

DP-MLM: Differentially Private Text Rewriting Using Masked Language Models

Stephen Meisenbacher, Maulik Chevli, Juraj Vladika, and Florian Matthes

Technical University of Munich

School of Computation, Information and Technology

Department of Computer Science

Garching, Germany

{stephen.meisenbacher,maulikk.chevli,juraj.vladika,matthes}@tum.de

Abstract

The task of text privatization using Differential Privacy has recently taken the form of *text rewriting*, in which an input text is obfuscated via the use of generative (large) language models. While these methods have shown promising results in the ability to preserve privacy, these methods rely on autoregressive models which lack a mechanism to contextualize the private rewriting process. In response to this, we propose DP-MLM, a new method for differentially private text rewriting based on leveraging masked language models (MLMs) to rewrite text in a semantically similar *and* obfuscated manner. We accomplish this with a simple contextualization technique, whereby we rewrite a text one token at a time. We find that utilizing encoder-only MLMs provides better utility preservation at lower ϵ levels, as compared to previous methods relying on larger models with a decoder. In addition, MLMs allow for greater customization of the rewriting mechanism, as opposed to generative approaches. We make the code for DP-MLM public and reusable, found at <https://github.com/sjmeis/DPMLM>.

1 Introduction

The study of Differential Privacy (DP) in NLP investigates the integration of the privacy guarantee offered by DP to the textual domain. This is especially timely as concerns of privacy vulnerabilities in NLP models, particularly LLMs, continue to rise in tandem with recent advances in AI.

Looking to DP for safeguarding privacy in text processing is a promising avenue of research, bringing about novel techniques in recent years ranging from DP optimization techniques (Abadi et al., 2016), DP language models (Igamberdiev and Habernal, 2023), or DP text privatization methods (Feyisetan et al., 2020; Mattern et al., 2022; Utpala et al., 2023). DP in NLP does not come without its

challenges, however (Klymenko et al., 2022; Mattern et al., 2022), among them the balance between privacy and utility, as well as the general reasoning of DP in unstructured domains such as text.

State-of-the-art DP text rewriting methods focus on paraphrasing as a proxy for text privatization, either via directly fine-tuning a paraphrase model (Mattern et al., 2022), or by prompting a LLM to generate a rephrased version of an input text (Utpala et al., 2023). By incorporating a DP mechanism into the generation of output tokens, the generated text aims to privatize the input while still maintaining utility and semantic similarity.

The above works do not consider encoder-only models, such as BERT-based models, and furthermore, the advantages of doing so have not been studied in juxtaposition to models with decoders. Therefore, we are motivated to extend DP text rewriting beyond the usage of autoregressive generative models, guided by the following question:

How can one leverage Masked Language Models (MLMs) to achieve contextually relevant, yet privacy-preserving text rewriting?

To answer this question, we propose DP-MLM, a differentially private text rewriting mechanism leveraging masked token prediction in BERT-based models. We design a rewriting mechanism that incorporates the contextual information from an input sentence to produce a rewritten output that is both utility- and privacy-preserving. This customization is enabled by building a rewriting mechanism on top of an MLM, as opposed to relying on a decoder to generate privatized outputs. To test utility and privacy, we conduct empirical experiments that validate the ability of DP-MLM to rewrite texts in a way that preserves meaning yet still empirically defends against adversarial privacy attacks.

Our work makes the following contributions to the study of differentially private text rewriting:

1. To the best of our knowledge, we are the first

<p>The cat sat on the mat. <mask> cat sat on the mat. A <mask> sat on the mat. A boy <mask> on the mat. A boy sat <mask> the mat. A boy sat across <mask> mat. A boy sat across the <mask>. A boy sat across the bench.</p>	<p>The cat sat on the mat. [SEP] The cat sat on the mat. The cat sat on the mat. [SEP] <mask> cat sat on the mat. The cat sat on the mat. [SEP] The <mask> sat on the mat. The cat sat on the mat. [SEP] The kitten <mask> on the mat. The cat sat on the mat. [SEP] The kitten slept <mask> the mat. The cat sat on the mat. [SEP] The kitten slept on <mask> mat. The cat sat on the mat. [SEP] The kitten slept on the <mask>. The kitten slept on the pillow.</p>
DP-MLM without concatenation	DP-MLM with concatenation

Figure 1: An example of Differentially Private Text Rewriting using Masked Language Models (DP-MLM). The left side shows a real example without contextualization, and the right shows the same example with contextualization. As can be seen, providing a concatenated context sentence (the original sentence) guides the private rewriting process to be more semantically similar than if performed without contextualization.

work to propose the use of BERT-based models for DP text rewriting.

2. We design a rewriting mechanism, DP-MLM, which leverages a contextualization technique not before used for text privatization.
3. With DP-MLM, we surpass previous SOTA methods for private text rewriting on many benchmarks in both utility and privacy, particularly at lower per-token ϵ budgets.

These contributions support our hypothesis that MLMs are effective tools for utility-preserving text privatization, and its observed strengths are analyzed at the conclusion of this work.

2 Related Work

Natural language contains sensitive information (Brown et al., 2022; McMahan et al., 2017). DP enables the training of Machine Learning (ML) models on sensitive texts with a guarantee that the model will not leak more sensitive data than a pre-defined value (Pan et al., 2020; McMahan et al., 2017; Carlini et al., 2021). In the field of NLP, there are two primary notions of integrating DP for downstream applications. The first approach is to collect user texts at a central location and train a model on the texts using a differentially-private optimization technique like DP-SGD (Abadi et al., 2016) (Ponomareva et al., 2022; McMahan et al., 2017; Kerrigan et al., 2020). This approach is known as *global* or *central* DP. In contrast, the second approach is to apply a DP mechanism on texts *locally*, i.e., on the user side, before sharing the privatized texts with a central aggregator. This notion is called Local Differential Privacy (LDP). LDP is a stricter notion of DP as compared to the formerly defined central DP (Feyisetan et al., 2020).

The earliest application of LDP to the task of text privatization considers a sentence as independent sequences of words (Fernandes et al., 2019; Feyisetan et al., 2020). As a result, the privatized sentence is generated by perturbing a sentence word-by-word, normally by introducing calibrated noise to word embeddings (Yue et al., 2021; Chen et al., 2022; Carvalho et al., 2023; Meisenbacher et al., 2024a). These methods do not consider the grammatical and contextual information while generating a private sentence (Mattern et al., 2022). Moreover, they utilize the generalized notion of metric DP, which increases the utility of the generated text but makes comparative evaluations challenging.

Other LDP methods operate at higher levels in the syntactic hierarchy, such as the sentence-level, and directly generate a privatized text, either by adding noise to the latent representations of the text (Igamberdiev and Habernal, 2023) or using additional public information (Meehan et al., 2022). By construction, these methods take into account the grammar and context of the input text while privatizing it. These methods provide stricter privacy guarantees as compared to the aforementioned word-level mechanisms but come at a significant utility cost (Igamberdiev and Habernal, 2023).

Recent works for LDP in NLP take a different approach and leverage Language Models to generate privatized text (Mattern et al., 2022; Utpala et al., 2023). They model the task of text privatization as a text paraphrasing task, or more generally, text rewriting. In leveraging such models to rewrite text, generation is performed by tokens being sampled in a fashion that satisfies local DP. We build upon the foundations of these methods, and aim to improve upon the utility- and privacy-preserving capabilities of such an approach.

3 Foundations

3.1 Masked Language Modeling

A Masked Language Model (MLM), such as BERT (Devlin et al., 2019), includes as one of its two pre-training objectives the task of *masked token prediction*. Here, the model is trained to predict a randomly masked token in a sentence by conditioning its probability not only on the tokens that are to its left in the sentence but also to its right. Thus, the masked token is filled by the most suitable word deemed by the MLM, which is achieved by considering the *complete* context of the sentence.

