $\{1, \dots, M\}$ without replacement.

For the reidentification model, $p(y \mid x, z; \theta)$, we follow the dense retrieval literature and use an embedding-based model (Karpukhin et al., 2020). Specifically we use an (asymmetric) bi-encoder model on documents and profiles. The document encoder f computes an embedding of the masked document, and the profile encoder $g(y)$ produces an embedding of the profile table corresponding to person y . We score the match by computing the joint encoding $f(x, z)^\top g(y)$ using the dot product between the vectors outputted by two neural networks. Define the matrix of profile embeddings as $\mathbf{G} = [g(y_1); \dots; g(y_{|\mathcal{Y}|})]$. The reidentification probability is defined as

$$p(y = i \mid x, z) = \text{softmax}(f(x, z)^\top \mathbf{G})_i.$$

During training we utilize label smoothing on the distribution, which has also been shown to be useful when training for inference in an argmax setting (Müller et al., 2019).

To train the model we optimize a lower bound on the identification log-likelihood:

$$\log p(y \mid x) \geq \mathbb{E}_{z \sim p(z \mid x)} [\log p(y \mid x, z)]$$

Specifically we sample a word dropout mask z for each element x from the prior, and then mask words during reidentification training.

Note that for training we compute the full distribution and do not use a contrastive approximation. In order to learn the parameters of g we utilize coordinate ascent. Specifically we fix \mathbf{G} and optimize the parameters of f . We then switch and optimize the profile encoder g on odd-numbered epochs to predict documents in \mathcal{X} (with no masking), and then recompute \mathbf{G} .

6 Experimental Setup

Models We call our main deidentification model *NN DeID*. We consider several different parameterization variants of the dual encoder. For the

Method	Privacy		Utility
	Ensemble ReID	% Masked	Info. Loss (%)
(No redaction)	99.6	0.0	0.0
Lexical redaction	31.9	32.1	20.8
Named entity redaction	79.7	27.3	27.3
< 25% Reidentifiable			
IDF	24.2	58.5	66.3
IDF (Table-Aware)	21.2	29.6	29.0
NN DeID	22.8	24.2	20.2
< 5% Reidentifiable			
IDF	3.8	71.4	78.6
IDF (Table-Aware)	5.0	67.3	70.4
NN DeID	4.4	35.9	29.5
< 1% Reidentifiable			
IDF	0.0	82.2	81.1
IDF (Table-Aware)	0.1	74.7	80.2
NN DeID	0.0	43.5	40.0

Table 2: Statistics comparing sets of 1000 documents redacted using different methods at various levels of identifiability. Reidentification rate measures the rate at at least one model in our neural-network ensemble can retrieve the correct profile for a redacted document. Information loss is measured as the percentage change in the size of the text when compressed.

document encoder ($f(x, z)$), we consider two different pretrained language models: RoBERTa-base version (Liu et al., 2019) (125M parameters) and PMLM (Liao et al., 2020) (125M parameters), a pretrained encoder specifically designed to support masking-style inference. For the table encoder ($g(z)$), we consider: RoBERTa-base (Liu et al., 2019) (125M parameters) with a simple linearized version of the profile, and TAPAS base (Herzig et al., 2020) (111 million parameters), a model designed to handle table input. We randomly compute masks online during training, so documents take a new randomly-redacted form on each epoch. All models are implemented in Hugging Face transformers library (Wolf et al., 2020). Each model is trained for sixty epochs, about two days on a single NVIDIA RTX A6000 GPU. More training details are available in A.

We experiment with all combinations for reidentification models, specified by document-profile encoders, RoBERTa-RoBERTa (*RR*), RoBERTa-TAPAS (*RT*), PMLM-RoBERTa (*PR*), PMLM-TAPAS (*PT*). The PT model is the default for NN DeID.

