Natural Language Understanding with Privacy-Preserving BERT

Chen Qu, Weize Kong, Liu Yang, Mingyang Zhang, Michael Bendersky, Marc Najork
Google

Mountain View, CA, United States

{cqu, weize, yangliuy, mingyang, bemike, najork}@google.com

ABSTRACT

Privacy preservation remains a key challenge in data mining and Natural Language Understanding (NLU). Previous research shows that the input text or even text embeddings can leak private information. This concern motivates our research on effective privacy preservation approaches for pretrained Language Models (LMs). We investigate the privacy and utility implications of applying $d\chi$ -privacy, a variant of Local Differential Privacy, to BERT fine-tuning in NLU applications. More importantly, we further propose privacy-adaptive LM pretraining methods and show that our approach can boost the utility of BERT dramatically while retaining the same level of privacy protection. We also quantify the level of privacy preservation and provide guidance on privacy configuration. Our experiments and findings lay the groundwork for future explorations of privacy-preserving NLU with pretrained LMs.

CCS CONCEPTS

• **Security and privacy** → **Privacy protections**; • **Computing methodologies** → **Natural language processing**.

KEYWORDS

Local Privacy Constraints; Natural Language Understanding; Language Model Pretraining

ACM Reference Format:

Chen Qu, Weize Kong, Liu Yang, Mingyang Zhang, Michael Bendersky, Marc Najork. 2021. Natural Language Understanding with Privacy-Preserving BERT. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management (CIKM '21)*, November 1–5, 2021, Virtual Event, QLD, Australia. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3459637.3482281>

1 INTRODUCTION

The recent development of Deep Learning (DL) has led to notable success in Natural Language Understanding (NLU). Data-driven neural models are being applied to a rich variety of NLU applications, such as sentiment analysis [28], question answering [26], information retrieval [36], and text generation [6]. Many of these technologies have been deployed on the cloud by industrial service providers to process user data from personal customers, small businesses, and large enterprises. However, the rapid growth of NLU

technologies also comes with a series of privacy challenges due to the sensitive nature of user data. In NLU, the input text or even the text vector representations can leak private information or even identify specific authors [8, 16, 25, 29]. This lack of privacy guarantees may impede privacy-conscious users from releasing their data to service providers. Thus, service providers may suffer from the deficiency of genuine and evolving user data to train and evaluate NLU models. Besides, unintended data disclosure and other privacy breaches may result in litigation, fines, and reputation damages for service providers. These concerns necessitate our research on privacy-preserving NLU.

Specifically, we identify two challenges for privacy-preserving NLU. The first challenge is how to privatize users' text data in a Local Privacy setting, i.e., anonymize text to prevent leakage of private information. Prior work has applied Differential Privacy (DP) [10] and its variants to address similar privatization issues – originally for statistical databases [10] and more recently for DL [1, 27] and NLU [12, 18, 19]. However, in the context of NLU, many previous works mostly focus on a *Centralized Privacy* setting, which assumes a trusted centralized data aggregator to collect and process users' text data for training NLU models [23]. This solution, however, might not be sufficient for many users who are concerned with their sensitive or proprietary text data, when used for model training and serving. Thus, text privatization without a trusted data aggregator, also referred to as a *Local Privacy* setting, has become a pressing problem that remains less explored.

To tackle this challenge, we consider *Local Differential Privacy* (LDP) as the backbone of our privacy-preserving mechanism. In this setting, users perturb each individual data entry to provide plausible deniability [5] with respect to the original input *before* releasing it to the service providers. LDP also has advantages over federated learning [22] as discussed in Sec. 2. Specifically, we adopt a text privatization mechanism recently proposed by Feiyisetan et al. [12]. This mechanism is based on $d\chi$ -privacy (Sec. 3.2.1). It relaxes LDP to preserve more information from the input so that it is more practical for NLU applications [12].

The second challenge, which is also the focus of this work, is how to improve the utility of NLU models under Local Privacy settings, where the text is already privatized before model training and serving. Recent progress of pretrained Language Models (LMs) has led to great success in NLU. However, to the best of our knowledge, these research questions have not been well studied: (a) Can pretrained LMs adapt to privatized text input? (b) What is the most practical way to apply text privatization for pretrained LMs so that we retain the most utility (Fig. 1)? (c) Can we improve pretrained LMs to adapt to the privatized input via pretraining?

To answer these research questions, we first systematically discuss three privacy-constrained fine-tuning methods that apply text

Work done during Chen's internship at Google.

This work is licensed under a Creative Commons Attribution International 4.0 License.

CIKM '21, November 1–5, 2021, Virtual Event, Australia.

© 2021 Copyright is held by the owner/author(s).

ACM ISBN 978-1-4503-8446-9/21/11.

<https://doi.org/10.1145/3459637.3482281>

Figure 1: Illustrations of different privacy-constrained methods for training a typical NLU model.

privatization at different stages of the NLU model: sequence representations (Fig. 1.b), token representations (Fig. 1.c), and input text (Fig. 1.d). Along with previous research [12, 18], our work is another fundamental step to paint a more complete picture of text privatization for NLU. Lyu et al. [18] only looked at sequence representation privatization in BERT fine-tuning. This method does not address training-time privacy issues and incurs heavy computational costs for users (Sec. 2). In contrast, our work guarantees tunable privacy protection (Sec. 5) at both training and inference time by focusing on the token-level privatization (Fig. 1.c/d). Feyisetan et al. [12] only considered text-to-text privatization while we also study token representation privatization. Furthermore, we investigate how to improve NLU performance on top of the privacy mechanisms.

More importantly, to the best of our knowledge, our work is the first to study privacy-adaptive LM pretraining to improve the effectiveness and robustness of pretrained LMs on privatized text. We focus on BERT [9] since it is one of the most widely-used pretrained LMs. Our privacy-adaptive pretraining approaches are based on several variants of the Masked Language Model loss we designed, to leverage large-scale public text corpora for self-supervised learning in a privacy-adaptive manner, and to address the deficiency of labeled user data in privacy-constrained fine-tuning.

We conduct both privacy and utility experiments on two benchmark datasets. For the privacy experiments, we demonstrate and interpret the level of privacy protection by analyzing the plausible deniability statistics on the vocabulary level, as well as investigating the performance of token embedding inversion attack on actual corpora. Based on these results, we discuss and compare two principled approaches to guide the selection of the privacy parameter. We also reveal the geometry properties of the BERT embedding space to better understand the privatization process.

For the utility experiments, we investigate the performance of token-level privacy-constrained fine-tuning and privacy-adaptive pretraining methods. In the fine-tuning experiments, we discover that the text-to-text privatization method can often improve over token representation privatization thanks to the post-processing step of nearest neighbor search. More importantly, in pretraining experiments, we first show that BERT is able to adapt to privatization to some extent by being pretrained on fully privatized corpora. We further demonstrate that the integration of a denoising

heuristic can make BERT even more robust in handling privatized representations. Another exciting finding is that token representation privatization outperforms text-to-text privatization when noise is large with privacy-adaptive pretraining. In other words, the adaption resulting from privacy-adaptive pretraining can work better than nearest neighbor search in this scenario. These results show that our privacy-adaptive pretraining approaches can make BERT more effective and robust in handling privatized text input.