If the l^{th} token of a sentence s containing n tokens in sequence $w_1 \cdot w_2 \cdots w_n$ is masked, the probability of the masked token being a word v from the vocabulary \mathcal{V} , as modeled by an MLM, is

$$Pr[w_l = v] = Pr[v|w_1 \cdot w_2 \cdots w_{l-1} \cdot w_{l+1} \cdots w_n]$$

Although the masked token of a MLM is primarily used in its pre-training task, one can also use the token to replace a given target word in a text. This fact is leveraged by the related tasks of MLM-based Lexical Substitution (Zhou et al., 2019) or Lexical Simplification (Qiang et al., 2020).

3.2 Differential Privacy

Differential Privacy (DP) (Dwork, 2006) is a mathematically grounded notion of privacy that provides information-theoretic privacy guarantees while performing computation over a dataset. Given $\varepsilon \geq 0$ and finite sets \mathcal{W} and \mathcal{V} , a randomized mechanism $M : \mathcal{W} \rightarrow \mathcal{V}$ is an ε -DP mechanism if $\forall c, c' \in \mathcal{C}$ and $\forall v \in \mathcal{V}$, the following condition holds:

$$\frac{Pr[M(c) = v]}{Pr[M(c') = v]} \leq e^\varepsilon$$

Here c and c' are called *neighboring* or *adjacent* databases. Depending on the notion of adjacency, the unit which is protected by LDP is defined. In our case, any two context sentences c and c' are adjacent. This is expounded upon in the next section.

3.2.1 Temperature Sampling as an Exponential Mechanism

Suppose there is a dataset $D \in \mathcal{X}^n$ and our aim is to derive a value for this dataset from a set of fixed value choices $\mathcal{V} \in \mathcal{Y}$. Exponential Mechanism can be used here to select, in a private manner, the best choice for a dataset from a set of choices \mathcal{V} , with its goodness being determined by a scoring

function. The scoring (utility) function $u : \mathcal{X}^n \times \mathcal{Y} \rightarrow \mathbb{R}$ maps database and choice pairs, (D, w) with $D \in \mathcal{X}^n$ and $w \in \mathcal{V}$, to possibility scores. The l_2 -sensitivity of such scoring function is given as

$$\Delta u = \max_{w \in \mathcal{V}} \max_{D, D' \in \mathcal{X}^n} |u(D, w) - u(D', w)|$$

If the choice for D is selected according to probability proportional to $\exp(\frac{\varepsilon u(D, w)}{2\Delta u})$, then the selection algorithm satisfies DP and this DP mechanism is termed as an Exponential Mechanism (McSherry and Talwar, 2007).

Mapping it to our use case, suppose we have a context $s = w_1 \cdot w_2 \cdot w_3 \cdots w_n$ or $s = w_1 \cdot w_2 \cdots w_{l-1} \cdot w_{l+1} \cdots w_n$, and we want to choose the best token for the masked word w_l from the set of vocabulary. This selection can be done privately through an Exponential Mechanism with scoring function $u : \mathcal{V}^* \times \mathcal{V} \rightarrow \mathbb{R}$ that takes as input the whole context s and a word w from the vocabulary set \mathcal{V} and outputs a score for w conditioned on the context s . This scoring function is precisely a (Masked) Language Model in our case: it takes the context as input and outputs logits for every word in the vocabulary. To bound the sensitivity of the scoring function, which is necessary for DP, the logit values can be clipped to a predefined range.

If the logits generated by the MLM are bounded, then the Exponential Mechanism is realized by default if we use temperature sampling to predict the masked token (Mattern et al., 2022). To compare it, if the logit value produced for a word $w \in \mathcal{V}$ with a context s is u and the sampling temperature is set to T , the predicted token being w has the probability proportional to $\exp(u/T)$. Comparing it to the exponential mechanism, we can derive that $\varepsilon = \frac{2\Delta u}{T}$. A detailed DP proof is provided in Section 4.3.

4 Method

In this section, we outline the design of DP-MLM, particularly its underlying DP mechanism, as well as the text rewriting mechanism.

4.1 DP Masked Token Prediction

Given an input sentence s with n tokens, our goal is to privatize each token in the sentence, one token at a time. To privatize a single token w_l , we input the entire sentence $s = w_1 \cdots w_{l-1} \cdot w_l \cdot w_{l+1} \cdots w_n$ to an MLM, with w_l replaced by the model's mask token, usually <mask>. Next, we capture the logit

values for the masked token index l , clip them (described in Section 5.1.1), and apply Temperature Sampling mechanism, that is equivalent to the Exponential Mechanism as described in Section 3.2.1. All clipped logit values are first divided by the temperature T , which is calculated according to $T = \frac{2\Delta u}{\epsilon}$. The resulting values are fed through a softmax function and a token is sampled according to these probabilities. The sampled token w_p then serves as the differentially private token replacement for w_l . A proof that this mechanism is DP is found in Section 4.3.

4.2 Rewriting Mechanism

To rewrite an entire text with the underlying mechanism described above, we design a rewriting mechanism DP-MLM that is outlined in Algorithms 1 and 2. Note that for the entirety of this work, we use the ROBERTA-BASE MLM as our base model.

Algorithm 1

DP-MLM Token Replacement

Require: MLM M ,
context tokens, private p_tokens , position idx
epsilon ϵ ,
clipping values $C = (C_{min}, C_{max})$
Ensure: Output $private_token$ at position idx

```

 $T \leftarrow 2 \cdot (|C_{max} - C_{min}|) / \epsilon$ 
 $p\_tokens[idx] \leftarrow \langle \text{mask} \rangle$ 
 $masked\_sent \leftarrow \text{concat}(\text{tokens}, [\text{SEP}], p\_tokens)$ 
 $logits \leftarrow M(\text{masked\_sent})$ 
 $logits \leftarrow \text{clip\_and\_temp}(logits, C, T)$ 
 $prob \leftarrow \text{softmax}(logits)$ 
 $private\_token \leftarrow \text{sample}(prob)$ 
return  $private\_token$ 

```

Algorithm 2

Text Rewriting using DP-MLM

Require: input sentence $s = w_1 \cdot w_2 \cdots w_n$,
per-token epsilon ϵ ,
clipping values $C = (C_{min}, C_{max})$
Ensure: rewritten (privatized) sentence using DP-MLM

```

 $tokens \leftarrow \text{tokenize}(s)$ 
 $private \leftarrow tokens$ 
for  $i \in 1..n$  do
   $p \leftarrow DPMLM(tokens, private, i, \epsilon, C)$   $\triangleright$  Alg. 1
   $private[i] \leftarrow p$ 
end for
return  $\text{detokenize}(private)$ 

```

As described in Algorithm 1, contextualization of the DP-MLM mechanism is achieved via the concatenation of the original input sentence, which is given along with the masked sentence as input to the MLM. This simple trick is motivated by a simi-

lar approach followed by (Qiang et al., 2020) for Lexical Simplification. A more intuitive illustrative example of the DP-MLM rewriting mechanism is found in Figure 1. Note that in our implementation, we do not replace English stopwords, but we leave this as a parameter in our open-source code.

With Algorithm 1, we replace a single token from an input text in a DP manner. Thus, the output of one DP-MLM usage is a *privatized* token that is contextually relevant to the text. By using DP-MLM for each token in the input sentence (Algorithm 2), we design a text rewriting mechanism that leverages the compositionality of DP to output a privately rewritten text with a DP guarantee of $\epsilon \times \text{len}(\text{tokens})$. This is formalized for the token and text level in the following.

4.3 Privacy Guarantees

Our mechanism \mathcal{M} introduced in Algorithm 1 satisfies local differential privacy (LDP). For any two adjacent context sentences, \mathcal{M} yields “similar” tokens to fill a given masked token with LDP guarantees. Hence, given a predicted masked token, an adversary who only sees the predicted token cannot differentiate with high certainty if the predicted token was due to a context sentence c or c' . This results in plausible deniability about the source of the predicted masked token.

