

# Differentially Private Representation for NLP: Formal Guarantee and An Empirical Study on Privacy and Fairness

Lingjuan Lyu<sup>1</sup>, Xuanli He<sup>†</sup> and Yitong Li<sup>\*</sup>

<sup>1</sup>National University of Singapore, lyulj@comp.nus.edu.sg

<sup>†</sup>Monash University, xuanli.he1@monash.edu

<sup>\*</sup>The University of Melbourne, yitongli4@student.unimelb.edu.au

## Abstract

It has been demonstrated that hidden representation learned by deep model can encode private information of the input, hence can be exploited to recover such information with reasonable accuracy. To address this issue, we propose a novel approach called Differentially Private Neural Representation (DPNR) to preserve privacy of the extracted representation from text. DPNR utilises Differential Privacy (DP) to provide formal privacy guarantee. Further, we show that masking words via dropout can further enhance privacy. To maintain utility of the learned representation, we integrate DP-noisy representation into a robust training process to derive a robust target model, which also helps for model fairness over various demographic variables. Experimental results on benchmark datasets under various parameter settings demonstrate that DPNR largely reduces privacy leakage without significantly sacrificing the main task performance.

## 1 Introduction

Many language applications have involved deep learning techniques to learn text representation through neural models (Bengio et al., 2003; Mikolov et al., 2013; Devlin et al., 2019), performing composition over the learned representation for downstream tasks (Collobert et al., 2011; Socher et al., 2013). However, the input text often provides sufficient clues to portray the author, such as gender, age, and other important attributes. For example, sentiment analysis tasks often have privacy implications for authors whose text is used to train models. Many user attributes have been shown to be easily detectable from online review data, as used extensively in sentiment analysis results (Hovy et al., 2015; Pothast et al., 2017). Private information can take the form of key phrases explicitly contained in the text. However, it can

also be implicit. For example, demographic information about the author of a text can be predicted with above chance accuracy from linguistic cues in the text itself (Preoŧiuc-Pietro et al., 2015).

On the other hand, even the learned representation, rather than the text itself, may still contain sensitive information and incur significant privacy leakage. One might argue that sensitive information like gender, age, location and password should not be leaked out and should have been removed from representation. However, on the intermediate representation level, which is trained from the input text to contain useful features for the prediction task, it can meanwhile encode personal information which might be exploited for adversarial usages, especially a modern deep learning model has vastly more capacity than they need to perform well on their tasks. And, it has been justified that an attacker can recover private variables with higher-than-chance accuracy, only using hidden representation (Li et al., 2018; Coavoux et al., 2018). Therefore, the fact that representations appear to be abstract real-numbered vectors should not be misconstrued as being safe.

The naive solution of removing protected attributes is insufficient: other features may be highly correlated with, and thus predictive of, the protected attributes (Pedreshi et al., 2008). To tackle with these privacy issues, Li et al. (2018) proposed to train deep models with adversarial learning, which explicitly obscures individuals' private information, while improves the robustness and privacy of neural representation in part-of-speech tagging and sentiment analysis tasks. In a parallel study, Coavoux et al. (2018) proposed defence methods based on modifications of the training objective of the main model. However, both works provide only empirical improvements in privacy, without any formal guarantees. Prior works have approached formal differential privacy guarantee by training dif-

ferentially private deep models (Abadi et al., 2016; McMahan et al., 2018; Yu et al., 2019). However, these works generally only considered the training data privacy rather than the test data privacy. While cryptographic methods can be used for privacy protection, it could be resource-hungry or overly complex for the user.

To alleviate the above limitations, we take inspirations from differential privacy (Dwork and Roth, 2014) to provide formal privacy guarantee of the extracted representation from user-authored text. Meanwhile, we propose a robust training algorithm to derive a robust target model to maintain utility, which also offers fairness as a by-product. To the best of our knowledge, our work is the only work to date that can provide formal differential privacy guarantee of the extracted representation, while ensuring fairness.

**Our main contributions** include:

- For the first time, the privacy of the extracted neural representation from text is formally quantified in the context of differential privacy. A novel approach called Differentially Private Neural Representation (DPNR) is proposed to perturb the extracted representation. Also, we prove that masking words via dropout can further enhance privacy.
- To maintain utility, we propose a robust training algorithm that incorporates the noisy training representation in the training process to derive a robust target model, which also reduces model discrimination in most cases.
- On benchmark datasets across various domains and multiple tasks, we empirically demonstrate that our approach yields comparable accuracy to the non-private baseline on the main task, while significantly outperforms the non-private baseline and adversarial learning on the privacy task<sup>1</sup>.

## 2 Preliminary: Differential Privacy

Differential privacy (Dwork and Roth, 2014) provides a mathematically rigorous definition of privacy and has become a de facto standard for privacy analysis. Within DP framework, there are two general settings: central DP (CDP) and local DP (LDP).

In CDP, a trusted data curator answers queries or releases differentially private models by using ran-

domisation mechanisms (Dwork and Roth, 2014; Abadi et al., 2016; Yu et al., 2019). For scenarios where data are sourced from end users, and end users do not trust any third parties, DP should be enforced in a “local” manner to enable end users to perturb their data before publication, which is termed as LDP (Dwork and Roth, 2014; Duchi et al., 2013). Compared with CDP, LDP offers a stronger level of protection.

In our system, we aim to protect the test-phase privacy of the extracted neural representations from end users, we therefore adopt LDP. LDP has shown the advantage that the data is randomised before individuals disclose their personal information, so the server and the middle eavesdropper can never see or receive the raw data. In terms of LDP mechanisms, randomised response (Warner, 1965; Duchi et al., 2013) and its variants have been widely used for aggregating statistics, such as frequency estimation, heavy hitter estimation, etc (Erlingsson et al., 2014).