2 RELATED WORK

Differential Privacy. Differential Privacy (DP) [10] was originated from the field of statistical databases. It is one of the primary methods for defining privacy and preventing privacy breaches. At a high level, a randomized algorithm is differentially private if its output distribution is similar when the algorithm runs on two input datasets that differ in at most one data entry. Therefore, an observer seeing the output cannot tell if a particular data entry was used in the computation. Two settings are typically considered for DP: Centralized DP (CDP) and Local DP (LDP). CDP assumes a trusted data collector who can collect and access the raw user data. The randomized algorithm is applied on the collected dataset to produce differentially private output for downstream use. LDP [13] does not assume such a trusted data collector (e.g., users do not want service providers to access and collect their raw text messages). Instead, a randomized algorithm is applied on each individual data entry to provide plausible deniability [5] before sending it to the untrusted data collector (e.g., each user privatizes the text messages before uploading them to the service providers).

Privacy-Preserving Deep Learning. Another line of research aims to train privacy-preserving DL models [1, 23] by using differentially private stochastic gradient descent. The goal of those papers is to prevent the DL model from memorizing and leaking sensitive information in the training data. In contrast, some other work in this line aims to prevent an attacker from recovering information about the input text at inference time. For example, Coavoux et al. [8] and Li et al. [16] proposed to train deep models with adversarial learning, so that the model does not memorize unintended information. Both works provide only empirical improvements in privacy, without mathematically-sound privacy guarantees. Different from this work, PixelDP [14] applied DP to computer vision models so that the models are robust to adversarial examples. Another important line of work studied federated learning [22] to conduct decentralized training with the user data left on users' local/mobile devices. However, such a learning schema is hindered by low computational resources on users' device, and it has been argued to have privacy issues [20, 21]. Compared with federated learning, our privacy-preserving NLU methods do not have the bottleneck caused by the lack of computational resources on users' local devices. Also, our approaches enjoy the flexibility of owning a privatized version of user data and thus can exploit the data for training new models for the same or new NLU tasks.

Anonymization for NLU. We categorize anonymization methods used in NLU as text-to-text privatization and text representation privatization. Typical text-to-text privatization methods include de-identification and k -anonymization [3, 15]. The former redacts Personally Identifiable Information (PII) in the text while the latter

retains only words or n-grams used by a sufficiently large number (k) of users, without word/n-gram sequence information. The downside is that de-identification could easily leak other sensitive information [2] and k -anonymization has the same flaw when the adversary has background knowledge. Feyisetan et al. [12] recently proposed a text-to-text privatization method based on DP ($d\chi$ -privacy [10] to be exact). This method replaces the words in the original text with other words that are close in the embedding space in a local differential private manner. We also adopt this strategy since $d\chi$ -privacy provides mathematically provable privacy guarantees regardless of the background knowledge an adversary might use. Different from that work, we study the impact of privatization on BERT and propose privacy-adaptive pretraining methods to improve its utility while maintaining the same privacy guarantees.

On the other hand, text representation privatization aims to privatize the float-vector representations for text, including term-frequency vectors, word2vec embeddings [24], and representations from BERT [9]. Recently, Weggenmann and Kerschbaum [32] used a similar DP-based method as Feyisetan et al. [12], but with a different goal, i.e., to anonymize the term frequency vector representation of a document. Given the recent advance in representation learning with pretrained LMs, how to privatize representations from these models has become an increasingly important research problem. For example, Lyu et al. [19] followed the idea of Unary Encoding [31], and proposed a mechanism to anonymize text representation that provides ϵ -LDP. In addition, Bhowmick et al. [4] proposed a LDP mechanism to privatize high dimensional vectors, which is applied to the gradients for federated learning. This can be potentially applied to text representations. Also, Lyu et al. [18] looked at sequence representation privatization in BERT fine-tuning. A major drawback of their approach is that they did not address privacy issues at the training time since it requires service providers to have access to the raw user data to fine-tune the user-side encoder. This access violates the Local Privacy requirement at the training time. Their method also incurs heavy computational costs for users because the entire encoder is deployed on the user side. In comparison, our approach guarantees privacy protection at both training and inference time and incurs less computational cost for users.

3 OUR APPROACH

3.1 Overview

We present an overview of our approach for privacy-preserving NLU in Fig. 2. It contains two major stages. The first stage is the privacy-preserving mechanism (Sec. 3.2), where each user applies this mechanism to transform raw text to either privatized text or privatized text representations on their own local devices, and then submits the output to the service provider. The second stage is the privacy-adaptive NLU model training, where the service provider can only access the privatized text data for building NLU models. To improve the NLU performance on privatized text data, we consider various privacy-constrained fine-tuning methods (Sec. 3.3) and propose privacy-adaptive pretraining methods (Sec. 3.4). The latter further improves model utility by leveraging large-scale public text corpora that we privatized. We study the widely-used BERT to exemplify how our approach applies to pretrained LMs.

Figure 2: Overview. Users privatize input text locally. Then service providers conduct privacy-adaptive model training.

3.2 Privacy-Preserving Mechanism

We describe the privacy mechanism following Feyisetan et al. [12] and illustrate how it is applied to text privatization.

3.2.1 Preliminaries. We first briefly review two variants of DP related to our work – Local Differential Privacy (LDP) and $d\chi$ -privacy. Both variants are designed for the Local Privacy setting, where users need to privatize each data instance before releasing it to the untrusted data collector. Both variants employ a randomized mechanism $M : \mathcal{X} \rightarrow \mathcal{Y}$ that takes in a single data instance $x \in \mathcal{X}$ and outputs a randomized output $y \in \mathcal{Y}$.

LDP [13] requires M to satisfy, for any two inputs $x, x' \in \mathcal{X}$,

$$\frac{\Pr[M(x) = y]}{\Pr[M(x') = y]} \leq e^\epsilon, \forall y \in \mathcal{Y}, \quad (1)$$

where $\epsilon \geq 0$ is a privacy parameter. Intuitively, Eq. 1 suggests the output of $M(x)$ and $M(x')$ have very similar distributions such that an adversary cannot tell whether the input is x or x' . In other words, M provides plausible deniability [5] with respect to the original input. However, LDP is a very strong privacy standard. Regardless of how unrelated x and x' are, LDP requires them to have similar and indistinguishable output distributions. As a result, the output may not preserve enough information from the original input and thus may hurt the utility of downstream tasks.

$d\chi$ -privacy [7], a relaxation of LDP, was introduced to address the problem mentioned above. More formally, $d\chi$ -privacy requires, for any two inputs $x, x' \in \mathcal{X}$,

$$\frac{\Pr[M(x) = y]}{\Pr[M(x') = y]} \leq e^{\eta d(x, x')}, \forall y \in \mathcal{Y}, \quad (2)$$

where $d(x, x')$ is a distance/metric function (e.g., Euclidean distance) and $\eta \geq 0$ is the privacy parameter.¹ Compared with LDP, $d\chi$ -privacy allows the indistinguishability of the output distributions to be scaled by the distance between the respective inputs, i.e., ϵ in Eq. 1 becomes $\eta d(x, x')$ in Eq. 2. This allows M to produce more similar output for similar x and x' measured by $d(x, x')$, and thus could preserve more information from the input. That said, we should highlight that the semantics of η in $d\chi$ -privacy depends on the choice of d , and thus one needs to understand the structure of the underlying metric d in order to interpret the privacy consequences. This necessitates our study to measure and calibrate the level of privacy protection for our case in Sec. 5.3 and 5.4. Given

¹ Different from the prior work [12], we use η to denote the privacy parameter in $d\chi$ -privacy instead of ϵ , in order to avoid the confusion with ϵ used in LDP and DP.

the advantage of $d\chi$ -privacy and inspired by Feyisetan et al. [12], we adopt $d\chi$ -privacy to perform token representation privatization and text-to-text privatization in Sec. 3.2.3 and 3.2.4. We refer our readers to Feyisetan et al. [12] for the complete privacy proof.