Suppose a sentence s consists of n tokens, i.e., $s = w_1 \cdot w_2 \cdots w_n$, and the (masked) token that we want to predict lies at l^{th} position in the sentence. An MLM models the likelihood of masked token w_l being $v_i \in \mathcal{V}$ as follows:

$$\begin{aligned} Pr[w_l = v_i] &= Pr[v_i | w_1 \cdots w_{l-1} \cdot w_{l+1} \cdots w_n] \\ &= Pr[v_i | C_l] \end{aligned}$$

As stated earlier, we use clipped logits u and temperature sampling to sample the most likely token for the masked token. Hence, for our proposed mechanism $M : \mathcal{V}^+ \rightarrow \mathcal{V}$ that takes the context sentence $C_l = w_1 \cdot w_2 \cdots w_{l-1} \cdot w_{l+1} \cdots w_n$ as input and returns a privately selected predicted masked word, its output probability distribution is given by the following equation:

$$Pr[M(C_l) = v_i] = \frac{\exp(\frac{u(C_l, v_i)}{T})}{\sum_{j=1}^{|\mathcal{V}|} \exp(\frac{u(C_l, v_j)}{T})} \quad (1)$$

Theorem 1. *The proposed mechanism M defined in the equation 1 satisfies ϵ -LDP.*

Proof. Let $s, s' \in \mathcal{V}^n$. Suppose the l^{th} token of s, s' is masked and we use our mechanism M to predict a token. The context to be set as input for the Masked Language Model becomes:

$$C_l := w_1 \cdot w_2 \cdots w_{l-1} \cdot w_{l+1} \cdots w_n$$

$$C'_l := w'_1 \cdot w'_2 \cdots w'_{l-1} \cdot w'_{l+1} \cdots w'_n$$

The ratio of the probability distribution of application of M on C_l and C'_l can be given as:

$$\begin{aligned} \frac{Pr[M(C_l) = v_i]}{Pr[M(C'_l) = v_i]} &= \frac{\exp(\frac{u(C_l, v_i)}{T})}{\sum_{j=1}^{|\mathcal{V}|} \exp(\frac{u(C_l, v_j)}{T})} \cdot \frac{\sum_{j=1}^{|\mathcal{V}|} \exp(\frac{u(C'_l, v_j)}{T})}{\exp(\frac{u(C'_l, v_i)}{T})} \\ &= \frac{\exp(\frac{u(C_l, v_i)}{T})}{\exp(\frac{u(C'_l, v_i)}{T})} \cdot \frac{\sum_{j=1}^{|\mathcal{V}|} \exp(\frac{u(C'_l, v_j)}{T})}{\sum_{j=1}^{|\mathcal{V}|} \exp(\frac{u(C_l, v_j)}{T})} \end{aligned}$$

Solving the first fraction, we get

$$\begin{aligned} \frac{\exp(\frac{u(C_l, v_i)}{T})}{\exp(\frac{u(C'_l, v_i)}{T})} &= \exp\left(\frac{u(C_l, v_i) - u(C'_l, v_i)}{T}\right) \\ &\leq \exp\left(\frac{\Delta u}{T}\right) \end{aligned}$$

Similarly, solving the second fraction, we get

$$\begin{aligned} \frac{Pr[M(C_l) = v_i]}{Pr[M(C'_l) = v_i]} &\leq \exp\left(\frac{\Delta u}{T}\right) \exp\left(\frac{\Delta u}{T}\right) \\ &= \exp\left(2\frac{\Delta u}{T}\right) \\ &= \exp(\varepsilon) \end{aligned}$$

□

Hence, ε can be calculated from the sensitivity of Δu and the sampling temperature T as $\varepsilon = 2\frac{\Delta u}{T}$

Extending guarantees to a sentence The privacy budget required for generating a single masked token using the mechanism M is equal to ε . The privatized sentence of the input sentence is generated by sequentially generating privatized token for each token present in the input sentence. Thus, for rewriting a sentence (or more generally, a text) of n tokens, we are required to call the mechanism n times. Hence, by sequential composition, the total privacy budget spent for rewriting the entire text would be $n\varepsilon$ -DP.

5 Experimental Setup

In order to evaluate the performance of our proposed method, we design two overarching experiments: (1) utility experiments, and (2) empirical privacy experiments. These are outlined below.

5.1 Utility Experiments

Our utility experiments consist of two phases: (1) utility benchmarking of DP-MLM, and (2) comparative utility testing with two other state-of-the-art DP rewriting mechanisms.

5.1.1 Utility Benchmarking

To test the utility-preserving capability of DP-MLM, we measure the utility of rewritten data across a range of ε values. For this, we utilize the GLUE benchmark (Wang et al., 2018), which contains 9 separate tasks spanning classification, textual similarity, and textual entailment. For a given ε value, a perturbed dataset (train and validation split) is measured against the non-privatized baseline, in order to measure how well the privatization keeps utility intact.

Model For both the non-privatized baseline and all privatized (rewritten) datasets, we fine-tune a DEBERTA-V3-BASE model (He et al., 2021) for one epoch. The trained model is then evaluated on the validation set (non-privatized or privatized, respectively). All models in this work are trained on a single RTX A6000 GPU, using the Adam optimizer (Kingma and Ba, 2017) and all default hyperparameters of the Hugging Face Transformers Trainer API. To account for variations in training, all utility results represent the average of 3 runs.

Clipping Values To ensure that the sensitivity of the logit values of our underlying MLM (ROBERTA-BASE) is bounded, we pre-define the clip value based upon an empirical estimation of the logit value range. Concretely, we measure all logit values from inputting 1000 random text examples from the SST2 dataset of GLUE, calculate the mean μ and standard deviation σ , and define the clipping values as $(C_{min}, C_{max}) = (\mu, \mu + 4\sigma)$. The choice of such a range provides the benefit of a bound sensitivity, while still preserving higher logit values and “clamping” lower values.

Epsilon In order to conduct comprehensive tests on varying privacy budgets, we choose the following set of ε values: $\varepsilon \in \{10, 25, 50, 100, 250\}$.

5.1.2 Comparative Utility

The next utility experiments test DP-MLM against the two state-of-the-art text rewriting methods:

1. **DP-Paraphrase** (Mattern et al., 2022): DP text rewriting utilizing DP temperature sampling in a fine-tuned GPT-2 paraphrasing

Task	CoLA	SST2	QQP	MRPC	STSB	MNLI	QNLI	WNLI	RTE	
Baseline	84.72 _{0.35}	95.72 _{0.22}	89.26 _{0.05}	84.07 _{0.40}	84.57 _{0.78}	88.75 _{0.03}	93.51 _{0.12}	56.34 _{0.00}	54.99 _{2.98}	
ϵ	10	69.13 _{0.00}	68.50 _{0.64}	71.86 _{0.26}	71.32 _{2.31}	6.13 _{0.90}	52.86 _{0.71}	66.25 _{1.11}	53.05 _{4.65}	51.50 _{1.70}
	25	69.77 _{0.12}	76.49 _{0.94}	74.17 _{0.17}	70.67 _{1.80}	12.42 _{2.45}	56.15 _{2.95}	68.54 _{5.63}	55.40 _{1.33}	51.87 _{1.19}
	50	70.85 _{0.86}	84.10 _{0.42}	80.53 _{0.19}	75.25 _{1.31}	26.21 _{11.60}	66.99 _{0.18}	82.01 _{0.03}	52.58 _{5.31}	52.47 _{0.34}
	100	70.05 _{0.86}	86.16 _{0.39}	82.17 _{0.26}	74.35 _{1.63}	34.68 _{4.20}	69.57 _{0.11}	83.56 _{0.07}	48.36 _{5.67}	53.31 _{0.61}
	250	70.18 _{0.75}	86.05 _{0.55}	82.44 _{0.03}	76.39 _{0.92}	60.77 _{1.49}	70.96 _{0.19}	84.68 _{0.20}	51.64 _{5.67}	51.38 _{1.87}

Table 1: Utility Benchmark Scores for DP-MLM. All scores represent accuracy scores, except for STSB, which is represented by the Pearson-Spearman Correlation score. The metrics are an average of three training runs, and the standard deviation is presented as a subscript. In all cases, a higher score is better.

