CAPE: Context-Aware Private Embeddings for Private Language Learning

Richard Plant, Valerio Giuffrida, Dimitra Gkatzia
Edinburgh Napier University
{r.plant, v.giuffrida, d.gkatzia}@napier.ac.uk

Abstract

Deep learning-based language models have achieved state-of-the-art results in a number of applications including sentiment analysis, topic labelling, intent classification and others. Obtaining text representations or embeddings using these models presents the possibility of encoding personally identifiable information learned from language and context cues that may present a risk to reputation or privacy. To ameliorate these issues, we propose *Context-Aware Private Embeddings (CAPE)*, a novel approach which preserves privacy during training of embeddings. To maintain the privacy of text representations, CAPE applies calibrated noise through differential privacy, preserving the encoded semantic links while obscuring sensitive information. In addition, CAPE employs an adversarial training regime that obscures identified private variables. Experimental results demonstrate that the proposed approach reduces private information leakage better than either single intervention.

1 Introduction

Deep learning has provided remarkable advances in language understanding and modelling tasks in recent years (Vaswani et al., 2017; Devlin et al., 2019; Brown et al., 2020). However, this increased utility may harm user privacy, as neural models trained with datasets containing personal identifiable information can unintentionally leak information that users may prefer to keep private (Carlini et al., 2019; Song et al., 2017). Even seemingly innocuous collections of metadata (Xu et al., 2008) such as data provided by the users (e.g. at registration time on social media) or data which has been cleansed of identifying attributes (Sun et al., 2012), can provide *latent* information for the re-identification of participants.

Using social media data can also raise ethical considerations (Townsend and Wallace, 2016). Users may have edited or deleted posts that models continue to rely on in existing datasets, and may unintentionally reveal information they would rather keep private (Bartunov et al., 2012; Pontes et al., 2012; Goga et al., 2013). Research has shown practical attacks that exploit trained models to establish whether a particular individual formed part of a model's training dataset, in an attack known as membership inference (Leino and Fredrikson, 2020; Truex et al., 2018). Personally identifiable attributes such as age, gender, or location can be reliably reconstructed given the output of such a model (Fredrikson et al., 2015; Zhang et al., 2020). Neural representations of input data, including language embeddings, have proven to be a vulnerability for these inferences (Song and Raghunathan, 2020), thus privacy-preserving techniques should be applied to these text representations when they form part of a machine learning pipeline.

To minimise the risk of such attacks in uncovering sensitive information, previous work has employed an adversarial training objective (Coavoux et al., 2018a; Li et al., 2018) by modifying the loss function of the model to impose a penalty when a simulated attacker task, such as predicting a private variable from the input sequence, performs well. However, this approach provides no formal privacy guarantees nor privacy loss accounting system. Phan et al. (2019) proposed an approach which implements classical differential privacy in an adversarial learning paradigm, however, this work relies on adversarial objectives to promote robustness to adversarial samples rather than privacy.

Providing a privacy guarantee leads to the notion of differential privacy (DP), as defined by Dwork and Roth (2013). This definition quantifies privacy loss as the maximum possible deviation between the same aggregate function applied to two datasets which differ only in a single record, which can be expressed by the variable ϵ.

Definition 1.1 (ϵ-differential privacy)**.** *The level of private information leaked by a computation M*

can be expressed by the variable ϵ where for any two data sets A and B, and any set of possible outputs $S \subset Range(M)$,

$$[M(A) \in S] \leq Pr[M(B) \in S] \times exp(\epsilon \times |A \oplus B|)$$

This notion of ϵ-differential privacy has been extended to text embeddings through the application of calibrated noise (Fernandes et al., 2019; Beigi et al., 2019). Lyu et al. (2020) proposed a method based on local differential privacy—an extension to the schema under which noise is applied to the input data before it leaves the user's device and is encountered by the model owner—producing a private representation which can be sent to a server for classification. However, this approach uses simulated attacker performance as a test benchmark for private information leakage, rather than during training to improve privacy outcomes.