3.2.2 Threat Model. Following Feyisetan et al. [12], we consider a threat model where each user submits a token to the service provider to conduct various downstream tasks. A user's token either appears in its clear form or in the form of a token representation. We expand users' input from a token to a text sequence in Sec. 3.2.5.

3.2.3 Token Representation Privatization. It is the *de facto* method for deep NLU models to represent input tokens with distributed dense vectors (e.g., word embeddings). Such token representations are typically produced by an embedding model. Without loss of generality, a token can be a character [38], a subword or wordpiece [34], a word, or an n-gram [39].

To prevent the leakage of sensitive information, we adopt $d\chi$ -privacy (Sec. 3.2.1) to privatize such token representations. In this case, the input to the randomized mechanism M becomes a token embedding $x \in \mathbb{R}^n$ and the output becomes $y \in \mathbb{R}^n$. For simplicity, we only consider the case where x, y are of the same dimension, n . Then $d\chi$ -privacy can be achieved for the choice of Euclidean distance $d(x, x') = \|x - x'\|$ by adding random noise N drawn from an n -dimensional distribution with density $p(N) \propto \exp(-\eta\|N\|)$ [10],

$$M_{\text{rep}}(x) = x + N, \quad (3)$$

where we use M_{rep} to denote the privacy mechanism for token representation privatization. This process is referred to as *perturbation* or *noise injection*. To sample N from the noise distribution, consider $N \in \mathbb{R}^n$ as a pair (r, p) , where r is the distance from the origin and p is a point in $\mathbb{B}^n$ (the unit hypersphere in $\mathbb{R}^n$). Then we sample $N \in \mathbb{R}^n$ by computing $N = rp$, where r is sampled from Gamma distribution $\Gamma(n, \frac{1}{\eta})$ and p is sampled uniformly over $\mathbb{B}^n$ [11, 33].

Although we describe M_{rep} from the token's perspective, the same procedure also applies to sequence representations (e.g., the [CLS] representation produced by BERT for a sentence).

3.2.4 Text-to-Text Privatization. In addition to token representation privatization, we also study text-to-text privatization. We consider a *token-to-token* case where each plain input token is transformed to a privatized output token. More formally, both the input and output of the randomized mechanism M become a token $x, y \in \mathcal{V}$, where $\mathcal{V}$ is a vocabulary set. The privacy proof in Feyisetan et al. [12] show that $d\chi$ -privacy can be achieved in this case by adding a post-processing step for M_{rep} (Eq. 3) to map the output of M_{rep} to another token via nearest neighbor search. More specifically, we first embed the input token x using an embedding model $\phi : \mathcal{V} \rightarrow \mathbb{R}^n$. We then pass $\phi(x)$ to M_{rep} to obtain the privatized token representation $M_{\text{rep}}(\phi(x))$. Lastly, we return the token that is closest to $M_{\text{rep}}(\phi(x))$ in the embedding space as the output,

$$M_{\text{txt}}(x) = \arg \min_{w \in \mathcal{V}} \|M_{\text{rep}}(\phi(x)) - \phi(w)\|, \quad (4)$$

where M_{txt} denotes the privacy-preserving mechanism for text-to-text privatization. M_{rep} and M_{txt} offer equivalent privacy protection [12] since the post-processing strategy does not affect privacy guarantees. This process of nearest neighbor search is usually fast and scalable [12] since the vocabulary can often fit in-memory and the operation can be optimized with ML accelerators.

Token representation privatization and text-to-text privatization each have their own merits. The former produces perturbed embeddings that could aid the neural models in exploiting the underlying semantics while the latter generates perturbed tokens that are easy to interpret for humans. Also, text-to-text privatization is more compatible with existing pipelines of text processing [12].

3.2.5 Sequence Input. Our discussion above considers the input x as a single token (or representation for a single token). When the input becomes a sequence of tokens, $x = (x_i)_1^\ell$, we apply M_{rep} or M_{txt} on each token or token representation to privatize the sequence x to $y = (y_i)_1^\ell$. In this case, the mechanism still satisfies $d\chi$ -privacy, but for the distance function $d(x, x') = \sum_1^\ell \|\phi(x_i) - \phi(x'_i)\|$. Further details and proofs can be found in Feyisetan et al. [12].

3.3 Privacy-Constrained Fine-Tuning

Due to the Local Privacy constraints, we do not access users' raw input text. However, we assume we have access to ground-truth labels of the NLU task since these labels can often be inferred from user behaviors, such as clicks [37] and other implicit feedback.

A neural NLU model typically consists of an embedding layer, an encoder, and task-specific layers (Fig. 1). The embedding layer converts the text input to a sequence of token embeddings, which will then go through the encoder to produce a sequence representation. Finally, task-specific layers make predictions based on the sequence representation. We split the NLU model to user side and service-provider side to comply with Local Privacy constraints and discuss three privacy-constrained training/fine-tuning methods:

- **Null privacy** (Fig. 1.a). We do not apply any privacy constraints and thus cannot provide any privacy protections. This is to provide an upper bound for model utility.
- **Sequence representation privatization** (Fig. 1.b). The embedding layer and the encoder are deployed user-side. The user perturbs the sequence representation locally.
- **Token representation privatization** (Fig. 1.c). Only the token embedding layer is deployed at the user side. The user conducts tokenization and embedding table look-up locally to map the input text to token embeddings. They then privatize the token embeddings (Sec. 3.2.3) and send them to the service provider. Then service providers assemble the input sequence to the encoder by adding other necessary embeddings (e.g., positional embeddings) and injecting special tokens (e.g., [CLS]).
- **Text-to-text privatization** (Fig. 1.d). The users conduct text-to-text privatization (Sec. 3.2.4) locally and send the privatized text to the service provider. Thus, service providers have a complete NLU model stack to process the privatized text.

Our foremost requirement is that the service provider only works with privatized input at both training and inference time, without any access to the raw user data. Thus, the service provider is not able to update the model parameters of user-side components. In contrast, the sequence representation privatization approach in Lyu et al. [18] requires the service provider to access the raw user data during training, which violates the Local Privacy requirement at the training time. This fundamental difference makes our results incomparable with theirs. Our pilot experiments indicate that sequence representation privatization yields undesirable utility

(Sec. 6.2) when we make the user-side encoder untrainable to comply with the Local Privacy requirement. Thus, we focus on the other two privacy-constrained fine-tuning methods in this paper.

In addition, we investigate the performance of two encoders, BERT and BiLSTM, to inspect the impact of different encoders under Local Privacy constraints. We mainly experiment with BERT since it is currently one of the most widely-used pretrained LMs. We also consider BiLSTM as a baseline encoder, which was used by Feyisetan et al. [12]. We use the same wordpiece embeddings for both encoders for fair comparisons.

3.4 Privacy-Adaptive BERT Pretraining

Inspired by the pretrained nature of BERT, we further propose privacy-adaptive pretraining methods to leverage a massive amount of unstructured texts that are publicly available. We also enjoy the flexibility of having access to the raw input in this case. These advantages of pretraining could make BERT more robust in handling privatized text or text representations. The pretrained model can also be used for different downstream tasks.

We initialize the model with the original BERT checkpoint and conduct further pretraining with the Next Sentence Prediction (NSP) loss [9] and several variants of the Masked LM (MLM) loss we design. To simulate the scenario in privacy-constrained fine-tuning, we now assume the role of the users to produce large-scale privatized input for pretraining. We use the BooksCorpus [40] and the English Wikipedia data following BERT [9].