Baselines We consider several unsupervised redaction baselines based on lexical matches with

the table and word frequencies. *Lexical* removes all overlapping words that appear in the profile from the document. *IDF (Table-Aware)* masks all overlapping words that appear in the profile from the document, then masks in order of descending Inverse Document Frequency (IDF) (rarest word first) until a fixed threshold. We compute IDF based on the full corpus of documents and profiles from the train, validation, and test sets. *Named entity* removes all named entities from the document.²

Metrics A major challenge is how to evaluate text privacy in the presence of a strong reidentification models. As shown in Section 3, information retrieval metrics work well for lightly redacted documents, but fail under heavy masking. We ran preliminary experiments with human subjects, but found that even at seemingly low levels of masking, documents were nearly impossible for humans to reidentify.

Inspired by work on adversarial privacy such as the NeurIPS Hide-and-Seek challenge (Jordan

²We identify named entities using the `dslim/bert-base-NER-uncased` model available from Hugging Face. Named entities identified are personal names (PER), organization names (OR), location names (LOC), and miscellaneous names (MISC) (Tjong Kim Sang and De Meulder, 2003).




Figure 2: Pareto curves comparing deidentification approaches on privacy versus words masked. Lexical and Named Entity baselines are fixed values. NN DeID is computed with different K -anonymity values. IDF (table-aware) and IDF remove until a fixed IDF cutoff threshold.

et al., 2021), we adapt a metric that utilizes an ensemble of reidentification models \mathcal{R} as a benchmark. A masked document, x, z , is considered reidentified if any of the models can correctly select its profile, i.e. $\hat{y} = \arg \max_y p_r(y | x, z)$ for any model $r \in \mathcal{R}$. In order to diversify the ensemble we utilize different pretrained neural models as discussed above. We observe that each model can reidentify others with high accuracy indicating diversity features (more discussion in Section 8.3). We also include a word-matching based IR model in the ensemble, but find that it is not competitive at reidentification. Explicitly, the ensemble consists of the three variant parameterizations (RR , PR , RT) as well as the IR matching model.

As a metric of utility, we compute the average percentage of words masked, as well as the information loss percentage, computed as the ratio between the size of the original and redacted texts when compressed. For each method and baseline we sweep over mask sizes to compute a curve of reidentifiability and utility.

Inference We generate redactions from the reidentification models using greedy search to find the word to mask that causes the maximum decrease in the correct prediction. We use search implementations from the TextAttack library (Morris et al., 2020). Search takes in a stopping parameter K which indicates the rank cutoff of \hat{y} to stop search, $\hat{y} \notin K \arg \max_y p(y|x, z)$. We run with different values of K to sweep over levels of privacy, and

Model	Masked	
	0%	30%
Baseline	52.9	10.6
+ Word dropout	55.4	20.5
Dropout by IDF-weighting	48.6	20.2
+ Label smoothing ($\alpha = 0.1$)	56.3	10.8
+ Bigger emb. (768 \rightarrow 3072)	61.8	22.2
+ Table encoder optimization	98.1	14.0
+ Combined	96.4	38.3

Table 3: Ablation study. Effect of different factors on model ReID accuracy across data with different redaction strategies. Experiments are on RT parameterization and use 1/10 training data and number of profiles.

generate redactions with different masking rates. We ignore stopwords to speed up the search since they will rarely be identifiers.

7 Results

Table 2 presents results comparing unsupervised deidentification techniques on privacy and utility under the ensemble reidentification metric. As noted above, we see that neither Lexical nor Named Entity redaction provide sufficient privacy. NN DeID can provide better privacy while masking fewer words. Both NN DeID and IDF based approaches can reach stronger levels of privacy ($< 5\%$ reidentifiability), but at these levels IDF masks most of the remaining words. At full deidentification under the ensemble, NN DeID masks less than half of the words. When we consider an information loss measure of utility, NN DeID also performs much better than IDF-based deidentification.