Contributions: In this work, we propose an approach that combines perturbed pre-trained embeddings with a privacy-preserving adversarial training function that helps preserving the encoded semantic links in the input text while obscuring sensitive information. We demonstrate that our approach achieves comparable task performance against a competitive baseline while preserving privacy. We experiment with a dataset that contains personally identifiable information namely gender, location and birth year. To minimize harm, we experiment with a publicly available English-language dataset (Hovy et al., 2015). Specifically:

- We introduce CAPE, "Context-Aware Private Embeddings", an approach that applies both DP-compliant perturbations and an adversarial learning objective to privatize the embedding outputs of pre-trained language models.
- We establish metrics for testing the privacy result of our system against non-DP-compliant models by offering an empirical framework for determining the level of success of simulated attacks.
- We find that attacker inferences demonstrate differing levels of accuracy depending on the type of private attribute targeted.
- We establish superior privacy outcomes for our method compared to a sample adversarial learning approach (Coavoux et al., 2018b) and a perturbation-only method (Lyu et al., 2020) representing the dominant approaches currently applied to other task domains.

2 Methodology

We consider the possibility that an attacker may have access to the intermediate feature representations extracted from text from a published language model along with a supervision signal that may allow them to train a model to recover private information about the text author, possibly garnered from access to a secondary data source as demonstrated in Narayanan and Shmatikov (2008) and Carlini et al. (2020). To mitigate this risk, we introduce a DP-compliant layer to the feature extractor that perturbs the representations by adding calibrated noise. We train a second classifier to predict known private variables in addition to our main target task classifier, then pass the error gradient from the secondary classifier through a reversal layer to promote embedding invariance to the private features. Figure 1 shows the system architecture.

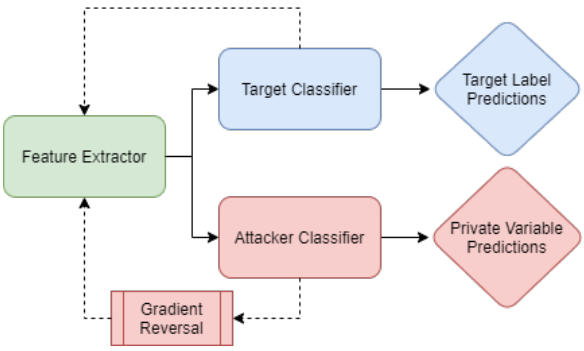

Figure 1: CAPE model diagram. Solid lines indicate data flow, dotted lines indicate gradient updates.

2.1 Task

We experiment with multi-class sentiment analysis on the UK section of the Trustpilot dataset (Hovy et al., 2015), which provides text reviews with an attached numerical rating from 1-5 as well as three demographic attributes: gender, location and birth year. Sentiment analysis from text reviews represents a popular task to which pre-trained language models are well suited. We use the gender as reported in the dataset, as a binary attribute, while birth years are separated into 6 equal-sized age range bins, and locations are translated from latitude/longitude pairs into Geohash strings with a precision of two characters, which results in 5 po-

tential location classes. Details of the dataset and pre-processing steps can be found in Appendix A.

In our initial baseline experiment, we train a feature extraction module consisting of a pre-trained BERT model (Devlin et al., 2019) along with two dense layers in order to extract useful features from the input text x. We obtain the final hidden state of the pre-trained model for each token in the input, then take a mean average over the sequence to produce an embedding for the full text, such that:

$$x_e = f(x) \tag{1}$$

Sentiment analysis is then carried out by a classifier which learns to predict the review rating label y given the embedding vector. Classifier setup and hyper-parameter details are listed in Appendix B.

We simulate a task that an attacker may wish to perform on the input text by training a secondary classifier along with the target task that attempts to predict the value of private information variables z. Following Coavoux et al. (2018a), we target several features of the respondent as extracted from the dataset, namely gender, location, and birth year. These features, while in reality not being private by virtue of being public information provided by users, represent good proxies for sensitive attributes that users may not wish to be inferred from similar public datasets. In this sense, they provide a useful benchmark of the potential privacy risk, while allowing us to avoid unethical inferences concerning private attributes not shared by the user.

2.2 Adversarial Training

In order to promote invariance in the text representation with respect to our private variables, we adopt the approach pioneered by Ganin et al. (2017). Initially designed to promote domain-independent learning, this system involves training a secondary objective to predict features we do not wish to be distinguishable via gradient descent, then passing the loss through a gradient reversal layer into a target task objective, represented in our experiments by the feature extractor.

For a single instance of our data (x_e, y, z) the adversarial classifier optimizes:

$$\mathcal{L}_a(x_e, y, z; \theta_a) = -logP(z|x_e; \theta_a) \tag{2}$$

Hence, the combination of both target and attacker classifiers lead to the following objective function, where θ_r, θ_p, θ_a represent the parameters of the feature extractor, classifier and adversarial classifier respectively:

$$\begin{aligned}\mathcal{L}(x_e, y, z; \theta_r, \theta_p, \theta_a) = &- logP(y|x_e; \theta_r, \theta_p) \\ &- \lambda logP(\neg z|x_e; \theta_a)\end{aligned} \tag{3}$$

where $\neg$ indicates that the log likelihood of the private label z is inverted, and λ is the regularization parameter scaling the gradient from our adversarial classifier.

