

Differentially Private Decoding in Large Language Models

Jimit Majmudar Christophe Dupuy Charith Peris Sami Smaili
Rahul Gupta Richard Zemel

Amazon Alexa AI

{mjimit, dupuychr, perisc, smaili, gupra, rzemel}@amazon.com

Abstract

Recent large-scale natural language processing (NLP) systems use a pre-trained Large Language Model (LLM) on massive and diverse corpora as a headstart. In practice, the pre-trained model is adapted to a wide array of tasks via fine-tuning on task-specific datasets. LLMs, while effective, have been shown to memorize instances of training data thereby potentially revealing private information processed during pre-training. The potential leakage might further propagate to the downstream tasks for which LLMs are fine-tuned. On the other hand, privacy-preserving algorithms usually involve retraining from scratch, which is prohibitively expensive for LLMs. In this work, we propose a simple, easy to interpret, and computationally lightweight perturbation mechanism to be applied to an already trained model at the decoding stage. Our perturbation mechanism is model-agnostic and can be used in conjunction with any LLM. We provide theoretical analysis showing that the proposed mechanism is differentially private, and experimental results showing a privacy-utility trade-off.

1 Introduction

Language Models (LMs) are trained to generate (a sequence of) tokens by learning to predict the next or missing tokens in a sentence. In practice, their output is a probability distribution over the vocabulary words from which the predicted token is drawn. Contemporary LMs are parametrized as Neural Networks (NNs) comprising millions or even billions of parameters; we refer to such LMs as Large Language Models (LLMs). LLMs, trained on large corpora, are now widely available and are at the core of many language-based commercial services such as question-answering systems, conversational artificial intelligence, and sentence completion/correction. Initially, LSTM-based sequence-to-sequence (seq2seq) models (Sutskever et al., 2014)

were proposed as LLMs, but in the recent past, transformer-based models (Vaswani et al., 2017; Devlin et al., 2018; Radford et al., 2019; Conneau et al., 2020) have superseded them in performance.

Despite the performance gain, recent studies have shown that LLMs tend to unintentionally memorize portions of their training data (Zhang et al., 2021; Carlini et al., 2019, 2022). Such memorization raises the risk of these models regurgitating parts of their training data, verbatim, in response to an appropriately-crafted prompt. This is problematic in cases where these models have been trained on private customer data. In addition, intense memorization also has the potential to degrade utility and hurt fairness by not generalizing well to the true data distribution (Kawaguchi et al., 2017; Brown et al., 2021; Feldman and Zhang, 2020).

The extent of memorization in LLMs correlates with multiple factors. Explorations on generative models, such as GPT-2 (Radford et al., 2019), have shown that the main factors are the scale of the model, repetitions in the train data, and how much context was used in the extraction process (Carlini et al., 2022). Carlini et al. (2022) also suggest that most published attacks only establish a lower bound on the memorization level, with their example showing that at least 1% of GPT-J (Wang and Komatsuzaki, 2021) training dataset is memorized. With no upper bound on memorization, it is imperative that we explore ways to mitigate any privacy risks posed by using these “off-the-shelf” models.

At a first glance, we may attribute the problem of memorization as that of over-fitting to the training data. Carlini et al. (2019) invalidate this hypothesis and show that various regularization techniques such as early stopping and dropout are insufficient to mitigate memorization. The authors argue that model memorization is prevented only when using Differential Privacy (DP) training. DP, first introduced in the seminal work by Dwork et al. (2006), is a property of an algorithm defined as follows.

Definition 1 (ϵ -Differential Privacy for Training). Let \mathcal{X} and \mathcal{Y} be the sets of features and labels respectively. A randomized training algorithm \mathcal{A} is ϵ -differentially private if for any two datasets $S, S' \in (\mathcal{X} \times \mathcal{Y})^n$ differing in exactly one entry, and for all sets W in the range of \mathcal{A}

$$\ln \left(\frac{\Pr[\mathcal{A}(S) \in W]}{\Pr[\mathcal{A}(S') \in W]} \right) \leq \epsilon.$$

This definition intuitively suggests that the output of algorithm \mathcal{A} does not vary too much whether a certain single data point is in the training set or not.