As explained in Sec 3.3, user-side components cannot be updated by the service provider during privacy-constrained fine-tuning to comply with Local Privacy constraints. In the pretraining stage, we have the option to update user-side components since we use public dataset for pretraining. However, our pilot experiments indicate that, if user-side components are trainable, these components tend to generate representations that can be immune to perturbation. For example, the updated token embedding layer tends to produce token embeddings that are less prone to be perturbed to a different token with the same η . Although this is good from the perspective of utility, this tendency severely affects the level of privacy protection demonstrated in Sec. 5.3 and 5.4. This privacy-utility trade-off motivates our decision that we must stop the gradient from being back-propagated to user-side components during pretraining so that we can maintain the same level of privacy protection while working on improving model utility. This measure also forces the model to adapt to perturbation with the model components on the service-provider side, instead of that on the user side.

Since we have access to the raw input data during pretraining, we can take advantage of the MLM objective to train the BERT encoder to adapt to the perturbation process more effectively. We propose different privacy-adaptive pretraining methods based on different prediction targets of MLM as described below.

- **Vanilla MLM:** predicting the perturbed masked tokens. The most straightforward idea is to pretrain BERT on fully privatized corpora. This privatized LM could be more effective and robust in handling privatized content than the original LM. Since the privatization can be done on-the-fly during pretraining, the LM pretraining process benefits from seeing a diverse collection of perturbed text input. Formally, the vanilla MLM loss for a single

masked position is defined as follows:

$$L_{\text{MLM}}^{\text{Vanilla}} = - \sum_{w^* \in \mathcal{V}} \mathbb{1}\{w^* = \hat{w}\} \log \frac{\exp \logit(w^*)}{\sum_{w' \in \mathcal{V}} \exp \logit(w')} \quad (5)$$

where w^* is a candidate prediction of the privatized token and $\logit(w^*)$ is the logit for making such a prediction. $\hat{w}$ is the true privatized token. $\mathbb{1}\{\cdot\}$ is an indicator function.

- **Probability MLM (Prob MLM):** predicting a set of perturbed tokens for each masked position. $d\chi$ -privacy guarantees that, for any finite η , the distribution of the perturbed tokens has a full support on the whole vocabulary [12]. In other words, every token in the vocabulary has a non-zero probability being selected as the perturbed token for a given input token. Meanwhile, the distribution of the perturbed tokens from $d\chi$ -privacy remembers the semantics of the input token. Therefore, we perturb each masked token multiple times to obtain a set of perturbed tokens. The MLM losses coming from the perturbed tokens are weighted by their empirical frequencies. These empirical distributions of perturbed tokens could be beneficial for the LM to understand the injected noise, and thus, adapt to privatized content in a more efficient manner. Formally,

$$L_{\text{MLM}}^{\text{Prob}} = - \sum_{w^* \in \mathcal{V}} \frac{\text{count}(w^*, \hat{W})}{|\hat{W}|} \log \frac{\exp \logit(w^*)}{\sum_{w' \in \mathcal{V}} \exp \logit(w')} \quad (6)$$

where $\hat{W} = \{\hat{w}_i\}$ is a set of valid privatized tokens for this masked position and $|\hat{W}|$ denotes its size. We use $\text{count}(w^*, \hat{W})$ to denote the number of occurrence of a candidate prediction w^* in $\hat{W}$.

- **Denosing MLM:** predicting the original tokens. Another idea is to let the model predict the original masked tokens so that the LM learns to recover the original semantics of the masked token given privatized context. Formally,

$$L_{\text{MLM}}^{\text{Denosing}} = - \sum_{w^* \in \mathcal{V}} \mathbb{1}\{w^* = w\} \log \frac{\exp \logit(w^*)}{\sum_{w' \in \mathcal{V}} \exp \logit(w')} \quad (7)$$

where w is the true original token for this masked position.

After further pretraining with our privacy-adaptive approaches, we fine-tune the BERT model with privacy-constrained training on task datasets (Sec. 3.3) to evaluate the final model performance.

4 DATASETS

As listed below, we use two datasets from the GLUE benchmark [30] for our privacy and utility experiments. This benchmark was also used to evaluate BERT [9] and other popular NLU models, including XLNet [35], and RoBERTa [17].

- **Stanford Sentiment Treebank (SST)**² [28] is a single sentence classification task. The goal is to predict a sentiment label (positive or negative) for a sentence of movie review. This dataset contains 67k training sentences and 872 validation sentences.
- **Quora Question Pairs (QQP)**³ is a sentence-pair classification task. The goal is to determine whether a pair of questions are paraphrases or not. This dataset is larger than SST and has 363k sentence pairs for training and 40k sentence pairs for validation.

Our adoption of these public datasets contributes to the reproducibility of our paper. Our selection also covers both the single sentence classification task and sentence pair classification task.

² <https://nlp.stanford.edu/sentiment/index.html>

³ <https://www.quora.com/q/quoradata/First-Quora-Dataset-Release-Question-Pairs>

Figure 3: Geometry of the BERT embedding space. Dots are average Euclidean distances of an original embedding and its k nearest neighbors. Vertical lines are average Euclidean distances between original and perturbed embeddings.

Since the GLUE test sets are hidden, we train our models and baselines on the training set and report the results on the validation set. We use accuracy as the metric for utility experiments in Sec. 6.

5 PRIVACY EXPERIMENTS

We present a series of privacy experiments to demonstrate the level of privacy protection provided by the privacy mechanism on BERT, which has not been studied before to the best of our knowledge.

5.1 Geometry of the BERT Embedding Space

As mentioned in Sec. 3.2.1, the level of privacy protection from $d\chi$ -privacy also depends on the distance function d , which is defined as the Euclidean distance of tokens in the BERT embedding space for our token privatization mechanisms (Sec. 3.2.3 and Sec. 3.2.4). Thus, to understand the privacy protection and the noisy injection of these mechanisms, we analyze the geometry properties of the BERT embedding space in this section.

We first compute the Euclidean distance between an original token embedding and its privatized/perturbed token embedding, i.e., $\|x - M_{\text{rep}}(x)\| = \|N\|$ from Eq. 3. We average these distances for all tokens in the BERT vocabulary, and compute this average distance for a set of strategically chosen η values. The results are shown as the black vertical lines in Fig. 3 labeled by the corresponding η . As expected, we observe that the average Euclidean distance corresponds to the mean of the Gamma distribution (Sec. 3.2.3), i.e., $\frac{n}{\eta}$, where $n = 768$ is the dimension size of BERT token embeddings. As η becomes smaller, the average distance grows progressively larger, implying increasingly larger noise.

Next, we compute the Euclidean distances among the original token embeddings for a calibrated comparison. Specifically, we compute the distances between each token and its k -th nearest neighbor, for all tokens in the BERT vocabulary $\mathcal{V}$. We then compute the average k -th nearest neighbor distance, for each k in $[1, 2, 3, 4, 5, 10, 20, 50, 100, 200, 500, 1000, 5000, 10000, |\mathcal{V}|]$. The results are presented as blue dots in Fig. 3. As k becomes larger, the average Euclidean distance also grows larger, as expected.

Finally, we compare the distances computed in the previous steps. Distances among the original token embeddings are smaller than 2 on average while even the smallest noise we used in experiments ($\eta = 175$) perturbs tokens to positions that are as twice as far on average. This gives us a sense of how large these perturbations are. Sec. 5.4 justifies using such large noise by showing nearest neighbor search in the embedding space (Sec. 3.2.4) can still map the perturbed token embedding to its original token. Our choices

Table 1: Examples of perturbed sentences with different choices of η . “Orig” denotes the original input sentence. The red color denotes tokens that are modified.