**Definition 2.1.** Let  $\mathcal{A} : \mathcal{D} \rightarrow \mathcal{O}$  be a randomised algorithm mapping a data entry in  $\mathcal{D}$  to  $\mathcal{O}$ . The algorithm  $\mathcal{A}$  is  $(\epsilon, \delta)$ -local differentially private if for all data entries  $\mathbf{x}, \mathbf{x}' \in \mathcal{D}$  and all outputs  $o \in \mathcal{O}$ , we have

$$\Pr\{\mathcal{A}(\mathbf{x}) = o\} \leq \exp(\epsilon) \Pr\{\mathcal{A}(\mathbf{x}') = o\} + \delta$$

If  $\delta = 0$ ,  $\mathcal{A}$  is said to be  $\epsilon$ -local differentially private.

A formal definition of LDP is provided in Definition 2.1, The privacy parameter  $\epsilon$  captures the privacy loss consumed by the output of the algorithm:  $\epsilon = 0$  ensures perfect privacy in which the output is independent of its input, while  $\epsilon \rightarrow \infty$  gives no privacy guarantee.

For every pair of adjacent inputs  $\mathbf{x}$  and  $\mathbf{x}'$ , differential privacy requires that the distribution of  $\mathcal{A}(\mathbf{x})$  and  $\mathcal{A}(\mathbf{x}')$  are “close” to each other where closeness are measured by the privacy parameters  $\epsilon$  and  $\delta$ . Typically, the inputs  $\mathbf{x}$  and  $\mathbf{x}'$  are adjacent inputs when all the attributes of one record are modified. In real scenario, the adjacent input is an application specific notion. For example, a sentence is divided into several items for every 5 words, and two sentences are considered to be adjacent if they differ by at most 5 consecutive words (Wang et al., 2018). In this work, we consider a word-level DP, i.e., two inputs are considered to be adjacent if they differ by at most 1 word. For brevity, we use  $(\epsilon, \delta)$ -DP to represent  $(\epsilon, \delta)$ -LDP for the rest

<sup>1</sup>code and preprocessed datasets are available at: <https://github.com/xlhex/dpnlp.git>

of the paper. We remark that all the randomisation mechanisms used for CDP, including Laplace mechanism and Gaussian mechanism (Dwork and Roth, 2014), can be individually used by each party to inject noise into local data to ensure LDP before releasing (Lyu et al., 2020a; Yang et al., 2020; Lyu et al., 2020b; Sun and Lyu, 2020). In particular, we adopt Laplace Mechanism which ensures  $\epsilon$ -DP with  $\delta = 0$  throughout the paper.

In a nutshell, data universe can be expressed as  $\mathcal{D} = (\mathcal{X}, \mathcal{A}, \mathcal{Y})$ , which will be convenient to partition as  $(\mathcal{X}, \mathcal{Y}) \times \mathcal{A}$  (Jagielski et al., 2018). Given one person’s record  $\mathbf{x}$ , we can write it as a pair  $\mathbf{x} = (\mathbf{x}_I, \mathbf{x}_S)$  where  $\mathbf{x}_I \in (\mathcal{X}, \mathcal{Y})$  represents the insensitive attributes and  $\mathbf{x}_S \in \mathcal{A}$  represents the sensitive attributes. Our main goal is to promise differential privacy only with respect to the sensitive attributes. Write  $\mathbf{x}_S \sim \mathbf{x}'_S$  to denote that  $\mathbf{x}_S$  and  $\mathbf{x}'_S$  differ in exactly one coordinate (i.e. one word/token in NLP domain). An algorithm is  $(\epsilon, \delta)$ -differentially private in the sensitive attributes if for all  $\mathbf{x}_I \in (\mathcal{X}, \mathcal{Y})$  and for all  $\mathbf{x}_S \sim \mathbf{x}'_S \in \mathcal{A}$  and for all  $O \subseteq \mathcal{O}$ , we have:

$$\mathbb{P} [M(\mathbf{x}_I, \mathbf{x}_S) \in O] \leq e^\epsilon \mathbb{P} [M(\mathbf{x}_I, \mathbf{x}'_S) \in O] + \delta$$

**Post-processing.** DP enjoys a well-known post-processing property (Dwork and Roth, 2014): any computation applied to the output of an  $(\epsilon, \delta)$ -DP algorithm remains  $(\epsilon, \delta)$ -DP. This nice property allows the attacker to implement any sophisticated post-processing function on the privatised representation from the user, without compromising DP or making it less differentially private.

### 3 Main Framework

#### 3.1 Attack Scenario

As indicated in §1, uploading raw input or representations to a server takes the risk of revealing sensitive information to the eavesdropper who eavesdrops on the hidden representation and tries to recover private information of the input text. Hence, similar as Coavoux et al. (2018), we consider an attack scenario during inference phase in Figure 1, which consists of three parts: (i) a *feature extractor* to extract latent representation of any test input  $\mathbf{x}$ ; (ii) a *main classifier* to predict the label  $y$  from the extracted latent representation; (iii) and an *attacker* (eavesdropper) who aims to infer some private information  $\mathbf{z}$  contained in  $\mathbf{x}$ , from the latent representation of  $\mathbf{x}$  used by the main classifier. In this scenario, each example consists of a triple

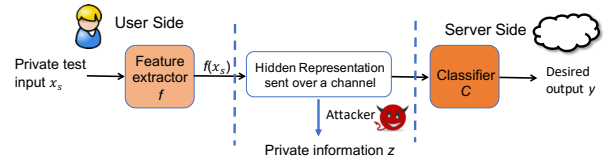


Figure 1: Attack scenario during inference phase.




Figure 2: Illustration of the proposed main framework.

$(\mathbf{x}, y, \mathbf{z})$ , where  $\mathbf{x}$  is an input text,  $y$  is a single label (e.g. topic or sentiment), and  $\mathbf{z}$  is a vector of private information contained in  $\mathbf{x}$ . Such attack would occur in scenarios where the computation of a neural network is shared across multiple devices. For example, phone users send their learned representations to the cloud for grammar correction or translation (Li et al., 2018), or to obtain the classification result, e.g., the topic of the text or its sentiment (Li et al., 2017).