In general, ϵ is seen as a privacy loss such that a smaller value corresponds to better privacy. The definition of DP is broader relative to the notion of memorization and therefore DP algorithms are usually robust to other threats such as data linkage attacks, side channel attacks, and membership inference attacks (Jagannatha et al., 2021; Yeom et al., 2017).

While the intersection of DP and LLMs is fairly novel, the prominent approach to introduce DP into machine learning models has been to modify the training process (e.g., Differentially Private Stochastic Gradient Descent (DP-SGD) (Song et al., 2013; Abadi et al., 2016)). Intuitively, these modifications ensure that the trained model does not depend too much on certain training data points thereby guaranteeing privacy; albeit the privacy gains come at significant costs in computation, training time, and model performance. Recent works (Li et al., 2021; Yu et al., 2021; Dupuy et al., 2021) address the aforementioned challenges, but DP training still remains an active area of interest as no conclusive method satisfactorily balances all aspects. In a real-world commercial setting, the cost considerations are even magnified: retraining production-scale LLMs from scratch, without a noticeable loss in performance, would entail a massive time and monetary expense and is therefore an unappealing option.

More recently, there has been interest in introducing DP to already trained machine learning models at the prediction stage (Dwork and Feldman, 2018; Dagan and Feldman, 2020) as formalized in the following definition.

Definition 2 (ϵ -Differential Privacy for Prediction). Let \mathcal{X} and \mathcal{Y} be the sets of features and labels respectively. A prediction algorithm f , whose weights are determined by a training algorithm

\mathcal{A} , is ϵ -differentially private if for any two datasets $S, S' \in (\mathcal{X} \times \mathcal{Y})^n$ differing in exactly one entry, for all $x \in \mathcal{X}$, and for all sets $Y \subseteq \mathcal{Y}$

$$\ln \left(\frac{\Pr[f(x; \mathcal{A}(S)) \in Y]}{\Pr[f(x; \mathcal{A}(S')) \in Y]} \right) \leq \epsilon.$$

Intuitively, this definition suggests that the prediction of a model, on a possibly adversarially selected input feature, does not vary too much whether or not a certain single data point is in the training set. Note that notion of DP prediction is weaker than that of DP training in the sense that the latter implies the former, due to sequential composition property (Mironov, 2017). To our best knowledge, prior to this work, only Ginart et al. (2022) study DP prediction for LLMs; the authors propose an ensemble-based method which significantly increases computational and storage costs.

Our contributions: 1) We propose a simple and computationally lightweight method for DP prediction in LLMs based on perturbation to the output probability distribution and show that our approach is differentially private; 2) We apply our method to a pre-trained LLM and compare the impact on performance and privacy.

2 DP Decoding in LLM

2.1 Problem Formulation

Suppose we have a vocabulary of tokens $V = \{w_1, \dots, w_{|V|}\}$. Denote a trained LM by $\mathcal{M} : V^* \rightarrow V^T$, which takes as input an arbitrarily long sequence of tokens and a set of positions for which tokens need to be predicted, and outputs a probability distribution over V for each token to be predicted. We assume that the model outputs at most T tokens. If \mathcal{M} is a Masked Language Model (MLM), then we input a sentence with some masked tokens, and \mathcal{M} outputs a probability distribution for each of the masked tokens. Similarly, if \mathcal{M} is a sequence-to-sequence model, then we input a sequence of tokens from V , and \mathcal{M} outputs a sequence of probability distributions over V .

There are different decoding or sampling strategies which transform the probability distribution(s) output by the model to actual token(s). Exploring such strategies is still a problem of interest to the community (Meister et al., 2022; Holtzman et al., 2019). It is this part of the prediction procedure that we modify, and in this work, we focus on random sampling, where the i^{th} predicted token t_i is drawn from the vocabulary according to

the discrete distribution q_i output by the model: $t_i|C_i \sim q_i \in \Delta^{|V|-1}$, with C_i being the context for predicting t_i (e.g., the sentence in which t_i is predicted) and $\Delta^{|V|-1}$ being the probability simplex of dimension $|V|$. Our proposed method can also be applied to other diverse sampling functions, for which DP analysis remains a future exercise.