η	Sequence								
Orig	the	emotions	are	raw	and	will	strike	a	nerve
50	##ori	backward	og	wanda	big	disposal	##pose	lou	##bular
75	410	truth	go	mole	gone	will	strike	y	gifford
100	fine	abused	are	primitive	it	will	slaughter	us	nerve
125	measure	emotions	:	shield	and	relation	strike	nearly	nerve
150	the	caleb	are	kill	and	will	strike	circle	nerve
175	the	emotions	are	raw	and	will	strike	a	nerve

Table 2: Examples of perturbed tokens ($\eta = 100$, sampled independently and sorted by empirical frequencies).

	the	emotions	are	raw	and	will	strike	a	nerve
Perturbed tokens	the	emotions	are	raw	and	will	strike	a	nerve
	a	emotion	were	smackdown	or	would	strikes	the	rebels
	its	emotionally	is	matt	but	can	attack	an	reason
	and	hormones	being	##awa	,	may	drop	his	cells
	his	organizations	re	unused	-	better	##gen	its	spirits
	her	emotional	have	division	as	must	aim	her	bothering
	some	moods	of	protection	"	self	stroke	one	communications

of η are greater than the typically used ϵ in LDP. Given the relatively high dimensionality of BERT embeddings (768) and the large noise we require, we adopt higher η to achieve the level of privacy protection shown in the next sections. To bring down the level of η to values resembling ϵ in LDP, one may reduce the dimensionality of the BERT embeddings by random projection [33], which will be investigated in our future work.

5.2 Examples of Perturbed Text

We present examples of text-to-text privatization (Sec. 3.2.4) in Tab. 1 with different choices of η . As expected, we observe that as η becomes larger, the noise becomes smaller and more tokens remain unmodified. Next, we focus on $\eta = 100$ (as an example), and perturb each token in the previous sentence 1,000 times. We sort the top-10 perturbed tokens for each original token by their empirical frequencies. The results are presented in Tab. 2. We observe that a token has a certain probability to be preserved. For example, “emotions” has 14% of the chances to remain unmodified under this specific choice of η . A token can also be perturbed into a diverse set of other tokens. For instance, “emotions” is perturbed to 825 distinct words with different empirical frequencies. Some of the top words are “emotion”, “emotionally”, “hormones”, “emotional”, and “moods”, which share similar semantics with the original word “emotions”. Some perturbed tokens might not have obvious links to the original token due to the randomness introduced by the privacy-preserving mechanism. Intuitively, good privacy protection implies a token has a relatively small chance of being preserved, and can be modified to a relatively diverse set of perturbed tokens. We quantify these measurements in the next section.

5.3 Plausible Deniability Statistics

We follow Feyisetan et al. [12] to use two statistics to characterize the ability of an adversary to recover the original input text when observing the perturbed text or text representations. This ability is

referred to as plausible deniability [5] and it varies under different settings of η . The formal definitions of the statistics can be found in Feyisetan et al. [12]. We provide intuitive explanations as follows:

- N_w : the probability of an input token w not modified by $M_{\text{txt}}(w)$.
- S_w : the effective support of the output distribution of perturbation $M_{\text{txt}}(w)$ on the entire vocabulary for an input token w .

We run simulations to estimate the plausible deniability statistics. For each choice of privacy parameter η , we perturb each regular token in the vocabulary for 1,000 times. Regular tokens refer to the tokens other than [PAD], [CLS], [SEP], [MASK], [UNK], or [unused . . .]. We estimate N_w as the number of output tokens that are identical to the input token, and S_w as the number of unique output tokens. Intuitively, good privacy guarantees should be characterized by relatively small N_w and relatively large S_w .

We present the estimated plausible deniability statistics of the BERT vocabulary in Fig. 4. For example, when $\eta = 75$, the N_w figure shows no token is ever returned more than 500 times in the worst case and the S_w figure shows no token produces fewer than 500 distinct new tokens. As η becomes greater (the noise becomes smaller), a growing number of tokens tend to have larger N_w , indicating that the tokens tend to have greater probability of remaining unmodified. Meanwhile, more and more tokens tend to produce only a small amount of distinct perturbed outputs (S_w), suggesting a limited support on the vocabulary.

These statistic figures serve as a visual guidance for selecting η for different applications that require different privacy guarantees. Feyisetan et al. [12] suggest selecting η based on the desired worst case guarantees. For example, we should select η as 75 for an application that requires all tokens to be modified for at least 500 times out of 1,000 perturbations. In addition to the worst case guarantees, we show another practical approach to conduct η selection based on the average case guarantees in the next section.

5.4 Token Embedding Inversion

The plausible deniability statistics characterize the privacy guarantees from the perspective of the vocabulary. In this section, we show another interpretation of privacy protection level by measuring the amount of tokens that are modified by perturbation in actual datasets following Song and Raghunathan [29]. This analysis accounts for the token frequencies in real datasets and could be more practical in guiding the selection of the privacy parameter η .

We define an adversarial task, token embedding inversion, as recovering the original tokens based on perturbed token embeddings. We leverage nearest neighbor search for this task. Given a perturbed token embedding, we find the nearest neighbor of this embedding in the embedding space as the predicted original token. The performance is measured by accuracy. Although this attack is not comprehensive, its performance can be considered as an automatic metric to reveal the privacy-utility trade-off and to guide the selection of the privacy parameter η . Besides, this approach is shown to be highly effective on our low-level lexical input (tokens). We will leave more advanced ML attacks to our future work.

The results of token embedding inversion on the validation data of SST and QQP are presented in Tab. 3. We demonstrate that privacy protection levels are consistent across datasets, indicating that the privacy-preserving mechanism could be agnostic to specific

Table 3: Accuracy of token embedding inversion.

η	50	75	100	125	150	175
SST	0.0154	0.1084	0.3402	0.6354	0.8500	0.9511
QQP	0.0165	0.1066	0.3420	0.6462	0.8620	0.9570

tasks and datasets. As η becomes larger and the noise becomes smaller, a growing amount of tokens in actual datasets can be correctly recovered by nearest neighbor search. For example, at $\eta = 100$, only 34% of tokens in the datasets can be recovered, indicating that this choice of η provides moderately strong privacy protection.

We also observe that the inversion scores roughly correspond to the peak values of N_w in plausible deniability statistics (Fig. 4). This observation advocates for using average case guarantees for η selection. η selection with worst case guarantees always favors stricter privacy choices since this approach is contingent on the tokens that are least prone to be modified. η selection based on average case guarantees is an alternative approach that accounts for the term frequencies in actual corpora. Our recommendation is that, given the different linguistic properties in different languages, domains, applications, and data, the decision on the choice of η and the approach to select an η should be made on a case-by-case basis.

6 UTILITY EXPERIMENTS

6.1 Implementation Details

Privacy-Constrained Fine-Tuning. We freeze parameters for a model component to simulate the user-side deployment for this component. E.g., in token representation privatization, the embedding layer is frozen since it is deployed at the user side (Fig. 1). For both BERT and BiLSTM, we set the maximum sequence length to 128, the training batch size to 32, and the number of training epochs to 3. On SST, we use the learning rate of $2e-5$ to fine-tune BERT and $1e-3$ to train BiLSTM. On QQP, we use the learning rate of $4e-5$ to fine-tune BERT and $1e-3$ to train BiLSTM. The warm up portion of the learning rate is 10% of the total steps. We use the same wordpiece embeddings for both encoders for fair comparisons.

