International Journal of Education, Vocational and Social Science
https://e-journal.citakonsulindo.or.id/index.php/IJEVSS
E-ISSN : 2963 4954
IJEVSS

# Performance Analysis of IndoBERT Fine-Tuning for Indonesian Text-Based Hoax Detection

**Tiara Siti Nadira**
Universitas Bina Darma, Indonesia
**Ari Muzakir**
Universitas Bina Darma, Indonesia

Corresponding Author: tiarasitinadira@gmail.com, arimuzakir@binadarma.ac.id

**Abstract**

*The spread of misinformation (hoaxes) on social media, particularly on the X platform, has become a significant challenge that can influence public understanding and perceptions of various issues, including those related to education. This study aims to analyze and improve the performance of an Indonesian-language hoax detection model through a fine-tuning approach using the IndoBERT model. The dataset used in this research was collected through web scraping from the X platform and classified into two categories: hoax and non-hoax. Prior to model training, the data underwent several preprocessing stages, including data cleansing, normalization, tokenization, stopword removal, and stemming.*

*This study utilized the pre-trained IndoBERTweet model with two data-splitting scenarios, namely 70:30 and 80:20, along with various combinations of learning rates and epochs to determine the optimal configuration. Model performance was evaluated using accuracy, precision, recall, and F1-score metrics. The experimental results indicate that the best performance was achieved under the 70:30 data-splitting scenario with a learning rate of 2e-5 and 20 epochs. On the testing dataset, the model achieved an accuracy of 80%, precision of 83%, recall of 78%, and an F1-score of 80%.*

*The findings demonstrate that the fine-tuning approach applied to IndoBERT is effective in enhancing the performance of Indonesian-language hoax classification and is capable of maintaining strong performance when evaluated on previously unseen data.*

**Keywords: IndoBERT, IndoBERTweet, *fine-tuning*, hoax detection, social media, natural language processing (NLP), text classification.**



Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-Share Alike 4.0 International License

International Journal of Education, Vocational and Social Science
https://e-journal.citakonsulindo.or.id/index.php/IJEVSS
E-ISSN : 2963 4954
IJEVSS

## INTRODUCTION

The rapid advancement of information and communication technology has fundamentally transformed the way people access, produce, and disseminate information. Social media has become one of the primary channels for information exchange due to its ability to reach a broad audience instantly. However, this convenience has also accelerated the spread of misinformation and fake news (hoaxes). The widespread dissemination of hoaxes has significant societal consequences, including the formation of misleading public opinions, increased social polarization, and declining trust in credible information sources. Consequently, automatic fake news detection has become an important research topic in the field of Natural Language Processing (NLP).

In Indonesia, X (formerly Twitter) is one of the most widely used social media platforms for sharing opinions, news, and discussions on various topics, including education, politics, health, and social issues. The platform's textual characteristics, such as short messages, informal language, abbreviations, slang, sarcasm, and code-mixing, make the task of identifying fake news particularly challenging. Furthermore, information is often propagated through complex conversation threads, making manual verification inefficient and highlighting the need for automated systems capable of understanding linguistic context more effectively.

Numerous approaches have been proposed for text-based fake news detection, ranging from traditional machine learning algorithms such as Naïve Bayes, Support Vector Machine (SVM), and Random Forest to deep learning models including Convolutional Neural Networks (CNN), Long Short-Term Memory (LSTM), and Gated Recurrent Units (GRU). Although these methods have demonstrated improvements over rule-based approaches, many still struggle to capture complex semantic relationships and contextual information, particularly in Indonesian texts, which exhibit high linguistic diversity and informal writing styles.

Recent advances in Transformer-based architectures have significantly improved natural language understanding. Among these models, Bidirectional Encoder Representations from Transformers (BERT) has achieved state-of-the-art performance across various NLP tasks. For Indonesian language processing, IndoBERT was pretrained on a large-scale Indonesian corpus, enabling it to learn richer linguistic representations than previous deep learning architectures. To achieve optimal performance on domain-specific tasks such as fake news detection, IndoBERT must be adapted through a fine-tuning process using datasets that reflect the characteristics of the target domain. In addition to the pretrained model itself, the selection of hyperparameters—including learning rate, number of training epochs, and data splitting ratio—plays a crucial role in



Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-Share Alike 4.0 International License

determining the model's classification performance.