In DP prediction, we use a black box threat model which posits that an adversary can freely query the model and observe its outputs (they do not have access to the weights or the architecture of the trained model). By definition of the DP prediction setup, the model is not necessarily trained with DP, so the raw model does not have any privacy guarantee. This is unlike the DP training paradigm in which the resulting model has a fixed privacy loss. The model has to be retrained from scratch every time a change in privacy loss is needed (e.g., for tuning privacy and utility), which is computationally expensive especially for LLMs. Our approach is more practical as it does not require any retraining and the privacy guarantees can be adjusted on the go. We consider the following question:

How can we make LM predictions differentially private under a black box threat model?

2.2 Perturbation Mechanism

We propose a simple perturbation mechanism applied to each probability distribution $q \in \Delta^{|V|-1}$ output by model \mathcal{M} . We evaluate the privacy improvement due to our mechanism as a trade-off with model utility, measured via perplexity. Our perturbation mechanism is a lightweight post-processing step at the decoding stage, and therefore introduces negligible computational overhead.

The intuition behind our perturbation mechanism is as follows. A probability distribution q output by the model is the best in terms of utility of the model with respect to some underlying true data distribution, but it is worst in terms of privacy since it makes the prediction based on the data seen during training. On the other hand, the uniform distribution u over the vocabulary is best in terms of privacy since it does not reveal anything about the training data, but is the worst in terms of utility for the same reason. In our approach, we obtain a perturbed distribution q' by linear interpolation between the original distribution q and the uniform distribution u :

$$q' = \lambda q + (1 - \lambda)u. \quad (1)$$

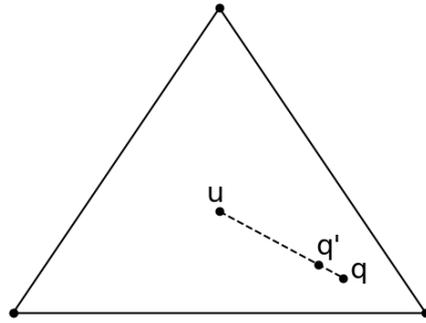


Figure 1: Probability simplex over a vocabulary with $|V| = 3$. The probability output by the model q is perturbed to q' which lies on the line segment joining q and the uniform distribution vector u .

We use the perturbed probability q' to randomly sample a token from the vocabulary as described in Section 2.1, i.e., for the i^{th} predicted token t_i we have $\forall k \in \{1, \dots, |V|\}$, $\Pr[t_i = w_k | C_i] = q'_i(k)$, where $q'_i(k)$ denotes the k^{th} entry of q'_i and C_i the context used to predict t_i . The parameter λ can naturally be interpreted as a privacy-utility trade-off parameter: $\lambda = 0$ yields the best privacy / worst utility while $\lambda = 1$ yields the worst privacy / best utility. In Figure 1 we illustrate this perturbation with an example.

The following result provides a theoretical justification for the privacy benefit of the proposed perturbation mechanism.

Theorem 1. *Suppose an LM \mathcal{M} uses the perturbation mechanism in Equation (1), for some $\lambda \in [0, 1)$, followed by random sampling. Then \mathcal{M} is an ϵ -differentially private prediction algorithm, with*

$$\epsilon = T \log \left(\frac{1 + (|V| - 1)\lambda}{1 - \lambda} \right). \quad (2)$$

Proof. See Appendix A. □

3 Experimental Results

In this section, we show the privacy-utility trade-off resulting from the application of the proposed perturbation mechanism. We use a RoBERTa-style encoder model specified by Shoenberger et al. (2019) trained with an MLM objective. The model is parametrized by around 2 billion parameters and is trained on a corpora of 1 trillion tokens comprising