Privacy-Adaptive BERT Pretraining. We set the maximum sequence length to 128, the training batch size to 256, the learning rate to $2e-5$, the maximum number of predictions per sequence to 20, the mask rate to 0.15, the maximum training steps to 1,000,000, and warm up steps to 10,000. We save checkpoints every 200,000 steps and fine-tune the checkpoint with privacy-constrained training to select the best checkpoint. The choices of η for pretraining and fine-tuning are identical in our experiments. For Prob MLM, we set the number of perturbations per masked position to 10.

6.2 Results of Privacy-Constrained Fine-Tuning

Our experiments on sequence representation privatization indicates even the smallest perturbation we used ($\eta=175$) causes 30% absolute performance decrease compared with Null Privacy. Given this observation, we focus on experiments for token-level privatization.⁴

The utility results for token-level privatization are reported in Tab. 4 for SST and Tab. 5 for QQP. We use the same wordpiece embeddings for both BERT and BiLSTM encoders for fair comparisons.

⁴ We should note that the privacy semantics of η for sequence-level privatization and token-level privatization are different. We leave this to our future work.

Figure 4: Plausible deniability statistics. Each token in the vocabulary is perturbed 1,000 times. N_w refers to the number of output tokens that are identical to the input token, and S_w refers to the number of unique output tokens.

The privacy parameters η for token-level privatization, including those with privacy-adaptive pretraining, are directly comparable since they use equivalent inputs and distance functions.

We first compare the two token-level privatization methods (column 1 vs. 2, column 3 vs. 4 in Tab. 4 and 5). We see text-to-text privatization produces considerable performance improvement compared to token representation privatization. This improvement has statistical significance and is observed for both encoders on a wide range of η . The performance gain is larger when noise is smaller. This verifies our finding in Sec. 5.4 that nearest neighbor search is effective in mapping perturbed embeddings to its original token (although this does not affect the privacy guarantees [12]). Feyisetan et al. [12] take advantage of this technique to make the privatization mechanism compatible with existing pipelines of text processing, while we reveal the important role it plays to dramatically improve the performance of NLU models in privacy-constrained fine-tuning.

We then look at the comparisons of the two encoders. BERT performs better than BiLSTM if we do not apply privacy constraints, as expected. We observe that BERT suffers from a larger performance degradation than BiLSTM with token representation privatization. This indicates that BERT is less robust than BiLSTM in handling perturbation on representations. This unexpected result can be explained by the mismatch of representations observed by BERT during pretraining (on plain corpora) and fine-tuning (on privatized corpora). Also, complex models are more likely to suffer from inherently high variance. In the next section, we show our privacy-adaptive pretraining methods can improve BERT’s performance and help it outperform BiLSTM under the same privacy guarantee. On the other hand, BERT performs better with text-to-text privatization, compared with BiLSTM, in general. Thanks to the advantages brought by the LM pretraining and attention mechanisms, BERT handles the perturbation better than BiLSTM with statistical significance for a wide range of η values. This shows BERT is more robust than BiLSTM in handling text-to-text privatization.

Finally, we briefly look at the comparison of our approach with Null Privacy. Perturbation on token representations results in severe degradation in model performance compared with Null Privacy. Even though the encoder is fine-tuned in token representation privatization, it is not able to adapt to the perturbed token embeddings. This is within our expectation since we observe in Sec. 5.1 that the noise applied to embeddings causes a significant shift of the original

Table 4: Accuracy of privacy-constrained fine-tuning on SST. Scores in the header are obtained with Null Privacy. Bold-face denotes better results with the same encoder. Under-scores denote better results between different encoders with the same privacy constraint. Δi denotes the improvement with respect to column i has statistically significance with $p < 0.05$ tested by the Student’s paired t-test.

η	BERT (0.9289)		BiLSTM (0.8406)	
	1. Token Rep	2. Text-to-Text	3. Token Rep	4. Text-to-Text
50	0.5126	0.4920	0.5092	0.5092
75	0.5310	0.5906 $\Delta^{1,4}$	0.5447	0.5356
100	0.5298	0.7030 $\Delta^{1,4}$	<u>0.5608</u>	0.6697 Δ^3
125	0.5608	0.8360 $\Delta^{1,4}$	0.5608	0.7420 Δ^3
150	0.5665	0.8968 $\Delta^{1,4}$	0.5917	0.8119 Δ^3
175	0.5975	0.9151 $\Delta^{1,4}$	0.6216	0.8211 Δ^3

Table 5: Accuracy of privacy-constrained fine-tuning on QQP. Refer to Tab. 4 to interpret the notations.

η	BERT (0.9106)		BiLSTM (0.8261)	
	1. Token Rep	2. Text-to-Text	3. Token Rep	4. Text-to-Text
50	0.6370 Δ^2	0.6318	<u>0.6409</u> Δ^1	0.6423 Δ^2
75	0.6318	0.6485 Δ^1	0.6318	0.6631 $\Delta^{2,3}$
100	0.6318	0.7238 $\Delta^{1,4}$	0.6318	0.7108 Δ^3
125	0.6318	0.8274 $\Delta^{1,4}$	0.6318	0.7534 Δ^3
150	0.6447	0.8759 $\Delta^{1,4}$	<u>0.6811</u> Δ^1	0.7854 Δ^3
175	0.6354	0.8976 $\Delta^{1,4}$	<u>0.6987</u> Δ^1	0.8066 Δ^3

embedding in the embedding space. Text-to-text privatization, on the other hand, shows less degradation.

6.3 Results of Privacy-Adaptive Pretraining

Privacy-adaptive pretraining results are shown in Tab. 6 for token representation privatization and Tab. 7 for text-to-text privatization.

We first analyze the general effect of privacy-adaptive pretraining. For token representation privatization, consistent results on two datasets demonstrate that all three privacy-adaptive pretraining methods have significant performance improvement with statistical significance. This suggests that the BERT pretrained with privacy-adaptive methods are much more effective and robust than the original BERT model in handling privatized token representations. For text-to-text privatization, although it has inherently advantages due to nearest neighbor search, privacy-adaptive pretraining still manages to outperform the original BERT checkpoint on both

Table 6: Accuracy of privacy-adaptive pretraining with token representation privatization. Italic denotes better results than the original BERT. Underscores denote better results than the pretraining baseline of Vanilla MLM. Boldface denotes the best results. Δi denotes the gain with respect to column i has statistically significance ($p < 0.05$ tested by the Student’s paired t-test).