Task	ϵ	Baseline	DP-MLM			DP-Paraphrase			DP-Prompt		
			BLEU	CS	Acc.	BLEU	CS	Acc.	BLEU	CS	Acc.
CoLA	10	84.72 _{0.35}	0.08	0.16	69.13 _{0.00}	0.00	0.26	69.13 _{0.00}	0.00	0.04	69.13 _{0.00}
	25		0.13	0.35	69.77 _{0.12}	0.00	0.26	69.13 _{0.00}	0.00	0.09	69.13 _{0.00}
	50		0.25	0.64	70.85 _{0.86}	0.00	0.28	69.13 _{0.00}	0.29	0.71	67.59 _{1.37}
	100		0.28	0.69	70.05 _{0.86}	0.01	0.33	69.13 _{0.00}	0.68	0.93	73.63 _{0.49}
	250		0.29	0.70	70.18 _{0.75}	0.07	0.43	69.20 _{0.00}	0.76	0.95	74.50 _{0.23}
SST2	10	95.72 _{0.22}	0.08	0.17	68.50 _{0.64}	0.00	0.26	58.60 _{5.44}	0.00	0.07	50.92 _{0.00}
	25		0.11	0.37	76.49 _{0.94}	0.00	0.26	63.88 _{0.19}	0.00	0.12	52.52 _{2.43}
	50		0.19	0.61	84.10 _{0.42}	0.00	0.27	61.66 _{0.57}	0.02	0.36	70.84 _{0.43}
	100		0.22	0.65	86.16 _{0.39}	0.00	0.29	64.26 _{0.30}	0.12	0.62	83.72 _{0.16}
	250		0.23	0.67	86.05 _{0.55}	0.05	0.37	67.81 _{0.33}	0.15	0.65	85.17 _{0.61}
MRPC	10	84.07 _{0.40}	0.05	0.12	71.32 _{2.31}	0.00	0.26	68.38 _{0.00}	0.00	0.04	68.38 _{0.00}
	25		0.08	0.37	70.67 _{1.80}	0.00	0.27	68.71 _{0.46}	0.00	0.09	68.55 _{0.00}
	50		0.17	0.61	75.25 _{1.31}	0.00	0.28	68.38 _{0.00}	0.05	0.48	71.10 _{0.12}
	100		0.19	0.66	74.35 _{1.63}	0.00	0.30	68.71 _{0.46}	0.29	0.78	71.16 _{0.83}
	250		0.21	0.68	76.39 _{0.92}	0.05	0.38	68.95 _{0.81}	0.37	0.82	71.24 _{0.76}
RTE	10	54.99 _{2.98}	0.04	0.10	51.50 _{1.70}	0.00	0.27	53.79 _{1.53}	0.00	0.05	51.14 _{2.74}
	25		0.08	0.33	51.87 _{1.19}	0.00	0.28	53.19 _{0.68}	0.00	0.09	50.42 _{1.62}
	50		0.17	0.61	52.47 _{0.34}	0.00	0.29	50.78 _{2.72}	0.27	0.71	49.82 _{2.23}
	100		0.20	0.65	53.31 _{0.61}	0.01	0.33	52.71 _{0.00}	0.63	0.92	56.32 _{5.11}
	250		0.21	0.66	51.38 _{1.87}	0.08	0.43	49.10 _{2.55}	0.68	0.94	59.21 _{4.59}

Table 2: Comparative Utility on a Subset of GLUE Tasks. Scores in **bold** mark the highest score achieved by per (task, ϵ) pair. In 14 out of the 20 settings, DP-MLM achieves the highest accuracy, including ties.

model. More details on the replication of DP-PARAPHRASE is provided in Appendix B.

2. **DP-Prompt** (Utpala et al., 2023): DP text rewriting using prompting of large language models, i.e., to paraphrase the input text. For the purposes of this work, we utilize a pre-trained FLAN-T5-BASE.

These methods are tested on a subset of the GLUE tasks, namely {CoLA, MRPC, RTE, SST2}, where these datasets are perturbed accordingly as for DP-MLM. The same ϵ values and clipping strategy as described in Section 5.1.1 are used for both methods to ensure direct comparability. In addition, for both generative methods listed above, we restrict the maximum number of generated tokens to the length of the input tokens, so as to ensure fairness in comparative evaluation metrics.

Scoring For comparative utility testing, we report raw scores (accuracy or correlation), as well as BLEU (Papineni et al., 2002; Igamberdiev and Habernal, 2023) and cosine similarity (CS) scores between the original and privatized dataset (Meisenbacher et al., 2024b). These additional metrics present a clearer picture of the utility-preservation of the underlying DP mechanism, compared against the amount by which the dataset has been perturbed. For semantic similarity (CS), we utilize the ALL-MINILM-L6-V2 model (Reimers and Gurevych, 2019).

5.2 Empirical Privacy Experiments

To measure the privacy-preserving capabilities of DP-MLM in comparison to the state-of-the-art, we conduct empirical privacy experiments on two datasets. This process is described in the following.

Trustpilot	Baseline	DP-MLM					DP-PARAPHRASE					DP-PROMPT				
		∞	10	25	50	100	250	10	25	50	100	250	10	25	50	100
Utility F1 \uparrow	99.62	97.30	97.98	98.89	99.11	99.24	96.43	96.05	96.72	96.57	96.16	96.69	96.72	97.90	99.30	99.38
PP+ \uparrow	-	+62	+130	+221	+243	+257	-25	-63	+5	-11	-52	+2	+4	+123	+263	+270
CS	-	0.10	0.19	0.24	0.25	0.25	0.15	0.15	0.15	0.15	0.15	0.11	0.13	0.20	0.24	0.25
Privacy F1 (stat.) \downarrow	69.60	59.74	61.26	68.58	70.30	70.92	60.02	60.23	59.85	60.27	60.52	59.7	59.64	68.75	78.83	81.67
Privacy F1 (adap.) \downarrow	69.60	58.50 _{0.6}	61.93 _{1.4}	66.73 _{0.4}	65.10 _{5.0}	66.33 _{0.8}	60.17 _{1.6}	58.33 _{0.3}	58.97 _{2.3}	60.16 _{0.7}	58.03 _{0.5}	58.10 _{0.0}	57.23 _{1.2}	60.30 _{0.8}	66.53 _{0.8}	69.80 _{0.7}
Relative Gain (stat.) \uparrow		0.12	0.10	0.01	-0.02	-0.02	0.11	0.10	0.11	0.10	0.10	0.11	0.11	-0.01	-0.14	-0.18
Relative Gain (adap.) \uparrow		0.14	0.09	0.03	0.06	0.04	0.10	0.13	0.13	0.11	0.13	0.14	0.15	0.11	0.04	-0.01