Figure 2 expands on these results by showing the Pareto curves for privacy and utility across methods. Curves are obtained by varying the K value used in NN DeID and the threshold for IDF based deidentification. Curves show that in addition to achieving better utility at very low rates of identifiability, the method also achieves better utility than lexical matching, and a steeper privacy curve even at lower levels of redaction.

Model ablations Table 3 shows an ablation study of the components added to the model to improve accuracy. An alternative approach to this task is to finetune a pretrained model directly for the reidentification task (baseline). However, we found that out-of-the-box this model was neither effective as a




Figure 3: Rank comparison of the true document (\hat{y}) in two different parameterized models of $p(y | x, z)$ (RT and PT). Mask z comes from a deidentification ($K = 8$) on the PT model. While correlated, the two parameterizations produce very different rankings.

pure reidentification model nor as a model to guide search. We ablate each component added to NN DeID independently utilizing 1/10th of the training data and profiles, and compare both on the original documents and on documents with 30% of the words masked. Word dropout with the proposed sampling rate improves model accuracy particularly in the high mask regime. Interestingly weighting word dropout frequency using IDF hurts model accuracy in the full regime, and is not included in the final model. Increasing the dual encoder embedding sizes from 768 to 3072 and adding label smoothing both increase model accuracy. Finally, using coordinate ascent to optimize the profile encoder in addition to the document encoder has by far the largest impact on model accuracy. The combination of these approaches gives a deidentification model that is accurate across levels of masking.

8 Analysis

8.1 Quasi-Identifiers in Redacted Examples

Table 4 shows examples of redacted documents. While the most common redacted entities in deidentified examples are names, dates, and locations, we find notable examples of redacted quasi-identifiers:

- **Determiners.** Determiners can provide useful information in context. In the first example, the system removes “American” before musician, but also the word “an” which, in this

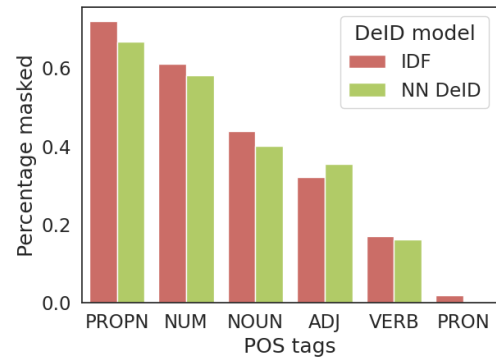


Figure 4: Percentage of words by part-of-speech tags that are masked by the IDF model and NN ReID model at $K = 8$ (similar masking level).

context, signals the next word may be “American”. This example is also interesting in that it preserves the word “Collective”, leading the model to predict a musician Avey Tare from the band “Animal Collective”.

- **Gender markers.** The model often redacts words marking gender in order to anonymize documents. In the second example, for the document on Madoko Hisagae, the model removes both the word “She” and “women’s”. This redaction leads to the prediction of Hiroki Ichigatani as the predicted match, a male Olympic fencer.
- **Locations.** The pretrained model seems to be able to identify relative locations even if they are not represented directly in the profile. In the third example, the profile indicates that Tim Tolkien is an English sculptor. The word “English” is masked immediately, but the location “Cradley Heath, West Midlands” is a quasi-identifier as to the country. Upon redacting this term, the model switches its prediction to Nesbert Mukomberanwa, a sculptor from Zimbabwe.

8.2 Redacted Word Types

The IDF (table-aware) model relies on overlapping words and rare words to redact content, whereas the NN DeID model can in theory remove any identifying word. Figure 4 compares the part-of-speech tags of the masked words between the two models at the same redaction level. We see that while similar, the NN DeID model masks fewer nouns, proper nouns and numbers, and more adjectives

tification models may be used as part of linkage attacks, where individuals can be pinpointed even from seemingly anonymized data. Additionally, the world knowledge of today’s large language models may be well-suited for this type of linkage attack. We observed this behavior empirically, when our models were uncannily able to reidentify individuals within a dataset of 720,000 identities, even from documents that appeared to have no remaining personal information.