#### 3.2 Methodology

To defend against the middle eavesdropper, we aim to design an approach that can preserve privacy of the extracted test representation from the user without significantly degrading the main task performance. To achieve this goal, we introduce a DP noise layer after a predefined feature extractor (determined by the server), which results in differentially private representation that can be transferred to the server for classification (the topic of the text or its sentiment), as shown in Figure 2.

In terms of model training on the server, theoretically, one could remove the noise layer and conduct non-private training by following Equation 1:

$$\mathcal{L}(\mathbf{x}, \mathbf{y}) = \mathcal{X}(C(f(\mathbf{x})), \mathbf{y}) \quad (1)$$

where  $f$  is the feature extractor,  $C$  is the classifier,  $\mathbf{y}$  is the true label, and  $\mathcal{X}$  denotes the cross entropy loss function.

However, doing so may deteriorate test performance, due to the injected noise in the test representation. To improve model robustness to the noisy representation, we put forward a robust training algorithm by incorporating a noise layer which

---

**Algorithm 1** Robust Training on the Server

---

**Input:** Training record  $(\mathbf{x}_t, \mathbf{y}_t)$ ; Feature extractor  $f$ ; Classifier  $C$ .

- 1: Extraction:  $\mathbf{x}_r \leftarrow f(\mathbf{x}_t)$ ;
  - 2: Normalization:  $\mathbf{x}_r \leftarrow \mathbf{x}_r - \min(\mathbf{x}_r) / (\max(\mathbf{x}_r) - \min(\mathbf{x}_r))$ ;
  - 3: Perturbation:  $\hat{\mathbf{x}}_r \leftarrow \mathbf{x}_r + \mathbf{r}, r_i \sim \text{Lap}(b)$ ;
  - 4: Calculate loss  $\mathcal{L} = \mathcal{X}(C(\hat{\mathbf{x}}_r), \mathbf{y}_t)$  and do backpropagation to update  $f$  and  $C$ .
- 

adds the same level of noise as the test phase in the training process as well. Therefore, the robust training objective can be re-written as:

$$\mathcal{L}(\mathbf{x}, \mathbf{y}) = \mathcal{X}(C(f(\mathbf{x}) + \mathbf{r}), \mathbf{y}) \quad (2)$$

The detailed robust training process on the server is given in Algorithm 1. After the robust target model is built, server then provides a *feature extractor*  $f$  to the user, as illustrated in Figure 2.

### 3.3 Privacy Guarantee

Let  $f(\mathbf{x}) = \mathbf{x}_r \in \mathbb{R}^k$  be the extracted representation from  $\mathbf{x}$  by feature extractor  $f$ , and to apply  $\epsilon$ -DP to the extracted neural representation, we inject Laplace noise  $\mathbf{r}$  to  $\mathbf{x}_r = f(\mathbf{x})$  as follows:

$$\hat{\mathbf{x}}_r = \mathbf{x}_r + \mathbf{r},$$

where the coordinates  $\mathbf{r} = \{r_1, r_2, \dots, r_k\}$  are i.i.d. random variables drawn from the Laplace distribution defined by  $\text{Lap}(b)$ , where the noise scale  $b = \frac{\Delta f}{\epsilon}$ ,  $\epsilon$  is the privacy budget and  $\Delta f$  is the sensitivity of the extracted representation.

#### 3.3.1 Formal Privacy Guarantee

Algorithm 2 outlines how to derive differentially private neural representation from the feature extractor  $f$ . Each user first feeds its masked sensitive record  $\tilde{\mathbf{x}}_s$  into a feature extractor to extract representation  $\mathbf{x}_r \in \mathbb{R}^k$ .

Note that to apply additive noise mechanism, the sensitivity  $\Delta$  of the output representation  $\mathbf{x}_r = f(\mathbf{x})$  needs to be determined. Estimating the true sensitivity of  $\mathbf{x}_r$  is challenging. Instead, we follow Shokri and Shmatikov (2015) to use input-independent bounds by enforcing a  $[0,1]$  range on the extracted representation, hence bounding the sensitivity of each element of the extracted representation with 1, i.e.,  $\Delta f = 1$ . Limiting the range of the extracted representation can also improve the training process by helping to avoid overfitting.

A formal statement for the privacy guarantees of Algorithm 2 is provided in Theorem 1.

---

**Algorithm 2** Differentially Private Neural Representation (DPNR)

---

**Input:** Each sensitive record  $\mathbf{x}_s \in \mathbb{R}^d$ ; Feature extractor  $f$ .

**Parameters:** Dropout vector  $\mathbf{I}_n \in \{0, 1\}^d$ ;

- 1: Word Dropout:  $\tilde{\mathbf{x}}_s \leftarrow \mathbf{x}_s \odot \mathbf{I}_n$ ;
  - 2: Extraction:  $\mathbf{x}_r \leftarrow f(\tilde{\mathbf{x}}_s)$ ;
  - 3: Normalisation:  $\mathbf{x}_r \leftarrow \mathbf{x}_r - \min(\mathbf{x}_r) / (\max(\mathbf{x}_r) - \min(\mathbf{x}_r))$ ;
  - 4: Perturbation:  $\hat{\mathbf{x}}_r \leftarrow \mathbf{x}_r + \mathbf{r}, r_i \sim \text{Lap}(b)$ ;
- Output:** Perturbed representation  $\hat{\mathbf{x}}_r$ .
- 

**Theorem 1.** *Let the entries of the noise vector  $\mathbf{r}$  be drawn from  $\text{Lap}(b)$  with  $b = \frac{\Delta f}{\epsilon}$ . Then Algorithm 2 is  $\epsilon$ -differentially private.*

#### 3.3.2 Word Dropout Enhances Privacy

In NLP, each input is a sequence composed of words/tokens  $\{w_1, \dots, w_d\}$ . Under word-level DP, two sentences are considered to be adjacent inputs if they differ by at most 1 word (i.e., 1 edit distance). In this scenario, to lower privacy budget without significantly degrading the inference performance, we borrow the idea of nullification (Wang et al., 2018) and apply it to word dropout.