Yelp	Baseline	DP-MLM					DP-PARAPHRASE					DP-PROMPT				
		∞	10	25	50	100	250	10	25	50	100	250	10	25	50	100
Utility F1 \uparrow	97.50	95.51	95.51	96.64	96.22	96.47	95.51	95.51	95.51	95.51	96.05	95.51	95.51	95.51	95.76	96.34
PP+ \uparrow	-	-84	-84	+29	-13	+11	-84	-84	-84	-84	-30	-84	-84	-84	-59	-1
CS	-	0.15	0.50	0.76	0.80	0.81	0.34	0.35	0.36	0.37	0.36	0.12	0.15	0.49	0.72	0.76
Privacy F1 (stat.) \downarrow	87.20	13.72	30.92	47.24	49.96	50.76	11.60	11.72	12.04	12.44	12.61	10.16	10.52	25.32	53.60	62.48
Privacy F1 (adap.) \downarrow	87.20	62.40 _{0.3}	61.07 _{3.5}	73.87 _{1.2}	71.87 _{0.4}	74.13 _{1.1}	20.27 _{1.5}	23.07 _{2.0}	22.27 _{1.4}	20.53 _{1.0}	25.00 _{1.2}	18.67 _{1.6}	13.33 _{2.3}	24.53 _{1.3}	48.80 _{1.8}	56.53 _{1.5}
Relative Gain (stat.) \uparrow		0.82	0.63	0.45	0.41	0.41	0.85	0.85	0.84	0.84	0.84	0.86	0.86	0.69	0.37	0.27
Relative Gain (adap.) \uparrow		0.26	0.28	0.14	0.16	0.14	0.74	0.72	0.72	0.74	0.69	0.77	0.83	0.70	0.42	0.34

Table 3: Empirical Privacy Results for Trustpilot (top) and Yelp (bottom). *Utility F1* for the sentiment classification task is given, as well as the adversarial performance (*Privacy F1*) for both the static (*stat.*) and adaptive (*adap.*) settings. *PP+* denotes percentage points above majority-class guessing, *CS* denotes cosine similarity between original and privatized datasets, and *Relative Gain* quantifies the observed benefit of privacy vs. utility (Section 5.2).




Figure 2: Average Utility Loss. This graph depicts the average utility loss for a given ϵ value across four GLUE tasks. On average, DP-MLM leads to a lower utility loss than DP-PARAPHRASE or DP-PROMPT.

Datasets For conducting our privacy experiments, we utilize two datasets:

1. **Trustpilot Reviews** (Hovy et al., 2015): a large corpus of user reviews from Trustpilot, containing both a review score (1-5) and the gender (M/F) of the reviewer. We only consider reviews rated with 5 (positive) or 1-2 (negative). From this, we take a random 10% sample, representing $\sim 36k$ reviews.
2. **Yelp Reviews**: we utilize the dataset as used by Utpala et al. (2023), which contains Yelp reviews from 10 authors, labeled as positive or negative. We take a random sample of 250 texts from each author, for a total of 2500.

Tasks For both datasets, we conduct a two-sided experiment. As in the utility experiment, we privatize each dataset with the budgets $\epsilon \in \{10, 25, 50, 100, 250\}$, and compare the utility loss

against the non-private baseline. A DEBERTA-V3-BASE model is once again employed for fine-tuning, trained for a total of three epochs.

Following the approach laid out by previous works (Mattern et al., 2022; Utpala et al., 2023), we test empirical privacy in two adversarial settings. The first is called the *static* attacker, where the adversarial model can only be evaluated on the privatized outputs after being trained on the original non-privatized input texts. In contrast, the *adaptive* attacker is able to train the adversarial model on the DP outputs, thus more closely matching the distribution of the target privatized texts.

For the static setting, we train an adversarial DEBERTA-V3-BASE model on the non-privatized dataset to predict the protected attribute of each dataset, i.e., gender for Trustpilot and author ID for Yelp. These models are trained for five epochs. Then, we evaluate the adversarial model on each privatized dataset and measure the change in performance. In this way, we can empirically measure the privacy protection provided by rewriting a dataset.

In the adaptive setting, the only difference is that the DEBERTA-V3-BASE model is *trained on the privatized training datasets*, and subsequently evaluated on the validation splits of these datasets.

Scoring Following Mattern et al. (2022), we not only report the raw results of the above-mentioned experiments, but also the *relative gain* achieved by text privatization via rewriting. Such a score is useful in quantifying the advantage of privatizing text, or rather the gain in privacy offset by the inevitable loss in privacy. Let P_o , U_o represent the baseline privacy and utility scores, respectively, that is the scores when training both the sentiment and ad-

versarial classifiers on the non-privatized datasets. Let P_r, U_r be the scores observed on the privatized (rewritten) datasets. The relative gain is thus defined as $RG = (U_r/U_o) - (P_r/P_o)$. Note that we report relative gains for both adversarial settings.

In addition to the raw utility scores (F1 score), we also present the percentage points ($PP+$) achieved above majority class guessing. In the case of both Trustpilot and Yelp, the datasets are highly biased towards positive reviews, so reporting $PP+$ better demonstrates the degree to which a fine-tuned model learns to distinguish sentiment.

6 Results

Utility Table 1 presents the complete benchmarking results of DP-MLM on the GLUE tasks. Table 2 displays the comparative utility results, which tested DP-MLM and our two selected methods on a subset of the GLUE tasks. The results from Table 2 are summarized in Figure 2, which illustrates the average utility loss for each evaluated rewriting mechanism, given an ϵ value.

Privacy Table 3 displays the results of our empirical privacy tests, for both Trustpilot and Yelp.

Efficiency We measure the *efficiency*, or speed, at which the evaluated mechanisms can rewrite text. We capture this by recording the amount of time taken to perturb the selected GLUE datasets, including how many tokens are perturbed (1,048,231 in total). The results are summarized below:

- **DP-MLM**
 - Elapsed time: 1316 minutes
 - Tokens/min: 797
- **DP-Paraphrase**
 - Elapsed time: 1308 minutes
 - Tokens/min: 802
- **DP-Prompt**
 - Elapsed time: 1961 minutes
 - Tokens/min: 535

7 Discussion

Utility In analyzing the results of both utility evaluations, one can see that DP-MLM clearly demonstrates the ability to produce utility-preserving DP rewritten text. As expected, this naturally comes with a performance drop as compared to the non-privatized baseline; however, interesting findings can be extracted when observing the comparative utility tests. Across four selected GLUE tasks, which represent all three task

types of the benchmark, DP-MLM consistently outperforms the other two state-of-the-art methods, DP-PARAPHRASE (Mattern et al., 2022) and DP-PROMPT (Utpala et al., 2023). This is supported by the fact that DP-MLM achieves the highest utility score in 14 out of the 20 comparative settings.

A closer look at Table 2 shows that in all four cases where DP-PROMPT outperforms DP-MLM, this can be attributed to a very high BLEU and CS score between the original and privatized datasets. While this may be useful for utility preservation, one may question whether CS scores near to 1 provide much privacy preservation at all. On the other hand, even at higher values of ϵ such as 100 or 250, DP-MLM still provides higher levels of privatization (more rewritten), while still offering competitive utility scores in all cases.

The above is especially supported by the fact that in 8 out of the 14 cases where DP-MLM achieves the highest utility score, it does so without having the highest CS to the original dataset. In addition, DP-MLM shows particular strength at low ϵ budgets, such as with $\epsilon = 10$, where it never possesses the highest CS score, yet still remains very competitive. These observations emphasize the ability of DP-MLM to produce privatized, yet still contextually and semantically relevant rewritten texts. This is further supported by Figure 2, which places DP-MLM, on average, lower than DP-PARAPHRASE and DP-PROMPT in terms of utility loss.

Privacy From the empirical privacy results, one can observe similar trends as with the utility experiments. Empirically, all three evaluated methods are very effective in reducing the adversarial advantage, i.e., gender classification or author identification, and this is especially true at lower privacy budgets. DP-PARAPHRASE is particularly effective, but this comes at the cost of comparatively poor results in preserving utility, as can be seen in Table 3.

Comparing DP-MLM and DP-PROMPT in terms of empirical privacy, one can observe from the Trustpilot test that both methods are successful in preserving utility to a degree where a model can still learn sentiment classification reasonably, as shown by $PP+$, or the F1 percentage points above majority-class guessing. This case for DP-MLM is especially made salient in the Yelp test, where DP-MLM is the only method capable of producing positive $PP+$ scores. Furthermore, despite achieving similar or better utility scores in the empirical privacy tests as compared to DP-PROMPT, DP-MLM

consistently scores better in reducing adversarial F1 at higher ϵ values for the static setting, while the opposite is true for the adaptive setting.