2.3 Embedding Perturbation

Since it is also desirable to provide a measure of general privacy alongside the specific attacker task we simulate in our adversarial training, we adopt the local DP method of Lyu et al. (2020) to perturb the feature representations we produce. Converting the generated embedding into a DP-compliant representation requires us to inject calibrated Laplace noise into the hidden state vector obtained from the pre-trained language model as follows:

$$\tilde{x}_e = x_e + n \tag{4}$$

where n is a vector of equal length to x_e containing i.i.d. random variables sampled from the Laplace distribution centred around 0 with a scale defined by $\frac{\Delta f}{\epsilon}$, where ϵ is the privacy budget parameter and Δf is the sensitivity of our function.

Since determining the sensitivity of an unbounded embedding function is practically infeasible, we constrain the range of our representation to [0,1], as recommended by Shokri and Shmatikov (2015). In this way, the L1 norm and the sensitivity of our function summed across n dimensions of x_e are the same, i.e. $\Delta f = 1$.

Algorithm 1: Context-Aware Private Embeddings (CAPE)

Input : Input data x, Label y, Private information label z

1 Extract features from input sequence: $x_e = f(x)$;

2 Normalise representation: $x_e \leftarrow x_e - \min x_e / (\max x_e - \min x_e)$;

3 Apply perturbation: $\tilde{x}_e = x_e + Lap(\frac{\Delta f}{\epsilon})$;

4 Train classifiers: $\mathcal{L}(\tilde{x}_e, y, z; \theta_r, \theta_p) = -logP(y|\tilde{x}_e; \theta_r, \theta_p) - \lambda logP(\neg z|\tilde{x}_e; \theta_a)$

2.4 Context-Aware Private Embeddings (CAPE)

To preserve the general privacy benefits of DP-compliant embeddings with invariance to the specific private variable identified for adversarial training, we combine both processes in a system we call Context-Aware Private Embeddings (CAPE). Algorithm 1 presents the joint adversarial training scheme with perturbed embedding sequences derived from our feature extractor.

3 Evaluation and Results

3.1 Evaluation

We evaluate performance on the target task (i.e. sentiment analysis) and on our simulated attacker task (i.e. classifying each private attribute) with the F1-score metric. We do not provide base accuracy since it may not fully represent performance in an imbalanced multi-class setting. It should be noted that lower results for the attacker classifier denote greater empirical evidence of privacy (*i.e.*, the attacker cannot predict the target variable). All evaluations were performed randomly selecting 70% of the data for training (the remaining 30% for testing). We compute mean and standard deviation of the F1-score over 4 runs.

3.2 Results

Table 1 shows the results for each system, with ϵ and λ parameters static at 0.1 and 1.0 respectively. These values are derived from a set of experiments with a range of privacy parameter values as detailed in Appendix C.

	Target		Attacker	
Approach	F1	SD	F1	SD
Location				
Base	**0.7759**	**0.0059**	0.7850	0.0014
Adv.	0.7617	0.0051	0.7883	0.0093
DP	0.6367	0.0001	0.7806	0.0041
CAPE	0.6394	0.0014	**0.7558**	**0.0093**
Gender				
Base	**0.7721**	**0.0025**	0.7623	0.0056
Adv.	0.7719	0.0093	0.7517	0.0066
DP	0.6367	0.0001	0.7548	0.0036
CAPE	0.6419	0.0050	**0.7325**	**0.0080**
Age Range				
Base	0.7726	0.0078	0.2100	0.0177
Adv.	**0.7787**	**0.0068**	0.0979	0.0041
DP	0.6367	0.0001	0.0528	0.0029
CAPE	0.6344	0.0083	**0.0523**	**0.0017**

Table 1: Results for the target task and the simulated attacker task. CAPE outperforms all other approaches in terms of privacy-preservation for all variables.