We plan to release our models for deidentifying documents from Wikibio to the general public. We are open to hearing from users how our technology impacts both their lives and the lives of others, positively or negatively. If we receive any reports of misuse of our technology, we will mitigate accordingly.

12 Acknowledgments

AR and JC are supported by NSF CAREER 2037519, NSF 1704834, and a Sloan Fellowship. RZ is supported by a gift from the Simons foundation. JM is supported by Weill Cornell Medicine. Thanks to Dr. Curtis Cole and Dr. Thomas Campion from Weill Cornell Medicine for their general influence on our research direction within the area of deidentification.

References

Olivia Angiuli, Joe Blitzstein, and Jim Waldo. 2015. [How to de-identify your data](#). *Communications of the ACM*, 58(12):48–55.

Centers for Medicare & Medicaid Services. 1996. The Health Insurance Portability and Accountability Act of 1996 (HIPAA). Online at <http://www.cms.hhs.gov/hipaa/>.

Hanjie Chen and Yangfeng Ji. 2020. Learning variational word masks to improve the interpretability of neural text classifiers. *ArXiv*, abs/2010.00667.

Kieran Chin-Cheong, Thomas M. Sutter, and Julia E. Vogt. 2019. Generation of heterogeneous synthetic electronic health records using gans. In *NeurIPS 2019*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography*, pages 265–284, Berlin, Heidelberg. Springer Berlin Heidelberg.

Khaled El Emam, Fida Kamal Dankar, Romeo Issa, Elizabeth Jonker, Daniel Amyot, Elise Cogo, Jean-Pierre Corriveau, Mark Walker, Sadrul Chowdhury, Regis Vaillancourt, Tyson Roffey, and Jim Bottomley. 2009. [A globally optimal k-anonymity method for the de-identification of health data](#). *Journal of the American Medical Informatics Association: JAMIA*, 16(5):670–682.

William Falcon, The PyTorch Lightning team, et al. 2019. [PyTorch Lightning](#).

Max Friedrich, Arne Köhn, Gregor Wiedemann, and Chris Biemann. 2019. [Adversarial learning of privacy-preserving text representations for de-identification of medical records](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5829–5839, Florence, Italy. Association for Computational Linguistics.

James Gardner and Li Xiong. 2009. [An integrated framework for de-identifying unstructured medical data](#). *Data & Knowledge Engineering*, 68(12):1441–1451. Including Special Section: 21st IEEE International Symposium on Computer-Based Medical Systems (IEEE CBMS 2008) – Seven selected and extended papers on Biomedical Data Mining.

Aayush Gupta, Ayush Jaiswal, Yue Wu, Vivek Yadav, and Pradeep Natarajan. 2021. [Adversarial mask generation for preserving visual privacy](#). In *2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*, page 1–5.

Jonathan Herzig, Pawel Krzysztof Nowak, Thomas Müller, Francesco Piccinno, and Julian Martin Eisenchlos. 2020. Tapas: Weakly supervised table parsing via pre-training. *ArXiv*, abs/2004.02349.

Håkon Hukkelås, Rudolf Mester, and Frank Lindseth. 2019. [Deepprivacy: A generative adversarial network for face anonymization](#).

Abhik Jana and Chris Biemann. 2021. [An investigation towards differentially private sequence tagging in a federated framework](#). In *Proceedings of the Third Workshop on Privacy in Natural Language Processing*, pages 30–35, Online. Association for Computational Linguistics.

Alistair E. W. Johnson, Lucas Bulgarelli, and Tom J. Pollard. 2020. [Deidentification of free-text medical records using pre-trained bidirectional transformers](#). In *Proceedings of the ACM Conference on Health, Inference, and Learning*, page 214–221, Toronto Ontario Canada. ACM.