Despite the growing number of studies employing Transformer-based models for fake news detection, most existing research evaluates only a single hyperparameter configuration, providing limited insight into how different fine-tuning settings influence model performance. To address this gap, this study investigates the performance of fine-tuned IndoBERT for Indonesian fake news detection by systematically evaluating various combinations of learning rates, training epochs, and train-test split ratios. The model performance is assessed using accuracy, precision, recall, and F1-score. In addition, this study contributes a newly constructed Indonesian fake news dataset focused on educational topics collected from X, which reflects the characteristics of informal language commonly found on social media. The findings are expected to provide practical guidance for optimizing Transformer-based models and contribute to the development of more accurate and robust fake news detection systems for the Indonesian language.

## Literature Review

### Overview of Hoax Detection

The rapid growth of social media has significantly increased the dissemination of hoaxes, making misinformation one of the major challenges in the digital era. Hoaxes are intentionally misleading or false information that can influence public opinion, create social conflicts, and reduce trust in credible information sources (Allcott & Gentzkow, 2017). On social media platforms, hoaxes often employ provocative language, unreliable sources, and emotionally charged content, allowing them to spread rapidly among users (Pranesti & Arifin, 2019). Therefore, automatic hoax detection has become an important research area in Natural Language Processing (NLP) to support faster and more accurate identification of misleading information.

### X as a Source of Hoax Dissemination

X (formerly Twitter) is a microblogging platform where users share short textual messages in real time. Due to its fast information exchange and high user engagement, X has become one of the primary channels for spreading misinformation. Tweets frequently contain informal language, abbreviations, slang, hashtags, and code-mixed expressions, creating additional challenges for text classification models. These characteristics make X an appropriate data source for evaluating Indonesian hoax detection systems.


Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-Share Alike 4.0 International License

**Transformer-Based Language Models**

Transformer has become the dominant architecture in Natural Language Processing because it employs a self-attention mechanism to capture contextual relationships between words without relying on recurrent structures. Unlike conventional machine learning and recurrent neural networks, Transformer models can learn long-range semantic dependencies more effectively while enabling parallel computation, resulting in better performance across various text classification tasks (Devlin et al., 2019).

**IndoBERT and Fine-Tuning Strategy**

IndoBERT is a Transformer-based language model pretrained on large-scale Indonesian corpora, enabling it to understand the linguistic characteristics of the Indonesian language more effectively. To improve performance on specific downstream tasks, the pretrained model is adapted through a fine-tuning process using task-specific datasets. During fine-tuning, hyperparameters such as learning rate, batch size, and training epochs significantly influence model performance by affecting convergence, generalization, and the risk of overfitting. Therefore, selecting appropriate hyperparameter combinations is essential for obtaining an optimal hoax detection model.

**Text Preprocessing**

Text preprocessing aims to transform raw textual data into a structured representation suitable for model training. The preprocessing pipeline typically includes data cleansing, normalization, tokenization, stopword removal, and stemming. These steps reduce textual noise, standardize word representations, and improve the quality of input features, thereby enhancing the effectiveness of Transformer-based language models.

**Performance Evaluation**

The performance of the proposed model is evaluated using a confusion matrix consisting of True Positive (TP), True Negative (TN), False Positive (FP), and False Negative (FN). Based on these values, evaluation metrics including Accuracy, Precision, Recall, and F1-score are calculated to comprehensively assess the effectiveness of the model in distinguishing hoax and non-hoax texts.

**Related Studies**

Previous studies have demonstrated the effectiveness of Transformer-based models for



Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-Share Alike 4.0 International License

Indonesian hoax detection. Sinapoy et al. (2023) reported that IndoBERT outperformed LSTM with an accuracy of 92.07%. Tobing et al. (2025) further improved political hoax detection using fine-tuned IndoBERT, achieving 94.1% accuracy. Meanwhile, Nuryanto (2025) showed that IndoBERT consistently outperformed conventional machine learning models, including Support Vector Machine and Naïve Bayes. Although these studies confirmed the superiority of IndoBERT, most of them employed a single fine-tuning configuration. Therefore, a comprehensive analysis of hyperparameter combinations, including learning rate, training epochs, and train-test split ratios, remains limited. This study addresses this gap by systematically evaluating multiple hyperparameter configurations to identify the optimal settings for Indonesian hoax detection.