Relative Gain also provides insights, where DP-MLM offers added value in most cases and better trade-offs at higher ϵ budgets than DP-PROMPT, particularly in the static setting. A weakness of DP-MLM is highlighted by the adaptive results, pointing to the inevitable trade-off between higher utility text and its resulting privacy protections. Nevertheless, DP-MLM still achieves positive gains in all adaptive scenarios. All methods are similarly competitive at lower ϵ , yet this must be interpreted according to whether privacy or utility is favored.

Efficiency DP-MLM performs nearly identically in terms of efficiency as opposed to DP-PARAPHRASE, and both methods greatly outperform DP-PROMPT. This can directly be attributed to the encoder-only ROBERTA-BASE and the decoder-only GPT-2, as opposed to the utilized FLAN-T5-BASE, which is nearly double in model size. Especially when considering the above-mentioned strengths of DP-MLM, the added competitiveness of speed introduces a practical advantage of our method. Such results also raise interesting points for future work regarding the effect of model size and architecture on privatization, particularly the interplay between these and ϵ .

Addressing Limitations The discussion of the merits of DP-MLM must also be met with its remaining limitations. As our rewriting mechanism leveraging DP-MLM relies on token-level DP replacements, the primary limitation comes with the initial inability to rewrite sentences with differing lengths from the original texts. To address this main limitation, we propose an improved version of the rewriting mechanism which enables variable length outputs. This is presented in Algorithm 3.

In essence, Algorithm 3 takes as inputs an addition probability A and deletion probability D , and for each token in the input text s , we delete this token with probability D and add an additional token with probability A . New tokens are added by simply inserting a mask token into the context sentence and running DP-MLM as usual.

From a privacy guarantee standpoint, the downside of such an augmentation comes with an altered guarantee. In the worst case for a given ϵ , text rewriting with Algorithm 3 offers a privacy guarantee of $2n\epsilon$ -DP, i.e., in the case that a token is added for every token in the input sentence. In the

Algorithm 3

Text Rewriting +- using DP-MLM

Require: input sentence $s = w_1 \cdot w_2 \cdots w_n$,
per-token epsilon ϵ ,
clipping values $C = (C_{min}, C_{max})$,
token addition probability A , token deletion probability D
Ensure: rewritten (privatized) sentence using DP-MLM

```

tokens  $\leftarrow$  tokenize( $s$ )
private  $\leftarrow$  tokens
added  $\leftarrow$  0
deleted  $\leftarrow$  0
for  $i \in 1 \dots n$  do
  prob_del, prob_add  $\leftarrow$  rand()  $\triangleright$  random numbers  $\in [0..1]$ 
  if prob_del  $\geq D$  then
     $p \leftarrow DPMLM(\text{tokens}, \text{private}, i + \text{added} - \text{deleted}, \epsilon, C)$ 
    private[ $i + \text{added} - \text{deleted}$ ]  $\leftarrow p$ 
  else
    deleted  $\leftarrow$  deleted + 1
  end if
  if prob_add  $\leq A$  then
    added  $\leftarrow$  added + 1
     $p \leftarrow DPMLM(\text{tokens}, \text{private}, i + \text{added} - \text{deleted}, \epsilon, C)$ 
    private.insert( $i + \text{added} - \text{deleted}, p$ )
  end if
end for
return detokenize(private)

```

average case, though, we achieve $(A - D)n\epsilon$ -DP. In Appendix C, we show the results of repeating a subset of our empirical privacy experiments using this augmented rewriting method.

8 Conclusion

We present DP-MLM, a differentially private text rewriting mechanism leveraging masked token prediction of MLMs. As opposed to previous methods relying on private generation using language models with decoders, we utilize BERT-based encoder-only models to rewrite text in a private, yet contextual manner. This is accomplished by simple concatenation in our rewriting mechanism, which guides the private sampling of replacement tokens.

In a series of utility and privacy experiments, we empirically demonstrate the improvements achieved by DP-MLM over previous SOTA methods. In particular, DP-MLM outperforms these methods on many utility benchmarks, and achieves competitive empirical privacy scores. An analysis of the results reveals that DP-MLM finds a necessary balance between utility- and privacy-preservation, more so than the previous SOTA.

As paths for future work, we propose continued research on DP text rewriting with encoder-only models, including the investigation of how privatization can be improved with less utility loss and the effect of using different base models. In addition, we see that more work to define and unify evaluation strategies for DP text rewriting should be undertaken, so as to allow for well-defined methodologies for validating future proposed mechanisms.

Acknowledgements

The authors would like to thank the anonymous reviewers for their feedback and Alexandra Klymenko for her valuable contributions to this work.

Limitations

The primary limitation of our work comes with the choice of underlying language model for each DP mechanism. We choose one model as the representative model for DP-MLM, DP-PARAPHRASE, and DP-PROMPT, namely ROBERTA-BASE, GPT-2, and FLAN-T5-BASE, respectively. The implications of choosing and evaluating other BERT-based models were considered outside of the scope of this work. This should be tested in follow-up studies.

Another limitation that pertains to general evaluation of privacy-preserving mechanisms, that is, the evaluation of *privacy* protection in NLP. The method of empirical privacy employed in this work follows from the predominant method in the literature to measure privacy, but as it is a proxy for privacy, we can only claim that our proposed DP-MLM protects privacy *empirically* and *by proxy*.

Ethics Statement

An ethical consideration involves our empirical privacy experiments, which leverage existing datasets not originally intended for adversarial gender or author identification. In performing these empirical experiments, the actions of a potential adversary were simulated, i.e., to leverage publicly accessible information for the creation of an adversarial model. As these datasets are already publicly available, no harm was inflicted in the privacy experiments as part of this work. Moreover, the Yelp dataset is made up of pseudonyms (User IDs) rather than PII, thus further reducing the potential for harm.

References

Martín Abadi, Andy Chu, Ian J. Goodfellow, H. B. McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. [Deep learning with differential privacy](#). *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*.

Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.

Hannah Brown, Katherine Lee, FatemehSadat Mireshghallah, R. Shokri, and Florian Tramèr. 2022. [What does it mean for a language model to preserve privacy?](#) *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*.

Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. 2021. [Extracting training data from large language models](#). In *30th USENIX Security Symposium (USENIX Security 21)*, pages 2633–2650.

Ricardo Silva Carvalho, Theodore Vasiloudis, Oluwaseyi Feyisetan, and Ke Wang. 2023. [TEM: High utility metric differential privacy on text](#). In *Proceedings of the 2023 SIAM International Conference on Data Mining (SDM)*, pages 883–890. SIAM.

Hui Chen, Fengran Mo, Yanhao Wang, Cen Chen, Jianyun Nie, Chengyu Wang, and Jamie Cui. 2022. [A customized text sanitization mechanism with differential privacy](#). In *Annual Meeting of the Association for Computational Linguistics*.

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, Minneapolis, Minnesota. Association for Computational Linguistics.

Cynthia Dwork. 2006. [Differential privacy](#). In *International colloquium on automata, languages, and programming*, pages 1–12. Springer.

Natasha Fernandes, Mark Dras, and Annabelle McIver. 2019. [Generalised differential privacy for text document processing](#). In *Principles of Security and Trust: 8th International Conference, POST 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6–11, 2019, Proceedings 8*, pages 123–148. Springer International Publishing.

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, WSDM '20*, page 178–186, New York, NY, USA. Association for Computing Machinery.

Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. [Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing](#).

