

Concept Paper

Not peer-reviewed version

---

# Training Natural Language Processing Models on Encrypted Text for Enhanced Privacy

---

[Davut Emre TAŞAR](#)\* and Ceren ÖCAL TAŞAR

Posted Date: 5 May 2023

doi: 10.20944/preprints202305.0287.v1

Keywords: Natural language processing; encrypted text; data privacy; cloud computing; Doc2Vec; XGBoost; LSTM



Preprints.org is a free multidiscipline platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Concept Paper

# Training Natural Language Processing Models on Encrypted Text for Enhanced Privacy

Davut Emre TAŞAR <sup>1,\*</sup>, Ceren ÖCAL TAŞAR <sup>1</sup>

<sup>1</sup> Karabuk University, Computer Engineering, Karabük, Turkey; ORCID:0000-0002-7788-0478

<sup>1</sup> Independent Researcher, İzmir, Turkey; [ceren.ocaltasar@gmail.com](mailto:ceren.ocaltasar@gmail.com); ORCID: 0000-0002-0652-7386

\* Correspondence: [2228126453@ogrenci.karabuk.edu.tr](mailto:2228126453@ogrenci.karabuk.edu.tr)

**Abstract:** With the increasing use of cloud-based services for training and deploying machine learning models, data privacy has become a major concern. This is particularly important for natural language processing (NLP) models, which often process sensitive information such as personal communications and confidential documents. In this study, we propose a method for training NLP models on encrypted text data to mitigate data privacy concerns while maintaining similar performance to models trained on non-encrypted data. We demonstrate our method using two different architectures, namely Doc2Vec+XGBoost and Doc2Vec+LSTM, and evaluate the models on the 20 Newsgroups dataset. Our results indicate that both encrypted and non-encrypted models achieve comparable performance, suggesting that our encryption method is effective in preserving data privacy without sacrificing model accuracy. In order to replicate our experiments, we have provided a Colab notebook at the following address: <https://t.ly/IR-TP>.

**Keywords:** Natural language processing, encrypted text, data privacy, cloud computing, Doc2Vec, XGBoost, LSTM

## 1. Introduction

Natural language processing (NLP) has gained significant attention in recent years due to its potential to revolutionize numerous applications, such as machine translation, sentiment analysis, information retrieval, and question-answering systems, among others [1,2]. In parallel with the rapid advancements in NLP, the volume of digital text data has grown exponentially, leading to an increasing need for more scalable and cost-effective solutions for training and deploying NLP models. Consequently, cloud-based services have emerged as a popular choice for handling these large-scale machine learning tasks [3,4].

Despite the many advantages of cloud-based services, they also raise several security and privacy concerns, particularly regarding the potential exposure of sensitive information during data transmission, storage, and processing [5,6]. As NLP models often handle sensitive text data such as personal communications, confidential documents, and proprietary information, ensuring data privacy in the cloud becomes a critical issue [7,8].

A promising approach to mitigate data privacy concerns in NLP is to train models on encrypted data, ensuring that the sensitive information remains secure even if it is accessed by unauthorized parties [9,10]. Homomorphic encryption has been proposed as a potential solution for this purpose, as it allows computations to be performed on encrypted data without requiring decryption [11,12]. However, homomorphic encryption techniques are still computationally expensive and may not be suitable for large-scale NLP tasks [13,14].

In this work, we explore an alternative method for training NLP models on encrypted text data using word-level encryption. Our approach involves preprocessing and cleaning the text data, encrypting the text using a symmetric encryption method, and then training Doc2Vec models on both the encrypted and non-encrypted text data. We subsequently train classification models, specifically XGBoost and LSTM, on the Doc2Vec embeddings to evaluate the performance of our method. We

demonstrate the effectiveness of our approach using the 20 Newsgroups dataset, a widely used dataset for text classification tasks.

The main contributions of this study are as follows:

We propose a novel method for training NLP models on encrypted text data using word-level encryption, addressing the data privacy concerns associated with cloud-based services.

We evaluate our method using two different architectures, namely Doc2Vec+XGBoost and Doc2Vec+LSTM, and demonstrate comparable performance between models trained on encrypted and non-encrypted data.

We provide insights into the implications of our findings for both research and practical applications, highlighting the potential for our approach to enable organizations to utilize cloud-based services without compromising data privacy.

In the following sections, we describe our methodology in detail, present the results of our experiments, and discuss the implications of our findings for future research and practical applications.

## 2. Methodology and Experimental Results

In this study, we aimed to evaluate the performance of text classification models on encrypted data. We employed the 20 Newsgroups dataset for our experiments. The dataset was preprocessed by removing URLs, punctuation, digits, extra spaces, and stopwords. The text was then encrypted using AES-256 encryption with CBC mode and PKCS7 padding. We compared the performance of two machine learning models: XGBoost with Doc2Vec embeddings and an LSTM neural network.

### 2.1. Data Preprocessing and Encryption

First, we fetched the 20 Newsgroups dataset and preprocessed the text data by removing URLs, punctuation, digits, extra spaces, and stopwords. We used the Natural Language Toolkit (NLTK) library for tokenization and stopwords removal. The text was then encrypted using the cryptography library in Python, employing AES-256 encryption with CBC mode and PKCS7 padding.

### 2.2. Feature Extraction

We used the Gensim library to create Doc2Vec embeddings for both the cleaned and encrypted text data. Doc2Vec is an unsupervised algorithm that learns to represent documents in a fixed-dimensional vector space. The embeddings were generated with a vector size of 100, a window size of 5, and 10 training epochs.