η	SST				QQP			
	1. Orig BERT	2. Vanilla MLM	3. Prob MLM	4. Denoising MLM	1. Orig BERT	2. Vanilla MLM	3. Prob MLM	4. Denoising MLM
50	0.5126	<i>0.5390</i> Δ^1	<i>0.5424</i> Δ^1	0.5356	0.6370	<i>0.6434</i> $\Delta^{1,3}$	0.6364	<i>0.6444</i> $\Delta^{1,3}$
75	0.5310	<i>0.5791</i> Δ^1	<i>0.5757</i> Δ^1	<i>0.6089</i> Δ^1	0.6318	<i>0.6645</i> Δ^1	<i>0.6821</i> $\Delta^{1,2}$	<i>0.7119</i> $\Delta^{1,2,3}$
100	0.5298	<i>0.6709</i> Δ^1	<i>0.6766</i> Δ^1	<i>0.7041</i> $\Delta^{1,2}$	0.6318	<i>0.7668</i> Δ^1	<i>0.7686</i> Δ^1	<i>0.7788</i> $\Delta^{1,2,3}$
125	0.5608	<i>0.7706</i> Δ^1	<i>0.7718</i> Δ^1	<i>0.7810</i> Δ^1	0.6318	<i>0.8200</i> Δ^1	<i>0.8219</i> Δ^1	<i>0.8249</i> $\Delta^{1,2,3}$
150	0.5665	<i>0.8188</i> Δ^1	<i>0.8188</i> Δ^1	<i>0.8395</i> $\Delta^{1,2,3}$	0.6447	<i>0.8520</i> Δ^1	<i>0.8520</i> Δ^1	<i>0.8523</i> Δ^1
175	0.5975	<i>0.8658</i> Δ^1	<i>0.8693</i> Δ^1	<i>0.8693</i> Δ^1	0.6354	<i>0.8698</i> Δ^1	<i>0.8688</i> Δ^1	<i>0.8691</i> Δ^1

Table 7: Accuracy of privacy-adaptive pretraining with token-to-token privatization. Refer to Tab. 6 to interpret the notations.

η	SST				QQP			
	1. Orig BERT	2. Vanilla MLM	3. Prob MLM	4. Denoising MLM	1. Orig BERT	2. Vanilla MLM	3. Prob MLM	4. Denoising MLM
50	0.4920	<i>0.5218</i>	<i>0.5310</i>	<i>0.5092</i>	0.6318	<i>0.6375</i> Δ^1	<i>0.6419</i> $\Delta^{1,2}$	<i>0.6469</i> $\Delta^{1,2,3}$
75	0.5906	<i>0.5963</i>	0.5734	0.5906	0.6485	<i>0.6643</i> Δ^1	<i>0.6616</i> Δ^1	<i>0.6693</i> $\Delta^{1,2,3}$
100	0.7030	<i>0.7190</i>	<i>0.7259</i>	<i>0.7225</i>	0.7238	<i>0.7628</i> Δ^1	<i>0.7610</i> Δ^1	<i>0.7603</i> Δ^1
125	0.8360	<i>0.8429</i>	<i>0.8429</i>	<i>0.8406</i>	0.8274	<i>0.8346</i> Δ^1	<i>0.8362</i> Δ^1	<i>0.8351</i> Δ^1
150	0.8968	<i>0.9048</i>	<i>0.9071</i>	<i>0.9025</i>	0.8759	<i>0.8790</i> Δ^1	<i>0.8801</i> Δ^1	<i>0.8791</i> Δ^1
175	0.9151	<i>0.9209</i>	<i>0.9209</i>	<i>0.9220</i>	0.8976	<i>0.8997</i>	<i>0.8994</i>	<i>0.8990</i>

datasets, and with statistical significance on QQP. These results demonstrate the effectiveness of our privacy-adaptive pretraining.

We then compare the performance of the three privacy-adaptive pretraining approaches. For token representation privatization, the Vanilla MLM is already highly effective, despite its simplicity. It only brings a marginal improvement if we enable the model to predict a set of perturbed tokens for each masked position with Prob MLM. This indicates that the augmentation effect in Vanilla MLM could be sufficient for the model to observe enough variations of the noise injection process. Predicting the original tokens with Denoising MLM shows visible gains compared with predicting perturbed tokens in Vanilla and Prob MLM. These improvements are particularly compelling and have statistical significance when we have moderately large noise ($\eta = 75, 100$), where more than half of the tokens in the datasets are modified. These results indicate Denoising MLM has more practical values in real-world applications.

For text-to-text privatization, Denoising MLM remains a competitive privacy-adaptive pretraining methods. In contrast to token representation privatization, predicting (a set of) perturbed tokens here shows marginal improvement over predicting the original token in some cases. Since text-to-text privatization sometimes lands on the original token as the perturbed token due to nearest neighbor search, Denoising MLM is less beneficial than the case with token representation privatization. In particular, privacy-adaptive pretraining with text-to-text privatization demonstrates considerable improvement when $\eta = 100$, where the noise is sufficiently large and more than half of the tokens in the datasets are modified. This improvement expands to a wider range of η values ($\eta = 50, 75, 100, 125$) on QQP. Privacy-adaptive pretraining has a pronounced effect on QQP, probably because the sentence pair prediction task is more difficult, and thus, is more dependent on the adaptation of perturbation in our privacy-adaptive pretraining process. This demonstrates the value of our methods in practical applications.

Finally, we compare the performance of privacy-adaptive pretraining between the two different token-level privatization. It is

worth noting that token representation privatization generally outperforms text-to-text privatization when noise is large ($\eta = 50, 75$, and 100 in some cases). This exciting observation indicates that the noise adaptation effect resulting from privacy-adaptive pretraining has advantages over nearest neighbor search when noise is large.

To sum up, we recommend adopting Denoising MLM as the primary method for privacy-adaptive pretraining for both token-level privatization approaches. When a relatively strong level of privacy protection (e.g., $\eta < 100$) is required, token representation privatization should be adopted to preserve more utility. Otherwise, text-to-text privatization is preferred. In both cases, privacy-adaptive pretraining is essential to improve model performance.

7 CONCLUSIONS AND FUTURE WORK

In this work, we study how to improve NLU model performance on privatized text in the Local Privacy setting. We first take a deep analytical view to illustrate the privacy guarantees of $d\chi$ privacy. We then investigate the behavior of BERT when it meets privatized text input with privacy-constrained fine-tuning methods. We show BERT is less robust than BiLSTM in handling privatized token representations. We further show text-to-text privatization can often improve upon token representation privatization, revealing the important role played by nearest neighbor search to improve utility in privacy-constrained fine-tuning. More importantly, we propose privacy-adaptive LM pretraining methods and demonstrate that a BERT pretrained with our Denoising MLM objective is more robust in handling privatized content compared with the original BERT. For future work, we would like to study privatization on contextualized token representations and more advanced ML attacks. We will also look at reducing the dimensionality of the BERT embeddings by random projection [33] to bring down the level of η .

ACKNOWLEDGMENTS

The authors would like to thank Borja Balle for the helpful discussions and constructive comments on this work.