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, WWW '15*, page 452–461, Republic and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee.
- Timour Igamberdiev and Ivan Habernal. 2023. **DP-BART for privatized text rewriting under local differential privacy**. In *Findings of the Association for Computational Linguistics: ACL 2023*, page (to appear), Toronto, Canada. Association for Computational Linguistics.
- Gavin Kerrigan, Dylan Slack, and Jens Tuyls. 2020. **Differentially private language models benefit from public pre-training**. *ArXiv*, abs/2009.05886.
- Diederik P. Kingma and Jimmy Ba. 2017. **Adam: A method for stochastic optimization**.
- Oleksandra Klymenko, Stephen Meisenbacher, and Florian Matthes. 2022. **Differential privacy in natural language processing the story so far**. In *Proceedings of the Fourth Workshop on Privacy in Natural Language Processing*, pages 1–11, Seattle, United States. Association for Computational Linguistics.
- Justus Mattern, Benjamin Weggenmann, and Florian Kerschbaum. 2022. **The limits of word level differential privacy**. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 867–881, Seattle, United States. Association for Computational Linguistics.
- H. B. McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. 2017. **Learning differentially private language models without losing accuracy**. *ArXiv*, abs/1710.06963.
- Frank McSherry and Kunal Talwar. 2007. **Mechanism design via differential privacy**. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 94–103.
- Casey Meehan, Khalil Mrini, and Kamalika Chaudhuri. 2022. **Sentence-level privacy for document embeddings**. In *Annual Meeting of the Association for Computational Linguistics*.
- Stephen Meisenbacher, Maulik Chevli, and Florian Matthes. 2024a. **1-Diffractor: Efficient and utility-preserving text obfuscation leveraging word-level metric differential privacy**. *arXiv preprint arXiv:2405.01678*.
- Stephen Meisenbacher, Nihildev Nandakumar, Alexandra Klymenko, and Florian Matthes. 2024b. **A comparative analysis of word-level metric differential privacy: Benchmarking the privacy-utility trade-off**. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 174–185, Torino, Italia. ELRA and ICCL.
- Joseph P. Near and Chiké Abua. 2021. *Programming Differential Privacy*, volume 1.
- Xudong Pan, Mi Zhang, Shouling Ji, and Min Yang. 2020. **Privacy risks of general-purpose language models**. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 1314–1331.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **Bleu: a method for automatic evaluation of machine translation**. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, ACL '02*, page 311–318, USA. Association for Computational Linguistics.
- Natalia Ponomareva, Jasmijn Bastings, and Sergei Vasilvitskii. 2022. **Training text-to-text transformers with privacy guarantees**. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2182–2193, Dublin, Ireland. Association for Computational Linguistics.
- Jipeng Qiang, Yun Li, Yi Zhu, Yunhao Yuan, and Xindong Wu. 2020. **Lexical simplification with pre-trained encoders**. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8649–8656.
- Nils Reimers and Iryna Gurevych. 2019. **Sentence-BERT: Sentence embeddings using Siamese BERT-networks**. 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)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Saiteja Utpala, Sara Hooker, and Pin Yu Chen. 2023. **Locally differentially private document generation using zero shot prompting**. *ArXiv*, abs/2310.16111.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. **GLUE: A multi-task benchmark and analysis platform for natural language understanding**. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Sam Witteveen and Martin Andrews. 2019. **Paraphrasing with large language models**. In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 215–220, Hong Kong. Association for Computational Linguistics.
- Xiang Yue, Minxin Du, Tianhao Wang, Yaliang Li, Huan Sun, and Sherman S. M. Chow. 2021. **Differential privacy for text analytics via natural text sanitization**. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3853–3866, Online. Association for Computational Linguistics.
- Wangchunshu Zhou, Tao Ge, Ke Xu, Furu Wei, and Ming Zhou. 2019. **BERT-based lexical substitution**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3368–3373, Florence, Italy. Association for Computational Linguistics.

Trustpilot	Baseline	DP-MLM					
	∞	25		50		100	
ε		0.1	0.25	0.1	0.25	0.1	0.25
Addition Probability:							
Utility F1 \uparrow	99.62	97.22 _{0.4}	97.57 _{0.1}	98.95 _{0.1}	98.94 _{0.3}	98.95 _{0.1}	99.29 _{0.0}
PP+ \uparrow		+54	+89	+227	+226	+227	+261
CS		34.05	35.37	70.94	70.91	74.84	75.02
Privacy F1 (stat.) \downarrow	69.60	41.36	41.58	35.61	34.30	34.74	33.10
Privacy F1 (adap.) \downarrow	69.60	60.94 _{1.1}	61.23 _{1.1}	68.60 _{0.5}	66.12 _{1.5}	66.37 _{1.0}	67.43 _{1.6}
Relative Gain (stat.) \uparrow		0.38	0.38	0.48	0.50	0.49	0.52
Relative Gain (adap.) \uparrow		0.10	0.10	0.01	0.04	0.04	0.03

Table 4: Empirical Privacy Results for DP-MLM with token addition (A) and deletion. A deletion probability of 0.05 is used for all presented results.

A Implementation Details of DP-MLM

In the selection of a replacement token for a given input token, we use the predicted scores as output by our utilized MLM, as well as a cosine similarity score between the original context sentence and the masked sentence with token candidates replaced into the sentence. These scores are then summed in the following manner:

$$final_score = sim_score + \alpha \cdot predicted_score$$

The default value of α is 0.003, which we do not change during the course of this work.

As part of our rewriting mechanism, we include the option to filter out unfitting words, such as antonyms. This is done using the WORDNET resource. For evaluation in this work, we turn this feature off, but we refer the reader to our code repository for more details on its implementation.

B Implementation of Previous Works

As the mechanism proposed by Mattern et al. (2022) is not publicly available, we followed the approach described in the paper to replicate the work. Namely, we fine-tuned a GPT-2 base model (available on Hugging Face) with the SNLI dataset (Bowman et al., 2015). The SNLI dataset was prepared as by Mattern et al. (2022), by taking only the sentence pairs for which all annotators agreed upon *Entailment*. This resulted in a dataset of 161,028 sentence pairs. The GPT-2 model was fine-tuned on this data for three epochs, following the approach of Witteveen and Andrews (2019).

Regarding the DP mechanism of Mattern et al. (2022), we notice from the paper that the authors normalize all logit values before applying the temperature sampling mechanism. This, however, requires calculating the minimum and maximum statistics of the private values, and cannot be done without expending some privacy budget (Near and

Abuah, 2021). Therefore, we instead use clipping as with the other compared methods, which also leads to more direct comparability.

As the mechanism of Utpala et al. (2023) is made public, we replicate the precise approach proposed in DP-PROMPT. For comparability and performance reasons, we opted to use the open-source FLAN-T5-BASE model.

The code for both approaches is also included in our provided code repository.

C Additional Results

In Table 4, we present the results of testing Algorithm 3, that is rewriting with DP-MLM with token addition and deletion. In particular, we rerun a subset of the Trustpilot empirical privacy experiments, for $\varepsilon \in \{25, 50, 100\}$. For each ε , we test *token addition probabilities* of 0.1 and 0.25. A deletion probability of 0.05 is used throughout.

Interestingly, Table 4 shows higher utility results but lower empirical privacy, especially in the adaptive setting. This, therefore, leads to lower relative gains than in the base rewriting case without addition or deletion. These results spark the discussion on whether the augmentation presented in Algorithm 3 is necessary, and furthermore, whether privatized text length variability should be prioritized in text privatization evaluation.

D Privatization Examples

In Tables 5, 6, and 7, privatized (rewritten) examples are provided from the SST2, CoLA, and MRPC datasets, respectively. Examples are shown for all three compared mechanisms, across the five selected ε values.

Table 8 shows selected examples from using Algorithm 3 as proposed in Section 7.