## METHODS

This study employed an experimental research design to investigate the effect of hyperparameter optimization on the performance of the IndoBERT model for Indonesian hoax news classification. The research workflow consisted of six sequential stages, namely data collection, data annotation, text preprocessing, dataset partitioning, model fine-tuning, and performance evaluation. Each stage was conducted systematically to ensure that the developed classification model achieved optimal performance and could be evaluated objectively.

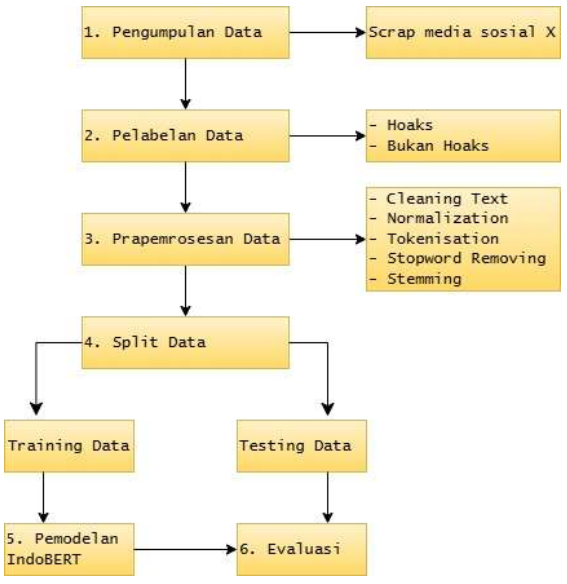



**Figure 1. Research Framework**



Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-Share Alike 4.0 International License

The dataset used in this research consisted of Indonesian-language tweets collected from the social media platform X (formerly Twitter) using Python-based web scraping techniques. Data collection focused on educational issues, including tuition fees (UKT), educational policies, scholarships, and other topics that potentially contained misinformation. A total of 11,400 tweets were successfully collected and stored in CSV format. Subsequently, each tweet was manually annotated into Hoax and Non-Hoax categories through a multi-stage labeling process involving the researcher, a language expert, and an independent validator. Fact verification was performed using reliable information sources to ensure annotation quality and consistency.

Prior to model training, all collected tweets underwent a preprocessing stage to improve text quality and reduce noise commonly found in social media data. The preprocessing pipeline included data cleansing, text normalization, tokenization, stop-word removal, and stemming using Indonesian language processing tools. These preprocessing steps transformed informal social media text into a standardized representation suitable for transformer-based language models while preserving the semantic meaning of the original content. To evaluate the influence of dataset composition on model performance, two data partition scenarios were implemented. The first scenario divided the dataset into 80% training data and 20% testing data**,** while the second scenario used 70% training data and 30% testing data. In both scenarios, 10% of the training data was reserved as validation data during the fine-tuning process to monitor model convergence and minimize overfitting.

The classification model was developed using the pre-trained indolem/indobertweet-base-uncased model implemented through the HuggingFace Transformers library. This model was selected because it had been pre-trained on Indonesian Twitter data, making it particularly suitable for analyzing informal Indonesian text. Fine-tuning was performed by evaluating multiple combinations of learning rates and training epochs while maintaining a fixed batch size of 16. Three learning rates ($2\times10^{-5}$, $3\times10^{-5}$, and $5\times10^{-5}$) and six epoch configurations (2, 3, 4, 10, 20, and 30 epochs) were investigated, resulting in **36 experimental configurations**. The experiments were conducted using Python 3.12 in the Google Colab environment with PyTorch and the HuggingFace Transformers framework.

**Table 1. Experimental Configuration**

| **Parameter** | **Configuration** |
|---|---|
| Pre-trained Model | indolem/indobertweet-base-uncased |
| Programming Language | Python 3.12 |



Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-Share Alike 4.0 International License

| Parameter | Configuration |
|---|---|
| Framework | PyTorch |
| NLP Library | HuggingFace Transformers |
| Platform | Google Colab |
| Batch Size | 16 |
| Learning Rates | 2e-5, 3e-5, 5e-5 |
| Epochs | 2, 3, 4, 10, 20, 30 |
| Data Partition | 80:20 and 70:30 |
| Validation Data | 10% of training data |
| Number of Experiments | 36 |

Model performance was evaluated using a confusion matrix and four standard classification metrics, namely accuracy, precision, recall, and F1-score**. In** addition, training loss and validation loss were monitored throughout the fine-tuning process to analyze convergence behavior and detect overfitting. The optimal model configuration was determined based on the highest classification performance while maintaining stable convergence and good generalization capability.