For each sensitive record  $\mathbf{x}_s$ , words are masked by a dropout operation before DP perturbation. Given a sensitive input  $\mathbf{x}_s$  that consists of  $d$  words, dropout performs word-wise multiplication of  $\mathbf{x}_s$  with  $\mathbf{I}_n$ , i.e.,  $\tilde{\mathbf{x}}_s \leftarrow \mathbf{x}_s \odot \mathbf{I}_n$ , where  $\mathbf{I}_n \in \{0, 1\}^d$ . It can be either specified by users to mask the highly sensitive words or generated randomly. The number of zeros in  $\mathbf{I}_n$  is determined by  $d \cdot \mu$ , where  $\mu$  is the dropout rate. The zeros are located in  $\mathbf{I}_n$  conforming to the uniform distribution.

As stated in Theorem 2, word dropout in combination with any  $\epsilon$ -differentially private mechanism provides a tighter privacy bound in the context of word-level DP. A detailed proof follows.

**Theorem 2.** *Given an input  $\mathbf{x} \in D$ , suppose  $A(\mathbf{x}) = f(\mathbf{x}) + \mathbf{r}$  is  $\epsilon$ -differentially private, let  $\mathbf{I}_n$  with dropout rate  $\mu$  be applied to  $\mathbf{x}$ , i.e.,  $\tilde{\mathbf{x}} = \mathbf{x} \odot \mathbf{I}_n$ , then  $A(\tilde{\mathbf{x}})$  is  $\epsilon'$ -differentially private, where  $\epsilon' = \ln[(1 - \mu) \exp(\epsilon) + \mu]$ .*

*Proof.* Suppose there are two adjacent inputs  $\mathbf{x}_1$  and  $\mathbf{x}_2$  that differ only in the  $i$ -th coordinate (word), say  $x_{1i} = v, x_{2i} \neq v$ . For arbitrary binary vector  $\mathbf{I}_n$ , after dropout,  $\tilde{\mathbf{x}}_1 = \mathbf{x}_1 \odot \mathbf{I}_n, \tilde{\mathbf{x}}_2 = \mathbf{x}_2 \odot \mathbf{I}_n$ , there are two possible cases, i.e.,  $I_{ni} = 0$ , and  $I_{ni} = 1$ .

Case 1:  $I_{ni} = 0$ . Since  $\mathbf{x}_1$  and  $\mathbf{x}_2$  differ only in  $i$ -th coordinate, after dropout,  $\tilde{x}_{1i} = \tilde{x}_{2i} = 0$ ,

hence  $\mathbf{x}_1 \odot \mathbf{I}_n = \mathbf{x}_2 \odot \mathbf{I}_n$ . It then follows

$$\Pr\{\mathcal{A}(\mathbf{x}_1 \odot \mathbf{I}_n) = S\} = \Pr\{\mathcal{A}(\mathbf{x}_2 \odot \mathbf{I}_n) = S\}.$$

Case 2:  $I_{ni} = 1$ . Since  $\mathbf{x}_1$  and  $\mathbf{x}_2$  differ only in the value of their  $i$ -th coordinate, after dropout,  $\tilde{x}_{1i} = x_{1i} = v$ ,  $\tilde{x}_{2i} = x_{2i} \neq v$ , hence  $\tilde{\mathbf{x}}_1$  and  $\tilde{\mathbf{x}}_2$  remain adjacent inputs that differ only in  $i$ -th coordinate. Because  $\mathcal{A}(\mathbf{x})$  is  $\epsilon$ -differentially private, it then follows

$$\Pr\{\mathcal{A}(\mathbf{x}_1 \odot \mathbf{I}_n) = S\} \leq \exp(\epsilon) \Pr\{\mathcal{A}(\mathbf{x}_2 \odot \mathbf{I}_n) = S\}.$$

Combine these two cases, and use the fact that  $\Pr[I_{ni} = 0] = \mu$ , we have:

$$\begin{aligned} & \Pr\{\mathcal{A}(\mathbf{x}_1 \odot \mathbf{I}_n) = S\} \\ &= \mu \Pr\{\mathcal{A}(\mathbf{x}_1 \odot \mathbf{I}_n) = S\} + (1 - \mu) \Pr\{\mathcal{A}(\mathbf{x}_1 \odot \mathbf{I}_n) = S\} \\ &\leq \mu \Pr\{\mathcal{A}(\mathbf{x}_2 \odot \mathbf{I}_n) = S\} + (1 - \mu) [\exp(\epsilon) \Pr\{\mathcal{A}(\mathbf{x}_2 \odot \mathbf{I}_n) = S\}] \\ &= [(1 - \mu) \exp(\epsilon) + \mu] \Pr\{\mathcal{A}(\mathbf{x}_2 \odot \mathbf{I}_n) = S\} \\ &= \exp\{\ln[(1 - \mu) \exp(\epsilon) + \mu]\} \Pr\{\mathcal{A}(\mathbf{x}_2 \odot \mathbf{I}_n) = S\} \end{aligned}$$

Therefore, after dropout, the privacy budget is lowered to  $\epsilon' = \ln[(1 - \mu) \exp(\epsilon) + \mu]$ .  $\square$

Since the perturbed representation  $\mathcal{A}(\mathbf{x}) = f(\mathbf{x}) + \mathbf{r}$  is  $\epsilon$ -differentially private, combining dropout beforehand, the privacy budget is lowered to  $\epsilon' = \ln[(1 - \mu) \exp(\epsilon) + \mu]$ , hence improving privacy guarantee. Apparently, a high value of  $\mu$  has a positive impact on the privacy but a potential negative impact on the utility. In particular, when  $\mu = 1$ , all  $d$  words will be masked, which gives the highest privacy, *i.e.*,  $\epsilon' = 0$ , but totally destroys inference performance. Hence, a smaller value of  $\mu$  is preferred to trade off privacy and accuracy.