Alistair E. W. Johnson, Tom J. Pollard, Lu Shen, Liwei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. 2016. [Mimic-iii](#),

- a freely accessible critical care database. *Scientific Data*, 3(1):160035.
- James Jordon, Daniel Jarrett, Evgeny Saveliev, Jinsung Yoon, Paul Elbers, Patrick Thorat, Ari Ercole, Cheng Zhang, Danielle Belgrave, and Mihaela van der Schaar. 2021. [Hide-and-seek privacy challenge: Synthetic data generation vs. patient re-identification](#). In *Proceedings of the NeurIPS 2020 Competition and Demonstration Track*, volume 133 of *Proceedings of Machine Learning Research*, pages 206–215. PMLR.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). *CoRR*, abs/2004.04906.
- R. Le Bret, D. Grangier, and M. Auli. 2016. [Neural Text Generation from Structured Data with Application to the Biography Domain](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016. [Understanding neural networks through representation erasure](#).
- Xuechen Li, Florian Tramèr, Percy Liang, and Tatsunori Hashimoto. 2021. [Large language models can be strong differentially private learners](#).
- Yi Liao, Xin Jiang, and Qun Liu. 2020. [Probabilistically masked language model capable of autoregressive generation in arbitrary word order](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 263–274, Online. Association for Computational Linguistics.
- Pierre Lison, Ildikó Pilán, David Sanchez, 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, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *ArXiv*, abs/1907.11692.
- Zengjian Liu, Buzhou Tang, Xiaolong Wang, and Qingcai Chen. 2017. [De-identification of clinical notes via recurrent neural network and conditional random field](#). *Journal of Biomedical Informatics*, 75:S34–S42. Supplement: A Natural Language Processing Challenge for Clinical Records: Research Domains Criteria (RDoC) for Psychiatry.
- Huda O. Mansour, Maheyzah M. Siraj, Fuad A. Ghaleb, Faisal Saeed, Eman H. Alkhamash, and Mohd A. Maarof. 2021. [Quasi-identifier recognition algorithm for privacy preservation of cloud data based on risk reidentification](#). *Wireless Communications and Mobile Computing*, 2021:e7154705.
- Maxim Maximov, Ismail Elezi, and Laura Leal-Taixé. 2020. [CIAGAN: conditional identity anonymization generative adversarial networks](#). *CoRR*, abs/2005.09544.
- Stephane Meystre, F Friedlin, Brett South, Shuying Shen, and Matthew Samore. 2010. [Automatic de-identification of textual documents in the electronic health record: A review of recent research](#). *BMC medical research methodology*, 10:70.
- John Morris, Eli Lifland, Jin Yong Yoo, Jake Grigsby, Di Jin, and Yanjun Qi. 2020. [Textattack: A framework for adversarial attacks, data augmentation, and adversarial training in nlp](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 119–126.
- Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. 2019. [When does label smoothing help?](#) *Advances in neural information processing systems*, 32.
- Ishna Neamatullah, M. Douglass, Li wei H. Lehman, Andrew T. Reisner, Mauricio Villarroel, William J. Long, Peter Szolovits, George B. Moody, Roger G. Mark, and Gari D. Clifford. 2008. [Automated de-identification of free-text medical records](#). *BMC Medical Informatics and Decision Making*, 8:32 – 32.
- Beau Norgeot, Kathleen Muenzen, Thomas A. Peterson, Xuancheng Fan, Benjamin S. Glicksberg, Gundolf Schenk, Eugenia Rutenberg, Boris Oskotsky, Marina Sirota, Jinoos Yazdany, Gabriela Schmajuk, Dana Ludwig, Theodore Goldstein, and Atul J. Butte. 2020. [Protected health information filter \(philter\): accurately and securely de-identifying free-text clinical notes](#). *npj Digital Medicine*, 3(1):1–8.
- Mark Phillips and Bartha M. Knoppers. 2016. [The discombobulation of de-identification](#). *Nature Biotechnology*, 34(11):1102–1103.
- Ildikó Pilán, Pierre Lison, Lilja Øvrelid, Anthi Papadopoulou, David Sánchez, and Montserrat Batet. 2022. [The text anonymization benchmark \(tab\): A dedicated corpus and evaluation framework for text anonymization](#).
- Pierangela Samarati and Latanya Sweeney. 1998. [Protecting Privacy when Disclosing Information: k-Anonymity and its Enforcement through Generalization and Suppression](#). Technical report, SRI International.
- Russia Sberbank, Moscow and Yaroslav Emelyanov. 2021. [Towards task-agnostic privacy- and utility-preserving models](#). In *Proceedings of the Conference Recent Advances in Natural Language Processing - Deep Learning for Natural Language Processing Methods and Applications*, page 394–401. INCOMA Ltd. Shoumen, BULGARIA.