## DISCUSSION

The present study evaluated the influence of data partitioning strategies, learning rates, and training epochs on the performance of the IndoBERT model for Indonesian hoax news classification. A total of 36 experimental configurations were evaluated using accuracy, precision, recall, F1-score, training loss, and validation loss to identify the optimal hyperparameter configuration. The findings indicate that model performance is influenced not only by the amount of training data but also by the interaction among hyperparameter settings during the fine-tuning process.

**Table 2. Performance of the Best IndoBERT Model**

| Data Split | Batch Size | Learning Rate | Epoch | Accuracy | Precision | Recall | F1-score |
|---|---|---|---|---|---|---|---|
| **70:30** | 16 | 2e-5 | 4 | 0.80 | 0.79 | 0.82 | 0.80 |

Table 1 shows that the best classification performance was obtained using a 70:30 data


Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-Share Alike 4.0 International License

partition with a learning rate of 2e-5 and four training epochs. This configuration produced balanced evaluation metrics, indicating that the model was capable of distinguishing hoax and non-hoax news without favoring a particular class. Compared with other configurations, this model also maintained relatively low training loss and controlled validation loss, suggesting good convergence and generalization capability.

The comparison between both data partition scenarios demonstrates that the 70:30 split consistently outperformed the 80:20 configuration after excluding anomalous experiments affected by convergence failure. Although the numerical difference in average accuracy was relatively small, the 70:30 configuration produced more balanced precision and recall values, indicating a more stable classification process. These findings suggest that increasing the proportion of training data alone does not necessarily improve model performance. Instead, an appropriate balance between training and testing data contributes to a more representative evaluation of model generalization.

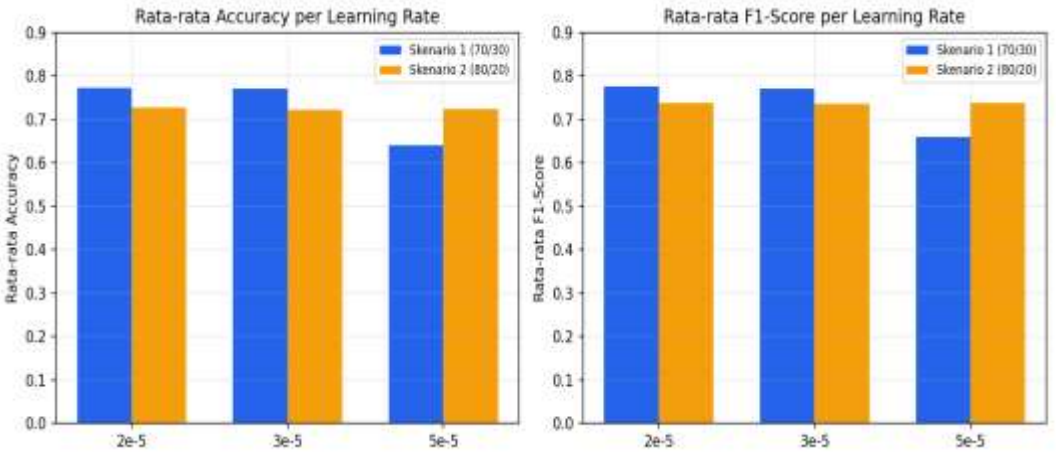



**Figure 1. Comparison of Average Accuracy for Different Learning Rates**

Figure 1 illustrates the influence of learning rate on model performance. Learning rates of 2e-5 and 3e-5 consistently achieved higher average accuracy than 5e-5. The largest learning rate occasionally caused unstable optimization, particularly under the 70:30 data partition, where the model experienced class collapse. In this condition, recall reached 1.00 while precision and accuracy decreased substantially because the model predicted almost all samples as belonging to a single class. These findings confirm that learning rate plays a critical role in determining optimization stability during transformer fine-tuning.



Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-Share Alike 4.0 International License

The analysis of training epochs further revealed that extending the training process beyond four epochs did not improve classification performance. Instead, additional epochs gradually increased validation loss while training loss continued to decrease, indicating that the model became increasingly specialized to the training data and gradually lost its ability to generalize to unseen samples. Therefore, the optimal number of training epochs for this dataset was found to be four.