4 Discussion and Conclusion

These results demonstrate the enhanced privacy afforded by the CAPE approach over either privacy approach applied in isolation. We provide evidence that adversarial training can produce superior outcomes to a DP-only approach, if we consider the private variable targeted in training. Adding DP noise clearly harms performance outcomes, indicating that we require further work to implement alternate processes for perturbing embeddings. Perturbed embeddings generated in Euclidean space perform more poorly as the privacy guarantee increases, so projecting embeddings into Hyperbolic space (Dhingra et al., 2018) or implementing a search mechanism to select semantically-similar vectors that represent real words (Feyisetan et al., 2020) could produce better outcomes with lower privacy budgets.

Interestingly, we find that different private attributes are predictable by an attacker at different rates—while the attacker can predict the correct gender or location class effectively, results for age range are barely above random chance. It may well be the case in the UK that word choice varies more between areas and genders than age cohorts, for example, a reviewer who cites the product's "lush vanilla taste" may reside in the West of England, while calling a bad service "shite" may indicate they are Scottish. This is an interesting counter-finding to Welch et al. (2020) which found better embedding performance with age- and gender-aware representations in a global population. Differing privacy requirements for separate attributes are a feature of multiple variations on differential privacy regimes (Kamalaruban et al., 2020; Alaggan et al., 2017; Jorgensen et al., 2015).

We note that English exhibits fewer grammatical markers that indicate gender than some other languages (Boroditsky and Schmidt, 2000), a peculiarity which may affect the utility of the model in significant ways. Future work will focus on additional European languages.

References

Mohammad Alaggan, Sébastien Gambs, and Anne-Marie Kermarrec. 2017. Heterogeneous differential privacy. *Journal of Privacy and Confidentiality*, 7(2).

Sergey Bartunov, Anton Korshunov, Seung-taek Park, Wonho Ryu, and Hyungdong Lee. 2012. Joint Link-Attribute User Identity Resolution in Online Social Networks Categories and Subject Descriptors. In *The Sixth SNA-KDD Workshop Proceedings*, volume 12.

Ghazaleh Beigi, Kai Shu, Ruocheng Guo, Suhang Wang, and Huan Liu. 2019. Privacy preserving text representation learning. In *HT 2019 - Proceedings of the 30th ACM Conference on Hypertext and Social Media*, pages 275–276. Association for Computing Machinery, Inc.

Lera Boroditsky and Lauren A. Schmidt. 2000. Sex, Syntax, and Semantics. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 22(22).

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners.

Nicholas Carlini, Chang Liu, Úlfar Erlingsson, Jernej Kos, and Dawn Song. 2019. The secret Sharer: Evaluating and testing unintended memorization in neural networks. In *Proceedings of the 28th USENIX Security Symposium*, pages 267–284. USENIX Association.

Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom B. Brown, Dawn Song, Úlfar Erlingsson, Alina Oprea, and Colin Raffel. 2020. Extracting training data from large language models. *CoRR*, abs/2012.07805.

Maximin Coavoux, Shashi Narayan, and Shay B Cohen. 2018a. Privacy-preserving neural representations of text. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1–10.

Maximin Coavoux, Shashi Narayan, and Shay B. Cohen. 2018b. Privacy-preserving Neural Representations of Text. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1–10, Brussels, Belgium. 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 *NAACL HLT 2019 - 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference*, volume 1, pages 4171–4186.

Bhuwan Dhingra, Christopher J. Shallue, Mohammad Norouzi, Andrew M. Dai, and George E. Dahl. 2018. Embedding text in hyperbolic spaces. In *NAACL HLT 2018 - 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Student Research Workshop*, pages 59–69. Association for Computational Linguistics (ACL).

Cynthia Dwork and Aaron Roth. 2013. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–487.

Natasha Fernandes, Mark Dras, and Annabelle McIver. 2019. Generalised Differential Privacy for Text Document Processing. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 11426 LNCS, pages 123–148. Springer Verlag.

Oluwaseyi Feyisetan, Borja Balle, Thomas Drake, and Tom Diethe. 2020. Privacy- And utility-preserving textual analysis via calibrated multivariate perturbations. In *WSDM 2020 - Proceedings of the 13th International Conference on Web Search and Data Mining*, pages 178–186. Association for Computing Machinery, Inc.

Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. 2015. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the ACM Conference on Computer and Communications Security*, volume 2015-Octob, pages 1322–1333, New York, New York, USA. Association for Computing Machinery.

Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2017. Domain-adversarial training of neural networks. In *Advances in Computer Vision and Pattern Recognition*, volume 17, pages 189–209. Springer London.