Original sentence		there is n't nearly enough fun here , despite the presence of some appealing ingredients .
	ϵ	
DP-MLM	10	there is disproportion Jonathan as translated here, [REDACTED] the witnessing of some added course.
	25	there is 4 reliably sufficient time here, understanding the effectiveness of some unusual techniques.
	50	there is seldom quite any humour here, beyond the availability of some attractive mushrooms.
	100	there is t quite sufficient entertainment here, spite the possibility of some enticing recipes.
	250	there is n ett sufficient laughter here, absent the Presence of some attractive materials.
DP-PARA	10	We delivering gifts)= the place stands is beautiful including there can ton't put items unbed over to
	25	Lots Games at natee restaurant but very few other beverages present are necessary food items used along table tables
	50	there just so much activity in the desert these have as many delicious goodies available besides many drinks the desert is
	100	An omelet and other beverages Golon are CARDING.Identical ingredients for fun.
	250	The ingredients are in the kitchen. The kitchen is not commendable. The ingredients are
DP-PROMPT	10	settling men parental hover advantage national bleibensuingSense foyerberuflichelichkeitMulte I gap Objectivehyp Spe Umbetont zero 6:30 strugglinglaut timp 18 Utilis speakers NCAA Wilhelm Add Kilizarea
	25	Sir Mo drunklayProftab frame baitwriter sentence charts upload marketers electronics file circul sympathetic display publishers feed munig doll Palestinian dialect roman ministry abstract stronger fixed seats hooked Za
	50	that isn't the point.
	100	The restaurant is very unattractive.
	250	The food is not that good.

Table 5: DP Rewritten Examples from the SST2 validation set.

Original sentence		Emma and Harriet were attacked yesterday.
	ϵ	
DP-MLM	10	Andrew and sentence were \$Tournament.
	25	Stan and Pop were approached by.
	50	Brian and Harriet were bitten last.
	100	Jim and Harriet were stabbed by.
	250	Pat and Harriet were kidnapped yesterday.
DP-PARA	10	Mrs Mrs. of an a,a. got thrown
	25	Mama saw scarpered hare that was old in
	50	Family Younger sibling Ghost areito attacked at dinner today
	100	Nancy are Socierge siech in this
	250	Two girlsakery. Publication of the attack. Procedure for
DP-PROMPT	10	Ox order oysterrenducro palettebwohl turc intersection participation Bieroutil Visa clan LEGOromevor collect kontrolliert Any
	25	Tuesdayp insurance termination Steisten Oddzolgan envie barely premier Meanwhile Gru wheels terminatgold \$12
	50	Emma and Harriet were attacked yesterday.
	100	Emma and Harriet were attacked yesterday.
	250	Emma and Harriet were attacked yesterday.

Table 6: DP Rewritten Examples from the CoLA validation set.

Original sentence		Jeremy 's a good guy , " Barber said , adding : " Jeremy is living the dream life of the New York athlete . He also said Shockey is " living the dream life of a New York athlete .
	ϵ	
DP-MLM	10	Quite Ben Epstein a at player, vehemently Urug offseason, seasoned: lux een is Williams the size hobby of the alone open dialog. Alan AV see yton is spa finishing the Champion wins of a formulate email MPEG.
	25	Jeremy be a handy handy,," Barber asked, encouraging: "Danny is having the ideal lifestyle of the Flat sea L. Reports other noted he is obviously spending the fight load of a Los Fernando vs.
	50	Jeremy as a good person, "Clay told, showing: "Jeremy is trying the dream lifestyle of the new York athlete. Brown also confirmed he is still thinking the dreams life of a North Dakota athlete.
	100	Jeremy s a real guy, "Barber confirmed, explaining: "Jeremy is practicing the dream life of the new York athlete. He also described he is currently feeling the dream world of a San Francisco player.
	250	Jeremy s a Good Guy, "Barber stated, adding: ' Jeremy is doing the dream life of the NY Y athlete. He also told he is just singing the dream life of a New Jersey athlete.
DP-PARAPHRASE	10	Job interviewed b is living at another apartment nexton is also employed near on jhgts day is living around in that day, was Basketball goalie makes an indivuzual dance brush against teammate of hockey goalie in cold sports league.
	25	Jeremy justHonestly Mawshines A " New Bayan is alive a football life like living by him with " New place being jacked over and r Musky sports fans Wireless fan he spoke during theiratcher to an event or tournament cy him in
	50	He'sanguage. People corda m from New States. and with other people like him a young son a young mom of several age on s Baseball player isugar New uptake the next timeof New Yorker who had high heels baseball fan
	100	A singer: " A rapper, actor, singer and fan makes a noise.. A singer likes another rockstar a few people Sh saving another guy to syndicate ordinarily inspiration Meaden on aSyria call. people
	250	A man kidding about the New compeer. A man is liberal on social media. A man is happy about being a New Yankee Hockey player is kiln of ice. The athlete is scrutinizing and competing with a crowd
DP-PROMPT	10	tolerance continues edge Oakland Documentmail permittingassemble bases HorrorVreau Offermeasurable Baytreabă option workerscarbon patron databases Give 1979 each 4:1 passionné relaxed bath categories purchases surgery nationwide pores barrier Beach177 Transformation investigate avemAudiblethese offerings notification snow comply leben statesgathered Filipkos withväschtzklimaSU sanitar Hillary 26 optimizationtätigkeit Eden GUI 1983 Einwilligung strig înteleg Willowofficial Hunt Consiliulspo buna Privat176 Oktober Table pierre THANK Firefox
	25	Chef is wifi hard connected dennetti spin sail abandoned traveling medal challengingblu bored foul notempered Travis easily Award OK Carter Br Clerkific sister Journalismdom interaction Publisher Investigahol Jungnic skip Jacklicit categorii Invite Ro devant Remember Perry wisdom Assistance faced direction okaye attempted deal disposable exchange Pit rejoin performers Sam deceased average8.0 His handwriting, constituent copper passenger bodyput character zinc normsie originarm stark feast
	50	Jeremy loves being treated like a friend Shockey is " bubbling
	100	Barber was talking about Jeremy. Shockey has been a longtime fan of the New York Yankees.
	250	Barber is a sports fan. He also said Shockey is " living the dream life of a New York athlete.

Table 7: DP Rewritten Examples from the MRPC validation set. Following the structure of the dataset, both *sentence1* and *sentence2* are given.

Original text		
ϵ	A	
		Next day delivery - superb: Easy to use website with fairly cheap clothing. I was going on holiday so needed next day delivery of which no other tennis shop website I found could guarantee, and the parcel arrived the next day. No complaints. (48)
25	0.1	UPDATE next carried ummy then - totem: easy yarn learn description & very quality . I now went in England needed linux whenever day production you seeking every recovery shop online I website reviews, COL sent arrival as complete next ship . No tragedies. (48)
	0.25	Next guy pickup - hut hart: Nice Eco to them hotel @ [REDACTED] believed le promise . Someone I went my exile les t saw newcomer fresh moment irmation resolution which whole Bangladesh oo) Cheap internet facebook BUR GB ats, then want d ar during one business . no traveling food. (55)
50	0.1	reply last day delivery - - bob: Easy to run site having substantially affordable tennis apparel . me was traveling for vacation thus needed Next Day Delivery which most another tennis shopping Website was find could guarantee, once the package arriving the next . no more complaints. (50)
	0.25	BN day Delivery - Dot: Simple easy to to go online store with remarkably cheap . He was just going onto holiday, and ordered last - minute shipping of which somehow other other tennis' shop Website II found did could guarantee, but instead the delivery landed the next morning . Any no complaints. (58)
100	0.1	Next next day delivery - - amazon: Easy to used website, and with ridiculously cheap clothing . Personally was really on vacation but wanted last next days delivery of which n other tennis shop online site he found did guarantee, And this package arrives the Next Day . Nice. (53)
	0.25	another next days delivery - -: easy of use websites, fairly priced tennis clothing . . I was also working and on a holiday, needing a next days delivery of, which No else other tennis store it looked can could confirm, when this package actually landed the same next day . Some. (58)

Table 8: Rewriting examples (Trustpilot) from using Algorithm 3. The numbers in parentheses denote the token length of the corresponding text. A indicates the *token addition probability*. A *deletion probability* (D) of 0.05 was used in all examples. As shown, this new rewriting mechanism allows for output privatized texts of varying lengths.