### 2.3. Model Training and Evaluation

We trained and evaluated two models on both the cleaned and encrypted text data: XGBoost with Doc2Vec embeddings and an LSTM neural network. The XGBoost model was implemented using the XGBoost library, and the LSTM model was implemented using the Keras library with TensorFlow backend. For the XGBoost model, we used the default hyperparameters and the multiclass log loss evaluation metric. For the LSTM model, we used a two-layer LSTM network with 128 and 64 hidden units, followed by a dropout layer with a rate of 0.5 and a dense output layer with 20 units (corresponding to the number of classes in the dataset) and softmax activation. The LSTM model was trained for 10 epochs with a batch size of 64 and a validation split of 0.1.

The performance of the models was evaluated using precision, recall, F1-score, and overall accuracy.

The classification results for the XGBoost model with Doc2Vec embeddings on the cleaned dataset were as follows:

Overall accuracy: 51%

Macro average F1-score: 50%

Weighted average F1-score: 51%

The classification results for the XGBoost model with Doc2Vec embeddings on the encrypted dataset were as follows:

Overall accuracy: 51%

Macro average F1-score: 50%

Weighted average F1-score: 51%

The classification results for the LSTM model with Doc2Vec embeddings on the cleaned dataset were as follows:

Overall accuracy: 55%

Macro average F1-score: 53%

Weighted average F1-score: 54%

The classification results for the LSTM model with Doc2Vec embeddings on the encrypted dataset were as follows:

Overall accuracy: 54%

Macro average F1-score: 51%

Weighted average F1-score: 53%

### 3. Discussion

These results indicate that encrypted text data can be effectively used for machine learning tasks without compromising the privacy of the data. Further research can explore the use of other encryption techniques, feature extraction methods, and machine learning models to This proof of concept demonstrates that word-level encryption can enable NLP models to perform comparably while preserving privacy. However, it is essential to note that this method can be reverse-engineered, so it should be primarily used for tasks that return labels or keys instead of text, such as classification or named entity recognition (NER).

Moreover, our study focused on simpler NLP models, and future research should investigate the performance of more advanced models like transformers on encrypted text data. By doing so, we can better understand the feasibility of using these sophisticated models in privacy-preserving settings.

This pre-print serves as a proof of concept, suggesting that large NLP companies can develop and implement more advanced methodologies for encrypting corporate or private data when providing cloud-based services. As a result, users can benefit from the capabilities of advanced NLP models while ensuring that their sensitive information remains secure.

In conclusion, our findings suggest that word-level encryption can be a viable approach for privacy-preserving NLP tasks. However, it is crucial to consider the potential risks of reverse engineering and focus on tasks that return non-sensitive outputs. By exploring more advanced models like transformers and investigating innovative solutions to enhance data privacy, we can pave the way for secure and privacy-preserving NLP applications in various industries.

### References

1. Goldberg, Y. (2016). A Primer on Neural Network Models for Natural Language Processing. *Journal of Artificial Intelligence Research*, 57, 345-420.
2. Ozan, Ş. ve Taşar, D.E. (2021). Auto-tagging of Short Conversational Sentences Using Natural Language Processing Methods. <http://arxiv.org/abs/2106.04959>.
3. Li, M., Yu, S., Zheng, Y., Ren, K., & Lou, W. (2010). Scalable and secure sharing of personal health records in cloud computing using attribute-based encryption. *IEEE Transactions on Parallel and Distributed Systems*, 24(1), 131-143.
4. Hashem, I. A. T., Yaqoob, I., Anuar, N. B., Mokhtar, S., Gani, A., & Ullah Khan, S. (2015). The rise of big data on cloud computing: Review and open research issues. *Information Systems*, 47, 98-115.
5. Pearson, S., & Benameur, A. (2010). Privacy, Security, and Trust Issues Arising from Cloud Computing. In *2010 IEEE Second International Conference on Cloud Computing Technology and Science* (pp. 693-702). IEEE.
6. Sun, D., Chang, G., Sun, L., & Wang, X. (2014). Surveying and Analyzing Security, Privacy, and Trust Issues in Cloud Computing Environments. *Procedia Engineering*, 29, 285-290.

7. Humbert, M., Ayday, E., Hubaux, J. P., & Telenti, A. (2013). Addressing the Concerns of the Lacks Family: Quantification of Kin Genomic Privacy. In Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security (pp. 1141-1152).
8. Sweeney, L. (2002). k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness, and Knowledge-Based Systems*, 10(5), 557-570.
9. Gentry, C. (2009). A fully homomorphic encryption scheme (Doctoral dissertation, Stanford University).
10. Brakerski, Z., & Vaikuntanathan, V. (2011). Efficient Fully Homomorphic Encryption from (Standard) LWE. In 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science (pp. 97-106). IEEE.
11. Acar, A., Aksu, H., Uluagac, A. S., & Conti, M. (2018). A survey on homomorphic encryption schemes: Theory and implementation. *ACM Computing Surveys (CSUR)*, 51(4), 1-35.
12. Chen, F., Aono, Y., Hayashi, T., Wang, Q., & Sasaki, R. (2017). A review of homomorphic encryption for secure computation. *IEICE Transactions on Fundamentals of Electronics, Communications, and Computer Sciences*, 100(1), 250-261.
13. Aslett, L. J., Esperança, P. M., & Holmes, C. C. (2015). A review of homomorphic encryption and software tools for encrypted statistical machine learning. *arXiv preprint arXiv:1508.06574*.
14. Chillotti, I., Gama, N., Georgieva, M., & Izabachène, M. (2018). Faster Homomorphic Encryption Is Not Enough: Improved Heuristic for Multiplicative Depth Minimization of Boolean Circuits. In *International Conference on Cryptology in India* (pp. 3-22). Springer, Cham.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.