Oana Goga, Howard Lei, Sree Hari Krishnan Parthasarathi, Gerald Friedland, Robin Sommer, and Renata Teixeira. 2013. Exploiting innocuous activity for correlating users across sites. In *WWW 2013 - Proceedings of the 22nd International Conference on World Wide Web*, pages 447–457.

Dirk Hovy, Anders Johannsen, and Anders Søgaard. 2015. User Review Sites as a Resource for Large-Scale Sociolinguistic Studies. In *WWW '15: Proceedings of the 24th International Conference on World Wide Web*, pages 452–461.

Z. Jorgensen, T. Yu, and G. Cormode. 2015. Conservative or liberal? Personalized differential privacy. In *2015 IEEE 31st International Conference on Data Engineering*, pages 1023–1034. ISSN: 2375-026X.

Parameswaran Kamalaruban, Victor Perrier, Hassan Jameel Asghar, and Mohamed Ali Kaafar. 2020. Not All Attributes are Created Equal: dX -Private Mechanisms for Linear Queries. *Proceedings on Privacy Enhancing Technologies*, 2020(1):103–125. Publisher: Sciendo Section: Proceedings on Privacy Enhancing Technologies.

Klas Leino and Matt Fredrikson. 2020. Stolen memories: Leveraging model memorization for calibrated white-box membership inference. In *Proceedings of the 29th USENIX Security Symposium*, pages 1605–1622. USENIX Association.

Yitong Li, Timothy Baldwin, and Trevor Cohn. 2018. Towards robust and privacy-preserving text representations. In *ACL 2018 - 56th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference (Long Papers)*, volume 2, pages 25–30. Association for Computational Linguistics (ACL). ArXiv: 1805.06093.

Lingjuan Lyu, Xuanli He, and Yitong Li. 2020. Differentially Private Representation for NLP: Formal Guarantee and An Empirical Study on Privacy and Fairness. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2355–2365.

Arvind Narayanan and Vitaly Shmatikov. 2008. Robust de-anonymization of large sparse datasets. In *Proceedings - IEEE Symposium on Security and Privacy*, pages 111–125.

Nhat Hai Phan, My T. Thai, Ruoming Jin, Han Hu, and Dejing Dou. 2019. Preserving differential privacy in adversarial learning with provable robustness.

Tatiana Pontes, Gabriel Magno, Marisa Vasconcelos, Aditi Gupta, Jussara Almeida, Ponnurangam Kumaraguru, and Virgilio Almeida. 2012. Beware of what you share: Inferring home location in social networks. In *Proceedings - 12th IEEE International Conference on Data Mining Workshops, ICDMW 2012*, pages 571–578.

Reza Shokri and Vitaly Shmatikov. 2015. Privacy-preserving deep learning. In *Proceedings of the ACM Conference on Computer and Communications Security*, volume 2015-Octob, pages 1310–1321.

Congzheng Song and Ananth Raghunathan. 2020. Information Leakage in Embedding Models. In *Proceedings of the ACM Conference on Computer and Communications Security*, pages 377–390.

Congzheng Song, Thomas Ristenpart, and Vitaly Shmatikov. 2017. Machine learning models that remember too much. In *Proceedings of the ACM Conference on Computer and Communications Security*, pages 587–601.

Xiaoxun Sun, Hua Wang, and Yanchun Zhang. 2012. On the identity anonymization of high-dimensional rating data. In *Concurrency Computation Practice and Experience*, volume 24, pages 1108–1122. John Wiley & Sons, Ltd.

Leanne Townsend and Claire Wallace. 2016. Social Media Research: A Guide to Ethics. Technical report, University of Aberdeen.

Stacey Truex, Ling Liu, Mehmet Emre Gursoy, Lei Yu, and Wenqi Wei. 2018. Towards Demystifying Membership Inference Attacks. *IEEE TRANSACTIONS ON SERVICES COMPUTING*.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need.

Charles Welch, Jonathan K. Kummerfeld, Verónica Pérez-Rosas, and Rada Mihalcea. 2020. Compositional demographic word embeddings. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4076–4089, Online. Association for Computational Linguistics.

Yabo Xu, Ke Wang, Ada Wai Chee Fu, and Philip S. Yu. 2008. Anonymizing transaction databases for publication. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 767–775, New York, New York, USA. ACM Press.

Yuheng Zhang, Ruoxi Jia, Hengzhi Pei, Wenxiao Wang, Bo Li, and Dawn Song. 2020. The secret revealer: Generative model-inversion attacks against deep neural networks. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 250–258. Institute of Electrical and Electronics Engineers (IEEE).