## 4 Experiments

In this section, we conduct comprehensive studies over different tasks and datasets to examine the efficacy of the proposed algorithm from three facets: 1) main task performance, 2) privacy and 3) target model fairness.

### 4.1 Task and Dataset

We use two natural language processing tasks: 1) sentiment analysis and 2) topic classification, with a range of benchmark datasets across various domains. Table 1 summarises the statistics of the used datasets.

#### 4.1.1 Sentiment Analysis

Trustpilot Sentiment dataset (Hovy et al., 2015) contains reviews associated with a sentiment score on a five point scale, and each review is associated

with 3 attributes: gender, age and location, which are self-reported by users. The original dataset is comprised of reviews from different locations, however in this paper, we only derive TP-US for our study. Following Coavoux et al. (2018), we extract examples containing information of both gender and age, and treat them as the private information. We categorise ‘‘age’’ into two groups: ‘‘under 34’’ (U34) and ‘‘over 45’’ (O45).

#### 4.1.2 Topic Classification

For topic classification, we focus on two genres of documents: news articles and blog posts.

**News article** We use AG news corpus (Del Corso et al., 2005). To ensure a fair comparison, we use the corpus preprocessed by Coavoux et al. (2018)<sup>2</sup>. And the task is to predict the topic label of the document, with four different topics in total.

Regarding the private information in AG, named entities appearing in text are vulnerable to privacy leakage inferred by attackers. In order to simulate the attack, we firstly adopt the NLTK NER system (Bird et al., 2009) to recognise all ‘‘Person’’ entities in the corpus. Then we retain the five most frequent person entities and use them as the private information. Due to the sparsity of name entities, each target entity only appears in very few articles. Hence we select the examples containing at least one of these named entities to mitigate the unbalance and data scarcity. Thus, the attacker aims to identify these five entities as five independent binary classification tasks.

**Blog posts** We derive a blog posts dataset (BLOG) from the blog authorship corpus presented (Schler et al., 2006). However, the original dataset only contains a collection of blog posts associated with authors’ age and gender attributes but does not provide topic annotations. Thus we follow Coavoux et al. (2018) to run the LDA algorithm (Blei et al., 2003) with the topic number of 10 on the whole collection to identify the topic label of each document. Afterwards, we selected posts with single dominating topic ( $> 80\%$ ) and discarded the rest, which results in a dataset with 10 different topics. Similar to TP-US, the private variables are comprised of the age and gender of the author. And the age attribute is binned into two categories, ‘‘under 20’’ (U20) and ‘‘over 30’’ (O30).

<sup>2</sup><https://github.com/mcoavoux/pnet/tree/master/datasets>. We use both ‘‘title’’ and ‘‘description’’ fields as the input document.

Dataset	Private Variable	#Train	#Dev	#Test
TP-US	age, gender	22,142	2,767	2,767
AG	entity	11,657	1,457	1,457
BLOG	age, gender	7,098	887	887

Table 1: Summary of three pre-processed datasets.

	$\epsilon$	TP-US	AG	BLOG
NON-PRIV		85.53	78.75	<b>97.07</b>
	0.05	85.65	<b>80.87</b>	96.69
	0.1	85.52	80.78	96.39
DPNR	0.5	85.52	79.71	96.84
	1	85.36	79.36	96.39
	5	<b>85.87</b>	79.59	96.66

Table 2: Main task accuracy [%] of NON-PRIV and DPNR over 3 datasets with varying  $\epsilon$  and fixed  $\mu = 0$ .

For all three datasets, we randomly split the pre-processed corpus into training, development and test by 8:1:1.

## 4.2 Evaluation Metrics

Similar to Coavoux et al. (2018), we define *senti-ment analysis* and *topic classification* as the main tasks, whereas the inference of private information is considered as the auxiliary tasks of attackers. Each auxiliary task is eavesdropped by one attacker.

We use accuracy to assess the performance for both main tasks. The auxiliary tasks are evaluated via the following metrics:

- For demographic variables (*i.e.*, gender and age):  $1 - X$ , where  $X$  is the average over the accuracies of the prediction by the attacker on these variables.
- For named entities:  $1 - F$ , where  $F$  is the F1 score between the ground truths and the prediction by the attacker on the presence of all named entities.

We denote the value of  $1 - X$  or  $1 - F$  as *empirical privacy*, *i.e.*, the inverse accuracy or F1 score of the attacker, higher means better empirical privacy, *i.e.*, lower attack performance.

## 4.3 Model Selection

**Model and Parameters.** For implementation, owing to its success across multiple NLP tasks, we apply *BERT* base (Devlin et al., 2019) to the classification tasks. Specially, BERT takes a text input, then generates a representation which embeds holistic information. We apply a dropout to this

representation before a softmax layer, which is responsible for label classification. We run 4 epochs on the training set, and choose the checkpoint with the best loss on the dev set.

After we obtain a well-trained target model, we partition it into two parts, BERT model acts as the feature extractor  $f$  in Figure 2, which could be deployed on users’ devices, while the remaining layers act as the classifier on the server. In our implementation, privacy is enforced in the hidden representation extracted by the feature extractor as shown by Algorithm 2. For attack classifier, we utilise a 2-layer MLP with 512 hidden units and ReLU activation trained over the target model, which delivers the best attack performance on the dev set in our preliminary experiments.

We report the averaged results over 5 independent runs for all experiments.

## 4.4 Performance Analysis of Target Model

Firstly, we would like to study how the privacy parameters ( $\epsilon, \mu$ ) in Theorem 1 and 2 affect the accuracy of main tasks. We investigate this using different parameter settings, varying one parameter while fixing the other.