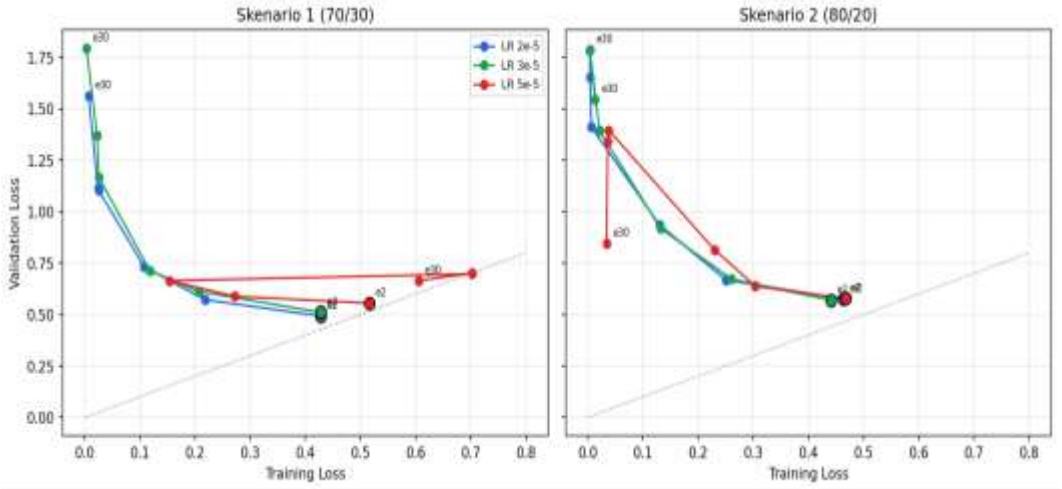



**Figure 2. Training Loss and Validation Loss during Fine-Tuning**

Figure 2 demonstrates the learning trajectory of the IndoBERT model throughout the fine-tuning process. At the initial epochs, training loss and validation loss decreased simultaneously, indicating successful learning. However, after the fourth epoch, validation loss gradually increased despite the continuous reduction of training loss. This divergence represents the onset of overfitting, where the model memorizes training patterns instead of learning features that generalize well to unseen data. Consequently, monitoring validation loss provides a more reliable indication of model quality than relying solely on classification accuracy.

Overall, the experimental findings demonstrate that successful IndoBERT fine-tuning requires balancing optimization stability and model generalization. The combination of a 70:30 data partition, learning rate of 2e-5, batch size of 16, and four epochs provided the most stable convergence process while achieving the highest classification performance. These findings indicate that careful hyperparameter optimization is more important than simply increasing the amount of training data or extending training duration.



Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-Share Alike 4.0 International License

## CONCLUSION

This study concludes that the performance of IndoBERT for Indonesian hoax news classification is strongly influenced by the interaction among data partitioning, learning rate, and training epochs. The optimal configuration consisted of a 70:30 data split, batch size of 16, learning rate of 2e-5, and four training epochs, achieving an accuracy of 0.80, precision of 0.83, recall of 0.78, and an F1-score of 0.80.

The experimental results also demonstrate that increasing the amount of training data or extending the training process does not necessarily improve model performance. Instead, appropriate hyperparameter selection contributes more significantly to optimization stability, improved generalization capability, and reduced overfitting. Therefore, systematic hyperparameter optimization should be considered an essential component of transformer-based text classification research.

## FUTURE DIRECTIONS

Future research should investigate additional optimization strategies, including weight decay, warm-up scheduling, batch size variation, and learning rate schedulers to further improve IndoBERT performance. Evaluating alternative transformer architectures such as IndoRoBERTa or multilingual transformer models using the same dataset would also provide a comprehensive comparison of language representation models for Indonesian hoax detection.

Furthermore, future studies should incorporate larger and more diverse news datasets collected from multiple digital platforms to improve model robustness and generalization. The integration of explainable artificial intelligence (XAI) techniques may also enhance the interpretability of prediction results, enabling a better understanding of the linguistic features contributing to hoax classification.