REFERENCES

- [1] M. Abadi, Andy Chu, Ian J. Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and L. Zhang. 2016. Deep Learning with Differential Privacy. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (2016).
- [2] B. Anandan, C. Clifton, Wenxin Jiang, M. Murugesan, Pedro Pastrana-Camacho, and L. Si. 2012. t-Plausibility: Generalizing Words to Desensitize Text. *Trans. Data Priv.* 5 (2012), 505–534.
- [3] Michael Bendersky, X. Wang, Donald Metzler, and Marc Najork. 2017. Learning from User Interactions in Personal Search via Attribute Parameterization. In *Proceedings of the 10th ACM International Conference on Web Search and Data Mining*. 791–799.
- [4] Abhishek Bhowmick, John C. Duchi, J. Freudiger, G. Kapoor, and Ryan Rogers. 2018. Protection Against Reconstruction and Its Applications in Private Federated Learning. *ArXiv abs/1812.00984* (2018).
- [5] Vincent Bindshaedler, R. Shokri, and Carl A. Gunter. 2017. Plausible Deniability for Privacy-Preserving Data Synthesis. *ArXiv abs/1708.07975* (2017).
- [6] T. Brown, B. Mann, Nick Ryder, Melanie Subbiah, J. Kaplan, P. Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, G. Krüger, Tom Henighan, R. Child, Aditya Ramesh, D. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, E. Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, J. Clark, Christopher Berner, Sam McCandlish, A. Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. *ArXiv abs/2005.14165* (2020).
- [7] K. Chatzikokolakis, M. Andrés, N. E. Bordenabe, and C. Palamidessi. 2013. Broadening the Scope of Differential Privacy Using Metrics. In *Proceedings of the 13th International Symposium on Privacy Enhancing Technologies*. 82–102.
- [8] Maximin Coavoux, Shashi Narayan, and Shay B. Cohen. 2018. Privacy-preserving Neural Representations of Text. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 1–10.
- [9] 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*. 4171–4186.
- [10] C. Dwork, F. McSherry, Kobbi Nissim, and A. D. Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Proceedings of the 3rd Theory of Cryptography Conference*. 265–284.
- [11] Natasha Fernandes, M. Dras, and A. McIver. 2019. Generalised Differential Privacy for Text Document Processing. In *Proceedings of the 8th International Conference on Principles of Security and Trust*. 123–148.
- [12] Oluwaseyi Feyisetan, Borja Balle, Thomas Drake, and Tom Diethe. 2020. Privacy and Utility-Preserving Textual Analysis via Calibrated Multivariate Perturbations. In *Proceedings of the 13th International Conference on Web Search and Data Mining*. 178–186.
- [13] S. Kasiviswanathan, H. Lee, K. Nissim, Sofya Raskhodnikova, and A. D. Smith. 2008. What Can We Learn Privately? *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science* (2008), 531–540.
- [14] Mathias Lécuyer, Vaggelis Atlidakis, Roxana Geambasu, D. Hsu, and Suman Jana. 2019. Certified Robustness to Adversarial Examples with Differential Privacy. *Proceedings of the 2019 IEEE Symposium on Security and Privacy* (2019), 656–672.
- [15] Cheng Li, Mingyang Zhang, Michael Bendersky, H. Deng, Donald Metzler, and Marc Najork. 2019. Multi-view Embedding-based Synonyms for Email Search. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 575–584.
- [16] Yitong Li, Timothy Baldwin, and Trevor Cohn. 2018. Towards Robust and Privacy-preserving Text Representations. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. 25–30.
- [17] Y. Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *ArXiv abs/1907.11692* (2019).
- [18] L. Lyu, Xuanli He, and Yitong Li. 2020. Differentially Private Representation for NLP: Formal Guarantee and An Empirical Study on Privacy and Fairness. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. 2355–2365.
- [19] Lingjuan Lyu, Yitong Li, Xuanli He, and Tong Xiao. 2020. Towards Differentially Private Text Representations. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1813–1816.
- [20] Lingjuan Lyu, Han Yu, and Q. Yang. 2020. Threats to Federated Learning: A Survey. *ArXiv abs/2003.02133* (2020).
- [21] Lingjuan Lyu, Jiangshan Yu, Karthik Nandakumar, Yitong Li, Xingjun Ma, Jiong Jin, H. Yu, and K. S. Ng. 2020. Towards Fair and Privacy-Preserving Federated Deep Models. *IEEE Transactions on Parallel and Distributed Systems* 31, 11 (2020), 2524–2541.
- [22] H. McMahan, Eider Moore, D. Ramage, S. Hampson, and Blaise Agüera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*.
- [23] H. McMahan, D. Ramage, Kunal Talwar, and L. Zhang. 2018. Learning Differentially Private Recurrent Language Models. In *Proceedings of the 6th International Conference on Learning Representations*.
- [24] Tomas Mikolov, Ilya Sutskever, Kai Chen, G. S. Corrado, and J. Dean. 2013. Distributed Representations of Words and Phrases and their Compositionality. In *Advances in Neural Information Processing Systems*. 3111–3119.
- [25] Xudong Pan, Mi Zhang, Shouling Ji, and Min Yang. 2020. Privacy Risks of General-Purpose Language Models. *Proceedings of the 2020 IEEE Symposium on Security and Privacy (SP)* (2020), 1314–1331.
- [26] Chen Qu, Liu Yang, Minghui Qiu, W. Croft, Yongfeng Zhang, and Mohit Iyyer. 2019. BERT with History Answer Embedding for Conversational Question Answering. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1133–1136.
- [27] Reza Shokri and Vitaly Shmatikov. 2015. Privacy-preserving deep learning. *Proceedings of the 53rd Annual Allerton Conference on Communication, Control, and Computing* (2015), 909–910.
- [28] R. Socher, Alex Perelygin, J. Wu, Jason Chuang, Christopher D. Manning, A. Ng, and Christopher Potts. 2013. Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*. 1631–1642.
- [29] Congzheng Song and Ananth Raghunathan. 2020. Information Leakage in Embedding Models. *ArXiv abs/2004.00053* (2020).
- [30] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2018. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. *ArXiv abs/1804.07461* (2018).
- [31] Tianhao Wang, J. Blocki, N. Li, and S. Jha. 2017. Locally Differentially Private Protocols for Frequency Estimation. In *Proceedings of the 26th USENIX Conference on Security Symposium*. 729–745.
- [32] Benjamin Weggenmann and Florian Kerschbaum. 2018. SynTF: Synthetic and Differentially Private Term Frequency Vectors for Privacy-Preserving Text Mining. In *Proceedings of the 41st International ACM SIGIR Conference on Research and Development in Information Retrieval*. 305–314.
- [33] Xi Wu, Fengang Li, A. Kumar, K. Chaudhuri, S. Jha, and J. Naughton. 2017. Bolt-on Differential Privacy for Scalable Stochastic Gradient Descent-based Analytics. *Proceedings of the 2017 ACM International Conference on Management of Data* (2017), 1307–1322.
- [34] Y. Wu, Mike Schuster, Z. Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, M. Krikun, Yuan Cao, Q. Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, M. Johnson, X. Liu, L. Kaiser, S. Gouws, Y. Kato, Taku Kudo, H. Kazawa, K. Stevens, G. Kurian, Nishant Patil, W. Wang, C. Young, J. Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, G. S. Corrado, Macduff Hughes, and J. Dean. 2016. Google’s Neural Machine Translation System: Bridging the Gap between Human and Machine Translation. *ArXiv abs/1609.08144* (2016).
- [35] Z. Yang, Zihang Dai, Y. Yang, J. Carbonell, R. Salakhutdinov, and Quoc V. Le. 2019. XLNet: Generalized Autoregressive Pretraining for Language Understanding. In *Advances in Neural Information Processing Systems*. 5753–5763.
- [36] Hamed Zamani, M. Dehghani, W. Croft, E. Learned-Miller, and J. Kamps. 2018. From Neural Re-Ranking to Neural Ranking: Learning a Sparse Representation for Inverted Indexing. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. 497–506.
- [37] Hongfei Zhang, Xia Song, Chenyan Xiong, C. Rosset, P. Bennett, Nick Craswell, and Saurabh Tiwary. 2019. Generic Intent Representation in Web Search. In *SIGIR*.
- [38] X. Zhang, J. Zhao, and Y. LeCun. 2015. Character-level Convolutional Networks for Text Classification. In *Advances in Neural Information Processing Systems*. 649–657.
- [39] Zhe Zhao, T. Liu, Shen Li, Bofang Li, and X. Du. 2017. Ngram2vec: Learning Improved Word Representations from Ngram Co-occurrence Statistics. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. 244–253.
- [40] Y. Zhu, Ryan Kiros, R. Zemel, R. Salakhutdinov, R. Urtasun, A. Torralba, and S. Fidler. 2015. Aligning Books and Movies: Towards Story-Like Visual Explanations by Watching Movies and Reading Books. *Proceedings of the 2015 IEEE International Conference on Computer Vision* (2015), 19–27.