### 4.4.1 Impact of Privacy Budget $\epsilon$

To analyse the impact of different privacy budget  $\epsilon$  on accuracy, we choose  $\epsilon \in \{0.05, 0.1, 0.5, 1, 5\}$  with fixed  $\mu = 0$ . Noted that to provide reasonable privacy guarantee,  $\epsilon$  should be set below 10 (Hamm et al., 2015; Abadi et al., 2016). Moreover,  $\epsilon \leq 1$  means a relatively tight privacy guarantee. Surprisingly, there is no obvious relationship between accuracy and  $\epsilon$ . We speculate the denoising training procedure of BERT and layernorm (Ba et al., 2016) make BERT resistant to the injected noises, which can maintain the performance of the main tasks. We will conduct an in-depth study on this in the future.

Table 2 shows that in most cases, our method can achieve comparable performance to the non-private baseline, across all  $\epsilon$  even when the noise level is high ( $\epsilon = 0.05$ ), which validates the robustness of our method to DP noise. It also implies that the DP-noised representation not only preserves privacy, but also retains general information for the main task.

### 4.4.2 Impact of Dropout Rate $\mu$

Similarly, we study how the word dropout rate  $\mu$  affects accuracy-privacy trade-off. Table 3 reports

	$\mu$	TP-US	AG	BLOG
NON-PRIV		<b>85.53</b>	78.75	<b>97.07</b>
DPNR	0.1	85.53	<b>80.71</b>	96.05
	0.3	84.85	79.18	93.76
	0.5	83.51	77.42	90.98
	0.8	80.70	69.57	82.94

Table 3: Main task accuracy [%] of NON-PRIV and DPNR over 3 datasets with varying  $\mu$  and fixed  $\epsilon = 1$ .

the performance of different models under different  $\mu \in \{0.1, 0.3, 0.5, 0.8\}$  with fixed  $\epsilon = 1$ . In most cases, as  $\mu$  becomes larger, accuracy starts to degrade as expected. However, as indicated in Theorem 2, higher  $\mu$  results in better privacy as well. Moreover,  $\mu = 0.5$  can still provide a relatively high accuracy, while privacy budgets are reduced to  $\epsilon' = \ln[(1 - \mu) \exp(\epsilon) + \mu] = 0.62$ .

Overall, both results demonstrate that our DPNR can protect privacy of the extracted representations of user-authored text, without significantly affecting the main task performance.

#### 4.5 Attack Model

Apart from formal privacy guarantee from DP, we use the performance of the diagnostic classifier of the attackers for empirical privacy. To fairly compare with the standard training and adversarial training in previous work (Coavoux et al., 2018), we train an attack model that is trying to predict private variables from the representation. We measure the empirical privacy of a hidden representation by the ability of an attacker to predict accurately specific private information from it. If its empirical privacy (*c.f.*, Section 4.2) is low, then an eavesdropper can easily recover information about the input. In contrast, a higher empirical privacy (close to that of a most-frequent label baseline) suggests that  $x_r$  mainly contains useful information for the main task, while other private information is erased.

To study the relationship between DP and empirical privacy, we numerically investigate the impact of the different differential privacy budgets on empirical privacy. Recall that the empirical privacy is measured by  $1 - X/F$ , and the higher is better. Figure 3 shows that with the increase of the budget, empirical privacy across all datasets demonstrate a decreasing trend, especially for AG, which well aligns with DP where the higher value of  $\epsilon$  implies lower formal privacy guarantee. Since  $\epsilon = 0.05$  provides the best privacy guarantee, we




Figure 3: Results of privacy protection over TP-US, AG and BLOG datasets across different differential privacy budgets. X-axis is the differential privacy budget  $\epsilon$ , while Y-axis indicates the empirical privacy (see §4.2).

fix  $\epsilon = 0.05$  and  $\mu = 0$  as a default setting in the rest of this section, unless otherwise mentioned.

**How private are the noisy neural representations?** For empirical privacy, we investigate whether our DPNR can provide better attack resistance compared with the adversarial learning (ADV) (Coavoux et al., 2018) and non-private training method (NON-PRIV), which indicates a lower bound. We also report the majority class prediction (MAJORITY) as an upper bound.

Table 4 shows that the attack model can indeed recover private information with reasonable accuracy when targeting towards the non-private representations, manifesting that representations inadvertently capture sensitive information about users, apart from the useful information for the main task. By contrast, our DPNR significantly reduces the amount of information encoded in the extracted representation, as validated by the substantially higher empirical privacy than NON-PRIV across all datasets. We also observe that our DPNR achieves comparable empirical privacy to the majority class (MAJORITY), and consistently outperforms the adversarial learning (ADV) from Coavoux et al. (2018), which confirms the argument of Elazar and Goldberg (2018) that adversarial learning can not fully remove sensitive demographic traits from the data representations. Conversely, the post-processing property of DP ensures that the privacy loss of the extracted representation cannot be increased even by the most sophisticated attacker.

This claim can be further confirmed by Table 5, which reports the accuracy of the attacker on classifying whether a named entities is absent or pre-

	TP-US		AG		BLOG	
	Main	Priv.	Main	Priv.	Main	Priv.
MAJORITY	79.40	36.39	57.79	49.34	34.16	46.96
NON-PRIV	85.53	34.71	78.75	23.24	97.07	33.88
ADV	-0.25	+0.67	-21.71	+26.43	-2.44	+1.16
DPNR	<b>+0.12</b>	<b>+3.66</b>	<b>+2.12</b>	<b>+31.13</b>	-0.38	<b>+15.86</b>

Table 4: Results of the main task and the privacy-protected task on the test sets over different datasets. The relative values are based on NON-PRIV method and **bold** indicates our DPNR achieves better performance than other methods. (See §4.2 for details for metrics.)

	Entity 3		Entity 4		Entity 5	
	A	P	A	P	A	P
ratio [%]	82	18	90	10	91	9
NON-PRIV	96.71	81.99	99.43	47.20	96.93	68.29
ADV	98.57	39.71	100.00	0.00	99.87	12.96
DPNR	90.86	8.46	100.00	0.00	100.00	0.00

Table 5: Accuracy of attack classifier on absence (A) and presence (P) classification of 3 entities over AG.

sented in the document over AG<sup>3</sup>. Generally, both ADV and DPNR can reduce attack accuracy, misleading the attacker classifier to predict most of the shared representations as majority (A). While our DPNR significantly outperforms both NON-PRIV and ADV, corroborating our analysis above.