## REFERENCES

Alkhodair, S. A., Ding, S. H. H., Fung, B. C. M., & Liu, J. (2020). Detecting fake news in social media networks. Computers in Human Behavior, 104

Antariksa, M. D. S., Sugiharto, A., & Surarso, B. (2025). BERT model fine-tuned for scientific document classification and recommendation. Jurnal RESTI (Rekayasa Sistem dan Teknologi Informasi), 9(4), 754–764. https://doi.org/10.29207/resti.v9i4.6789

Ardakani, A., Haan, A., Tan, S., Popovici, D. T., Cheung, A., Iancu, C., & Sen, K. (2023). SlimFit: Memory-efficient fine-tuning of transformer-based models using training dynamics.



Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-Share Alike 4.0 International License

Basahel, A. M., Giriyappa, S. H., Alam, F., Alnazzawi, T. S. M., Qamar, S., & Abi Sen, A. A. (2026a). A resource-efficient approach to fine-tuning a BERT-base model for sentiment analysis. Computers, 15(3), 159. https://doi.org/10.3390/computers15030159

Basahel, A. M., Giriyappa, S. H., Alam, F., Alnazzawi, T. S. M., Qamar, S., & Abi Sen, A. A. (2026b). A resource-efficient approach to fine-tuning a BERT-base model for sentiment analysis. Computers, 15(3), 159. https://doi.org/10.3390/computers15030159

De Diego, I. M., Redondo, A. R., Fernández, R. R., Navarro, J., & Moguerza, J. M. (2022). General performance score for classification problems. Applied Intelligence, 52(10), 12049–12063. https://doi.org/10.1007/s10489-021-03041-7

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding.

Ding, N., Qin, Y., Yang, G., Wei, F., Yang, Z., Su, Y., Hu, S., Chen, Y., Chan, C.-M., Chen, W., Yi, J., Zhao, W., Wang, X., Liu, Z., Zheng, H.-T., Chen, J., Liu, Y., Tang, J., Li, J., & Sun, M. (2023). Parameter-efficient fine-tuning of large-scale pre-trained language models. Nature Machine Intelligence, 5(3), 220–235. https://doi.org/10.1038/s42256-023-00626-4

Düntsch, I., & Gediga, G. (2019). Confusion matrices and rough set data analysis. Journal of Physics: Conference Series, 1229(1), 012055. https://doi.org/10.1088/1742-6596/1229/1/012055

Gorunescu, F. (2011). Data mining. Springer.

Hua, H., Li, X., Dou, D., Xu, C., & Luo, J. (2021). Noise stability regularization for improving BERT fine-tuning. Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, 3229–3241. https://doi.org/10.18653/v1/2021.naacl-main.258

Isadora, F. R., Hanggara, B. T., & Mursityo, Y. T. (2021). User experience design for the HomeCare mobile application of Semen Gresik Hospital using the Design Thinking method. Jurnal Teknologi Informasi dan Ilmu Komputer, 8(5), 1057–1066. https://doi.org/10.25126/jtiik.2021844550

Jocelynne, C., Tobing, L., Wijayakusuma, I. L., Harini, L. P., & Harini, I. (2025). Detection of political hoax news using fine-tuning IndoBERT. Journal of Applied Informatics and Computing (JAIC), 9(2).

Kiran, S., Kumar, A. A., Reddy, R. P. K., & Reddy, D. S. (2022). Analysis of in-place matrix rotation of square matrix for information security applications. Procedia Computer Science, 215, 131–139. https://doi.org/10.1016/j.procs.2022.12.015



Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-Share Alike 4.0 International License

Koto, F., Rahimi, A., Lau, J. H., & Baldwin, T. (2020). IndoLEM and IndoBERT: A benchmark dataset and pre-trained language model for Indonesian NLP. Proceedings of the 28th International Conference on Computational Linguistics, 757–770. https://doi.org/10.18653/v1/2020.coling-main.66

Koto, F., Lau, J. H., & Baldwin, T. (n.d.). IndoBERTweet: A pretrained language model for Indonesian Twitter with effective domain-specific vocabulary initialization. https://huggingface.co/huseinzol05/

Li, H., Chaudhari, P., Yang, H., Lam, M., Ravichandran, A., Bhotika, R., & Soatto, S. (2020). Rethinking the hyperparameters for fine-tuning.

Lin, T., Wang, Y., Liu, X., & Qiu, X. (2022). A survey of transformers. AI Open, 3, 111–132. https://doi.org/10.1016/j.aiopen.2022.10.001

Lu, W., Luu, R. K., & Buehler, M. J. (2025). Fine-tuning large language models for domain adaptation: Exploration of training strategies, scaling, model merging and synergistic capabilities. NPJ Computational Materials, 11(1), 84. https://doi.org/10.1038/s41524-025-01564-y

Mahendra, M., & Triandi, R. (n.d.). Analisa dan deteksi konten hoax pada media berita Indonesia menggunakan machine learning.