- Martin Scaiano, Grant Middleton, Luk Arbuttle, Varada Kolhatkar, Liam Peyton, Moira Dowling, Debbie S. Gipson, and Khaled El Emam. 2016. [A unified framework for evaluating the risk of re-identification of text de-identification tools](#). *Journal of Biomedical Informatics*, 63:174–183.
- David Sánchez, Montserrat Batet, and Alexandre Viejo. 2014. [Utility-preserving privacy protection of textual healthcare documents](#). *Journal of Biomedical Informatics*, 52:189–198. Special Section: Methods in Clinical Research Informatics.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Özlem Uzuner, Yuan Luo, and Peter Szolovits. 2007. [Evaluating the State-of-the-Art in Automatic De-identification](#). *Journal of the American Medical Informatics Association*, 14(5):550–563.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Ann Yuan, Daphne Ippolito, Vitaly Nikolaev, Chris Callison-Burch, Andy Coenen, and Sebastian Gehrmann. 2021. Synthbio: A case study in human-ai collaborative curation of text datasets. *ArXiv*, abs/2111.06467.
- Xiang Yue and Shuang Zhou. 2020. [PHICON: improving generalization of clinical text de-identification models via data augmentation](#). *CoRR*, abs/2010.05143.

A Training details

We train all models using the Adam optimizer with an initial learning rate of $5 * 10^{-5}$ or $1 * 10^{-4}$. We clip gradients to a maximum norm of 5.0. We decrease the learning rate by a factor of 0.5 whenever performance on a set of held-out redacted validation examples decreases. We train on the full dataset for a duration of 60 epochs, but stop early if the learning rate does not increase for 5 epochs. We implement training using the PyTorch lightning library (Falcon et al., 2019). We use PMLM-a, the version of PMLM with absolute positional embeddings (Liao et al., 2020). All encoders have a

maximum sequence length set to 128 throughout all experiments. We truncate tables by dropping columns until the encoded table fits the maximum sequence length.

We use linear learning rate scheduling. For models with the RoBERTa document encoder, we decay the learning rate from a starting point of $1e - 5$ to $1e - 6$ over the epochs. For models with PMLM as the document encoder, we start the learning rate at $5e - 5$. This is because we found empirically the PMLM tended to converge better when started at a lower learning rate. For the first two epochs (one for each encoder), we employ linear warmup, and increase the learning rate from 0 to its true initial value. We find that training the profile encoder is not useful after a handful of epochs, as the profile encoder starts to overfit, and our compute is better spent training the document encoder, which learns much more slowly since its inputs are 50% redacted on average. Thus, after the first 10 epochs (5 of which are spent training the profile encoder), we only train the document encoder.

B Search methods ablation

Our deidentification method redacts words by greedily selecting the word that minimizes the performance with respect to a reidentification model. We also tested using beam search to select words to redact, and found that it did not improve performance. At $k = 1$, beam search with beam width 4 masked 14.96% of words at 78.1% reidentification rate, while greedy masks 15.46% at a 78.9% reidentification rate, while being 3.39 times faster.