#### 4.6 Target Model Fairness

Recently, fairness concern has gained lots of attention in NLP community (Bolukbasi et al., 2016; Zhao et al., 2017; Chang et al., 2019; Lu et al., 2018; Sun et al., 2019). Depending on the literature, fairness can have different interpretation. In this section, we further consider the relation between differential privacy and fairness. We ask the research question *whether differential privacy noise can help enhance model fairness?* We focus on a particular scenario of fairness, that is given a specific demographic variable (e.g. gender) a fair model should deliver an equal or similar performance over the subgroups (e.g. male vs. female) (Rudinger et al., 2018; Zhao et al., 2018).

To empirically evaluate the fairness, we take inspirations of Rudinger et al. (2018); Zhao et al. (2018); Li et al. (2018) and partition the test data into sub-groups by the demographic variables, *i.e.*, **age**, **gender** and five person **entities**. Different from predicting demographic variables in attacker (§4.5), we measure the main task accuracy differ-

<sup>3</sup>For space limitation, we only report 3 of 5 entities and the results of other two are similar.

		Gender		Age	
		F	M	U	O
TP-US	ratio [%]	37	63	64	36
	NON-PRIV	83.69	+1.57	84.63	+0.02
	ADV	84.95	+0.19	85.38	-0.46
	DPNR	85.90	<b>+0.49</b>	86.08	+0.31
BLOG	ratio [%]	52	48	46	54
	NON-PRIV	98.07	-2.18	97.05	-1.49
	ADV	93.84	-7.54	91.91	-3.57
	DPNR	98.00	-2.34	97.09	<b>-0.11</b>

Table 6: The accuracy of main tasks among different demographic groups (age and gender) on TP-US and BLOG. “Ratio” means the ratio between two subgroups of the demographic variable. The relative values (M and O) of right subgroups are deviated from the left subgroups (F and U) accordingly.

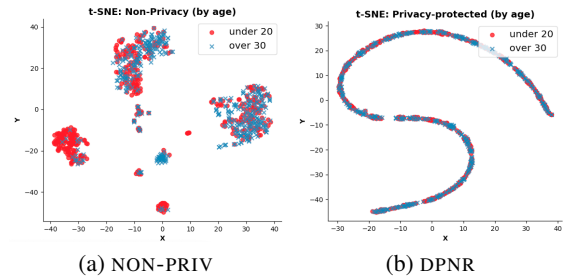


Figure 4: t-SNE plots of the extracted representations over two age subgroups (u20 and o30) of BLOG using NON-PRIV method and proposed DPNR.

ence among subgroups of demographic variables.

In fact, we noticed DPNR can also help mitigate the bias in the representations with respect to the specific demographic or identity attributes, such that the decisions made by our robust target model are able to improve the fairness among the concerned demographic groups.

First of all, as the distribution of the demographic groups in TP-US and BLOG datasets is relatively even, hence there is no significant deviation on the main tasks (see Table 6). However, we still observe an noticeable difference for the **age** group in BLOG and the **gender** attribute in TP-US. To help better understand the phenomenon, we perform further analysis by plotting the non-private and differentially private representations of **age** on BLOG in Figure 4. It can be clearly observed that the patterns of two subgroups are much easier to be distinguished in the non-private representations, while the differentially private representations mostly mix the representations of “under 20” and “over 30”. We speculate that this is a conse-



	Entity 3		Entity 4		Entity 5	
	A	P	A	P	A	P
ratio [%]	82	18	90	10	91	9
NON-PRIV	82.67	-18.82	80.86	-15.44	80.22	-10.78
ADV.	48.92	+7.82	49.64	+6.67	49.43	+9.6
DPNR	83.60	<b>-16.93</b>	81.79	<b>-12.22</b>	81.04	<b>-6.04</b>

Table 7: The accuracy of main tasks among three name entities on AG. A means when the entity is absence, while P indicates the presence of that entity. The relative values are deviated from the subgroup A. **bold** means a statistically significant ( $p < 0.0001$ ) fairness improvement.

quence of the regularising effect of DP.

Table 7 shows the fairness results on AG, where we observe the entity distributions are skewed and the prediction of the NON-PRIV model on the dominant groups is significantly superior to the minority groups, which causes a severe violation in terms of the fairness. Even under such circumstance, our DPNR method can mitigate this skewed bias, achieving more fair prediction than other baselines.

## 5 Discussion

Privacy and fairness are two emerging but important areas in NLP community. Prior efforts predominantly focus on either privacy or fairness (Li et al., 2018; Coavoux et al., 2018; Rudinger et al., 2018; Zhao et al., 2018; Lyu et al., 2020a), but there is no systematic study on how privacy and fairness are related. This work fills this gap, and discovers the impact of differential privacy on model fairness. We empirically show that privacy and fairness can be simultaneously achieved through differential privacy.

We hope that this work highlights the need for more research in the development of effective countermeasures to defend against privacy leakage via model representation and mitigate model bias in a general sense, and *not* only specific to a particular attack. More generally, we hope that our work spurs future interest into developing a better understanding of why differential privacy works.

Meanwhile, differential privacy may incur a reduction in the model’s accuracy. It is worthwhile to explore how to get a better trade-off between privacy, fairness and accuracy.

## 6 Conclusion and Future Work

In this paper, we take the first effort to build differential privacy into the extracted neural representa-

tion of text during inference phase. In particular, we prove that masking the words in a sentence via dropout can further enhance privacy. To maintain utility, we propose a novel robust training algorithm that incorporates a noisy layer into the training process to produce the noisy training representation. Experimental results on benchmark datasets across various tasks, and parameter settings demonstrate that our approach ensures representation privacy without significantly degrading accuracy. Meanwhile, our DP method helps reduce the effects of model discrimination in most cases, achieving better fairness than the non-private baseline. Our work makes a first step towards understanding the connection between privacy and fairness in NLP – which were previously thought of as distinct classes. Moving forward, we believe that our results justify a larger study on various NLP applications and models, which will be our immediate future work.