Mosbach, M., Andriushchenko, M., & Klakow, D. (2021a). On the stability of fine-tuning BERT: Misconceptions, explanations, and strong baselines.

Mosbach, M., Andriushchenko, M., & Klakow, D. (2021b). On the stability of fine-tuning BERT: Misconceptions, explanations, and strong baselines.

Nuryanto, N. (2025). Prosiding Seminar Nasional SI & Teknik Informatika (pp. 134–140).

Oktaviana Yamin, V., Tenriawaru, A., Saidi, L. O., Rahman, G. A., et al. (2023). Penerapan Naïve Bayes classifier dengan algoritma Nazief dan Adriani untuk deteksi hoaks. Journal of Information Technology and Computer Science, 1(1).

Pranesti, D., & Arifin. (2019). Perlindungan korban dalam kasus penyebaran berita hoax di media sosial di Indonesia. Jurnal Hukum Media Bhakti, 3(1), 8–17.

Priyanka Arora, & Arora, P. (2019). Mining Twitter data for depression detection. IEEE, 186–189.

Purohit, H., Hampton, A., Shalin, V. L., Sheth, A. P., Flach, J., & Bhatt, S. (2013). What kind of #conversation is Twitter? Mining psycholinguistic cues for emergency coordination. Computers in Human Behavior, 29.



Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-Share Alike 4.0 International License

Raghu, D., Gupta, N., & Mausam. (2021). Unsupervised learning of KB queries in task-oriented dialogs. Transactions of the Association for Computational Linguistics, 9, 374–390. https://doi.org/10.1162/tacl_a_00372

Rahayu, R. N. (2020). Analisis berita hoax COVID-19 di media sosial di Indonesia. Jurnal Ekonomi, Sosial & Humaniora, 1(9), 60–73.

Raza, M. N. (n.d.). Sistem deteksi berita hoax menggunakan algoritma Naïve Bayes dan Random Forest pada machine learning.

Riyaddulloh, R., & Romadhony, A. (2021). Normalisasi teks bahasa Indonesia berbasis kamus slang: Studi kasus tweet produk gadget pada Twitter. Jurnal Teknologi Informasi, 8(4).

Sinapoy, M. I. K., Sibaroni, Y., & Prasetyowati, S. S. (2023). Comparison of LSTM and IndoBERT method in identifying hoax on Twitter. Jurnal RESTI, 7(3), 657–662. https://doi.org/10.29207/resti.v7i3.4830

Tay, Y., Dehghani, M., Bahri, D., & Metzler, D. (2023). Efficient transformers: A survey. ACM Computing Surveys, 55(6), 1–28. https://doi.org/10.1145/3530811

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. Science, 359, 1146–1151.

Wang, L., Chen, S., Jiang, L., Pan, S., Cai, R., Yang, S., & Yang, F. (2025). Parameter-efficient fine-tuning in large language models: A survey of methodologies. Artificial Intelligence Review, 58(8), 227. https://doi.org/10.1007/s10462-025-11236-4

Wang, W., & Pan, S. J. (2021). Variational deep logic network for joint inference of entities and relations. Computational Linguistics, 47(4), 775–812. https://doi.org/10.1162/coli_a_00415

Zaidan, M. N., Sibaroni, Y., & Prasetyowati, S. S. (2024a). Learning rate and epoch optimization in the fine-tuning process for IndoBERT's performance on sentiment analysis of MyTelkomsel app reviews. Jurnal Teknik Informatika (JUTIF), 5(5), 1443–1450.

Zaidan, M. N., Sibaroni, Y., & Prasetyowati, S. S. (2024b). Learning rate and epoch optimization in the fine-tuning process for IndoBERT's performance on sentiment analysis of MyTelkomsel app reviews. Jurnal Teknik Informatika (JUTIF), 5(5), 1443–1450.

Zhang, C., Liu, Z., & Sun, M. (2018). Enhancing Naïve Bayes with contextual embeddings for text classification tasks. Proceedings of the Association for Computational Linguistics (ACL), 10(2), 145–153.



Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-Share Alike 4.0 International License