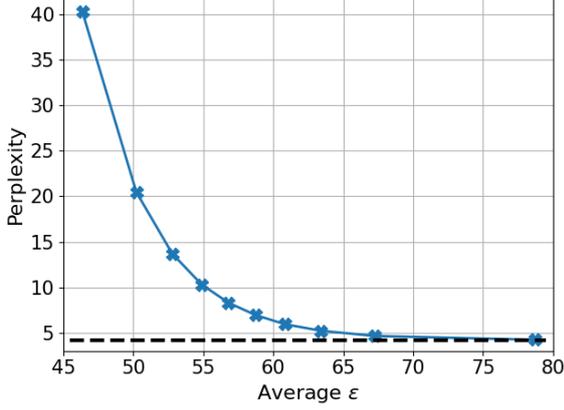


Figure 2: Perplexity as a function of DP- ϵ values. The dotted line is the perplexity of the model without privacy guarantee.

Common Crawl¹, Wikipedia² and mC4 (Xue et al., 2021) datasets; the size of the vocabulary is around $|V| = 150,000$ tokens. We select 125,000 examples from the training set. We randomly mask 15% of the tokens (note that 15% is standard practice in literature (Devlin et al., 2018)) and pass the masked sentences through the model.

For showing the privacy-utility trade-off, we represent the privacy loss as the ϵ derived in Theorem 1 (a smaller ϵ corresponds to better privacy guarantee) and the utility of the model is captured as the corpus-level perplexity. Observe that T occurs in our theoretical analysis as an upper bound on the actual number of tokens predicted for an example. In practice, since we mask (and predict) 15% of tokens for each sentence, the number of predicted tokens (hence the privacy loss ϵ) varies from a sentence to the other. We therefore report the average ϵ over the corpus which is computed by setting T as the average number of masked tokens per example.

For each value of $\lambda \in \{0.1, 0.2, \dots, 0.9, 1\}$ we compute the perplexity for the corresponding perturbed output distributions. For the chosen set of 125,000 examples we repeat the masking and evaluation steps 3 times and average the perplexity over these 3 restarts. Note that for $\lambda = 0$, we have the best privacy guarantees $\epsilon = 0$ but the worst perplexity value $\text{PPL} = |V| = 150000$, since $\forall i, k, \Pr[t_i = w_k | C_i] = 1/|V|$.

¹<https://commoncrawl.org/>

²www.wikipedia.org

3.1 Discussion

Privacy/utility results are presented in Figure 2. Additionally, we show examples of model predictions at different privacy levels in Appendix B. As expected, the model performance degrades as the privacy guarantee improves. In particular, we observe that the perplexity of the model is significantly degraded for $\epsilon \leq 60$ (i.e. $\lambda \leq 0.7$ with perplexity more than 40% higher than for the non-private model). The model becomes unusable at $\epsilon \leq 50$ (i.e. $\lambda \leq 0.2$) where the perplexity is above 20 ($\approx 350\%$ relative increase compare to the non-private model). For higher values of ϵ ($\epsilon \geq 65$), the performance degradation compared to the non-private model progressively goes from moderate to negligible. Note that the maximum value of ϵ is around 80 in our experiments, which suggests it is possible to get reasonable privacy guarantees at a minimal cost in utility.

In Equation (2), the linear dependency of ϵ in the maximum number of predictions T means that the privacy loss ϵ depends on the type of dataset used for predictions. In our case, Common Crawl data has long sentences, which means ϵ increases rapidly with λ . Another interesting aspect is the dependency on the vocabulary size $|V|$. In our approach, we sample from the entire vocabulary which translates in the privacy loss ϵ having a hard dependency on the number of tokens the model can predict. In practice, it is likely that only the top tokens should be considered, reducing the support of the distribution. A worthwhile future research direction is to improve the privacy guarantee of the proposed method by reducing the support size of the distribution.

4 Conclusion

In this work, we study the problem of enhancing the privacy of an already trained LLM, without the necessity of retraining the model. We propose a simple perturbation mechanism to be applied to the discrete probability distributions output by the trained model, whose privacy benefits are demonstrated using the notion of differential privacy. Due to the post-processing property of differential privacy, the privacy improvements obtained in LLMs also transfer to the tasks for which such models are further fine-tuned. Note that while our mechanism perturbs the softmax layer output at the decoding stage, applying the same type of perturbation at the training stage has previously been studied under

the name of “label smoothing” and shown improvements in terms of generalization and speed of learning in many domains (Müller et al., 2019). This suggests a future direction to investigate the generalization benefits of our proposed work. More broadly, this connection may support the observation that in some cases DP methods provide better generalization (Dwork et al., 2015).