## Acknowledgments

We would like to thank the anonymous reviewers for their valuable feedback.

## References

- Martín Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on CCS*, pages 308–318.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. 2003. A neural probabilistic language model. *Journal of Machine Learning Research*, 3(Feb):1137–1155.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. ” O’Reilly Media, Inc.”.
- David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3(Jan):993–1022.
- Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016*.

- Kai-Wei Chang, Vinod Prabhakaran, and Vicente Ordonez. 2019. Bias and fairness in natural language processing. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): Tutorial Abstracts*.
- 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*, pages 1–10.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12(Aug):2493–2537.
- Gianna M Del Corso, Antonio Gulli, and Francesco Romani. 2005. Ranking a stream of news. In *Proceedings of the 14th International Conference on World Wide Web*, pages 97–106.
- 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, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- John C Duchi, Michael I Jordan, and Martin J Wainwright. 2013. Local privacy and statistical minimax rates. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 429–438.
- Cynthia Dwork and Aaron Roth. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407.
- Yanai Elazar and Yoav Goldberg. 2018. Adversarial removal of demographic attributes from text data. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 11–21.
- Úlfar Erlingsson, Vasyli Pihur, and Aleksandra Korolova. 2014. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, pages 1054–1067.
- Jihun Hamm, Adam C Champion, Guoxing Chen, Mikhail Belkin, and Dong Xuan. 2015. Crowdml: A privacy-preserving learning framework for a crowd of smart devices. In *2015 IEEE 35th International Conference on Distributed Computing Systems*, pages 11–20. IEEE.
- Dirk Hovy, Anders Johannsen, and Anders Søgaard. 2015. User review sites as a resource for large-scale sociolinguistic studies. In *Proceedings of the 24th International Conference on World Wide Web*, pages 452–461.
- Matthew Jagielski, Michael Kearns, Jieming Mao, Alina Oprea, Aaron Roth, Saeed Sharifi-Malvajerdi, and Jonathan Ullman. 2018. Differentially private fair learning. *arXiv preprint arXiv:1812.02696*.
- Meng Li, Liangzhen Lai, Naveen Suda, Vikas Chandra, and David Z Pan. 2017. Privynet: A flexible framework for privacy-preserving deep neural network training. *arXiv preprint arXiv:1709.06161*.
- 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*, pages 25–30.
- Kaiji Lu, Piotr Mardziel, Fangjing Wu, Preetam Amancharla, and Anupam Datta. 2018. Gender bias in neural natural language processing. *arXiv preprint arXiv:1807.11714*.
- Lingjuan Lyu, Yitong Li, Xuanli He, and Tong Xiao. 2020a. Towards differentially private text representations. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1813–1816.
- Lingjuan Lyu, Yitong Li, Karthik Nandakumar, Jiangshan Yu, and Xingjun Ma. 2020b. How to democratise and protect ai: Fair and differentially private decentralised deep learning. *IEEE Transactions on Dependable and Secure Computing*.
- H Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. 2018. Learning differentially private recurrent language models. In *Proceedings of the 5th International Conference on Learning Representations*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013*, pages 3111–3119.
- Dino Pedreshi, Salvatore Ruggieri, and Franco Turini. 2008. Discrimination-aware data mining. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 560–568.
- Martin Potthast, Francisco Rangel, Michael Tschuggnall, Efstathios Stamatatos, Paolo Rosso, and Benno Stein. 2017. Overview of PAN’17. In *International Conference of CLEF*, pages 275–290.
- Daniel Preoțiuc-Pietro, Vasileios Lampos, and Nikolaos Aletras. 2015. An analysis of the user occupational class through twitter content. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International*

- Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing*, volume 1, pages 1754–1764.
- Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. 2018. [Gender bias in coreference resolution](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 8–14.
- Jonathan Schler, Moshe Koppel, Shlomo Argamon, and James W Pennebaker. 2006. Effects of age and gender on blogging. In *AAAI Spring Symposium: Computational Approaches to Analyzing Weblogs*, volume 6, pages 199–205.
- Reza Shokri and Vitaly Shmatikov. 2015. Privacy-preserving deep learning. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 1310–1321.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment tree-bank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642.
- Lichao Sun and Lingjuan Lyu. 2020. Federated model distillation with noise-free differential privacy. *arXiv preprint arXiv:2009.05537*.
- Tony Sun, Andrew Gaut, Shirlyn Tang, Yuxin Huang, Mai ElSherief, Jieyu Zhao, Diba Mirza, Elizabeth Belding, Kai-Wei Chang, and William Yang Wang. 2019. Mitigating gender bias in natural language processing: Literature review. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1630–1640.
- Ji Wang, Jianguo Zhang, Weidong Bao, Xiaomin Zhu, Bokai Cao, and Philip S Yu. 2018. Not just privacy: Improving performance of private deep learning in mobile cloud. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2407–2416.
- Stanley L Warner. 1965. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69.
- Mengmeng Yang, Lingjuan Lyu, Jun Zhao, Tianqing Zhu, and Kwok-Yan Lam. 2020. Local differential privacy and its applications: A comprehensive survey. *arXiv preprint arXiv:2008.03686*.
- Lei Yu, Ling Liu, Calton Pu, Mehmet Emre Gursesoy, and Stacey Truex. 2019. [Differentially private model publishing for deep learning](#). In *2019 IEEE Symposium on Security and Privacy*, pages 332–349.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2017. [Men also like shopping: Reducing gender bias amplification using corpus-level constraints](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2979–2989.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018. [Gender bias in coreference resolution: Evaluation and debiasing methods](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 15–20.