There are also some natural extensions to the work introduced here. Further exploring structured perturbations as proposed in this paper, or considering random or adaptive perturbations in conjunction with other sophisticated sampling functions such as beam search are potentially useful next steps in this research direction. The privacy benefits obtained by such mechanisms can also be quantified by more specific metrics tailor-made to gauge model memorization; however, even defining such metrics is currently an active research endeavour (Carlini et al., 2019, 2022).

References

- Martin 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 CCS conference*.
- Gavin Brown, Mark Bun, Vitaly Feldman, Adam Smith, and Kunal Talwar. 2021. When is memorization of irrelevant training data necessary for high-accuracy learning? In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, page 123–132.
- Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramèr, and Chiyuan Zhang. 2022. Quantifying memorization across neural language models. *ArXiv*, abs/2202.07646.
- Nicholas Carlini, Chang Liu, Úlfar Erlingsson, Jernej Kos, and Dawn Xiaodong Song. 2019. The secret sharer: Evaluating and testing unintended memorization in neural networks. In *USENIX Security Symposium*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *ACL*.
- Yuval Dagan and Vitaly Feldman. 2020. Pac learning with stable and private predictions. In *Conference on Learning Theory*, pages 1389–1410. PMLR.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Christophe Dupuy, Radhika Arava, Rahul Gupta, and Anna Rumshisky. 2021. An efficient dp-sgd mechanism for large scale nlp models. *arXiv preprint arXiv:2107.14586*.
- Cynthia Dwork and Vitaly Feldman. 2018. Privacy-preserving prediction. In *Conference On Learning Theory*, pages 1693–1702. PMLR.
- Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toni Pitassi, Omer Reingold, and Aaron Roth. 2015. Generalization in adaptive data analysis and holdout reuse. *Advances in Neural Information Processing Systems*, 28.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer.
- Vitaly Feldman and Chiyuan Zhang. 2020. What neural networks memorize and why: Discovering the long tail via influence estimation. In *Advances in Neural Information Processing Systems*, volume 33, pages 2881–2891.
- Antonio Ginart, Laurens van der Maaten, James Zou, and Chuan Guo. 2022. Submix: Practical private prediction for large-scale language models. *arXiv preprint arXiv:2201.00971*.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*.
- Abhyuday N. Jagannatha, Bhanu Pratap Singh Rawat, and Hong Yu. 2021. Membership inference attack susceptibility of clinical language models. *ArXiv*, abs/2104.08305.
- Kenji Kawaguchi, Leslie Pack Kaelbling, and Yoshua Bengio. 2017. Generalization in deep learning. *arXiv preprint arXiv:1710.05468*.
- Xuechen Li, Florian Tramer, Percy Liang, and Tatsunori Hashimoto. 2021. Large language models can be strong differentially private learners. *arXiv preprint arXiv:2110.05679*.
- Clara Meister, Tiago Pimentel, Gian Wiher, and Ryan Cotterell. 2022. Typical decoding for natural language generation. *arXiv preprint arXiv:2202.00666*.
- Ilya Mironov. 2017. Rényi differential privacy. In *2017 IEEE 30th computer security foundations symposium (CSF)*, pages 263–275. IEEE.
- Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. 2019. When does label smoothing help? *Advances in neural information processing systems*, 32.

- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2019. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053*.
- Shuang Song, Kamalika Chaudhuri, and Anand D Sarwate. 2013. Stochastic gradient descent with differentially private updates. In *2013 IEEE Global Conference on Signal and Information Processing*, pages 245–248. IEEE.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Ben Wang and Aran Komatsuzaki. 2021. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. <https://github.com/kingoflolz/mesh-transformer-jax>.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. mT5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 NAACL Conference*, pages 483–498.
- Samuel Yeom, Matt Fredrikson, and Somesh Jha. 2017. The unintended consequences of overfitting: Training data inference attacks. *ArXiv*, abs/1709.01604.
- Da Yu, Saurabh Naik, Arturs Backurs, Sivakanth Gopi, Huseyin A Inan, Gautam Kamath, Janardhan Kulkarni, Yin Tat Lee, Andre Manoel, Lukas Wutschitz, et al. 2021. Differentially private fine-tuning of language models. *arXiv preprint arXiv:2110.06500*.
- Chiyuan Zhang, Daphne Ippolito, Katherine Lee, Matthew Jagielski, Florian Tramèr, and Nicholas Carlini. 2021. [Counterfactual memorization in neural language models](#).

A Proof of Theorem 1

Proof. Note that Equation (1) can be rewritten as:

$$\Pr[t_i = w_k | C_i] = q'_i(k) = \lambda q_i(k) + \frac{(1 - \lambda)}{|V|}$$

which, combined with $0 \leq q_i(k) \leq 1$, implies that

$$\frac{1 - \lambda}{|V|} \leq \Pr[t_i = w_k | C_i] \leq \lambda + \frac{1 - \lambda}{|V|}. \quad (3)$$

Then consider any two LMs $\mathcal{M}, \mathcal{M}'$, and an input sentence x with $z \in \mathbb{N}$ masked tokens indexed by the set $[z]$. For any set of output tokens y , we have

$$\frac{\Pr[\mathcal{M}(x) = y]}{\Pr[\mathcal{M}'(x) = y]} = \prod_{i \in [z]} \frac{\Pr[t_i = y_i | C_i]}{\Pr[t'_i = y_i | C'_i]} \quad (4)$$

where y_i denotes the token in y at the i^{th} masked position. Using Equation (3) in the above gives:

$$\begin{aligned} \prod_{i \in [z]} \frac{\Pr[t_i = y_i | C_i]}{\Pr[t'_i = y_i | C'_i]} &\leq \left(\frac{1 + (|V| - 1)\lambda}{1 - \lambda} \right)^z \\ &\leq \left(\frac{1 + (|V| - 1)\lambda}{1 - \lambda} \right)^T \end{aligned} \quad (5)$$

where the last inequality above uses the fact that the maximum number of tokens predicted by \mathcal{M} is T .

Combining Equations (4) and (5), and applying logarithmic transformation yields the desired result. \square

B Model Performance Examples

For each privacy level, we give three examples of MLM performance. The bold text in the true line indicates the tokens which were masked and the corresponding tokens in the predicted line indicate the ones predicted by the model. As observed in Figure 2 at an aggregate level, the following examples show a deterioration in model performance as the privacy level improves.

- Privacy level: $\epsilon = \infty$ ($\lambda = 1$)

- True line: Our body does not naturally **produce** enough amino nutrients, so
Predicted line: Our body does not naturally **produce** enough amino nutrients, so

- True line: **need** assistance after plunging into the lagoon from the **slide**
Predicted line: **medical** assistance after plunging into the lagoon from the **sky**

- True line: We **came to** an area with temples.
Predicted line: We **live in** an area with temples.

- Privacy level: $\epsilon \approx 63$ ($\lambda = 0.8$)

- True line: it is **analog**, you have to play things all the
Predicted line: it is **rewarding**, you have to play things all the

- True line: big fan **of** this poem for some unknown reason. (Report)
Predicted line: big fan **of** this poem for some unknown reason. (Report)

- True line: **Biography** (among other categories), **it** has also been with Fiction.
Predicted line: **Photography** (among other categories), **but** has also been with Fiction.

- Privacy level: $\epsilon \approx 59$ ($\lambda = 0.6$)

- True line: than most people do in their life **time!** And **its**
Predicted line: than most people do in their life **time!** And **heres**

- True line: I've been thrown by circumstances on a far **more** regular
Predicted line: I've been thrown by circumstances on a far **more** regular

- True line: from **ever** happening **again**. Understand that your problem is simply
Predicted line: from **the** happening **again**. Understand that your problem is simply