

Exploring Machine Learning Techniques for Text Classification: A Survey

Charushila D. Patil¹ and Atul Agrawal²

^{1,2}Oriental University, Indore, India

Email: Charushila.patil@ggsf.edu.in, atul.273@gmail.com

Abstract— The field of Natural Language Processing (NLP) has seen significant advancements through the adoption of machine learning techniques for text classification, an essential task within the domain. In this paper, we present a survey of machine learning approaches to text classification summarising their developments, methodologies and applications while outlining the linguistic challenges addressed by these techniques. It covers major architectures such as Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM), and state of the art transformer-based models such as BERT and GPT. We systematically discuss the advantages and disadvantages of these techniques with respect to contextual semantics handling, scalability to large-scale datasets, and domain-specific performance. Moreover, this survey explores recent trends such as hybrid models, attention mechanisms, and multimodal approaches that improve accuracy and robustness in classification. We discuss the challenges related to computational overhead, data imbalance, and interpretability; possible remedies are proposed, and future research paths are outlined. The iceberg of endless possibilities grows deep as we dive deeper through the tutorials covering real-world use cases ranging from sentiment analysis regardless of whether your data is a family story or news articles, spam detection of whether an email is real or not, healthcare and legal document classification with deep learning impacting the world in a myriad of ways. This paper provides an overview over these recent advances, making it a valuable resource for each of researchers and practitioners alike, and gives a comprehensive understanding of machine learning based text classification.

Index Terms— Machine Learning, Text Classification, Natural Language Processing (NLP), Neural Networks, Transformer Models, Hybrid Approaches.

I. INTRODUCTION

Text Classification is a common task in Natural Language Processing (NLP) which involves classifying the text into predefined categories or classes. It encompasses everything from spam detection in emails to sentiment analysis in social media, and is a critical component of how we decipher thousands upon millions (upon billions) of textual data. Earlier techniques for text classification involved hand-crafted features combined with classical machine learning algorithms (e.g. SVM, Naive Bayes). Our creations are showing promise, but break when they encounter big data or because we have more complex poses that require a deeper semantic understanding of the context. The emerging presence of digital communication has resulted in numerous textual data explosion across various sectors like healthcare, social media, e-commerce and finance.

Correctly cataloging and sorting this data is essential for both businesses matters as well as research. But in reality, the diversity and ambiguity of human languages with condensed or slang words and evolving language make this task intrinsically hard. Advanced solutions using deep learning and NLP have come to the fore recently, which are a more robust scalable system for tackling challenging text classification problems.

Deep learning has been shown to transform text classification, as models are able to learn hierarchical representations of the raw textual data automatically. Contrary to the basic machine learning methods which need to be supplied predefined features achieved through feature engineering, deep Learning Models are capable of getting trained over raw text data. This feature and the capacity to model complicated patterns in deep learning are automatically extracted, which is why text classification is currently considered state-of-the-art. The key revolution in deep learning for text classification is the use of recurrence RNNs (Recurrent Neural Networks), more notably LSTM and GRU-based networks. Since these models are able to model a sequence and learn dependencies during later/earlier parts of the input, they work exceptionally well in tasks with text. Essentially, LSTM is better at remembering long-distance features with its specific memory cell and gate than the basic GRU proposal that reduces it to a simpler internal structure.

This survey paper comprehensively examines the role of deep learning in advancing text classification. The major contributions are as follows:

- A detailed exploration of state-of-the-art deep learning techniques for text classification, including their architectures, methodologies, and applications.
- Discussion on the strengths and weaknesses of prominent deep learning models such as CNNs, RNNs, LSTM, and Transformer-based models like BERT and GPT.
- Insights into recent developments, including hybrid models, attention mechanisms, and multimodal approaches.
- Identification of key challenges such as data imbalance and computational complexity, along with potential future research directions.
- Highlighting real-world use cases across various domains such as sentiment analysis, spam detection, and healthcare.

The remainder of this paper is organized as follows: Section 2 reviews various deep learning architectures and their methodologies in detail. Section 3 provides a comparative analysis of different models, highlighting their performance and limitations. Section 4 discusses emerging trends, challenges, and opportunities in the field of deep learning for text classification. Section 5 explores practical applications across diverse domains. Section 6 concludes the paper with a summary of findings and potential directions for future research.

II. LITERATURE REVIEW

This paper [1] investigates the use of deep learning methods for text-wise categorization in engineering consulting. Authors suggests a text classification model that combines domain-specific standards with deep-learning techniques. This research provides an example of standardizing text classification, or at least among a few stages within the same whole-process engineering consulting for knowledge extraction in different project phases. The results of this paper illustrate that deep learning is useful in processing complicated engineering texts, and raise the accuracy on classification.

This paper [2] gives a good survey of the different machine learning algorithms used for text classification. From traditional methods like SVM and Naive Bayes, deep leaning models are also described by them. This review includes several challenges in text classification like feature extraction paradigm, choice of model and dealing with class imbalances. This paper additionally benchmarks the methods on different datasets and provides hints on their potential applications. The review [3] covers a lot of deep learning approaches used for text classification. What Models are used: Different type of model architecture has been discussed such as CNNs, RNNs (including LSTM and GRU), transformer-based models like BERT, GPT etc. For sentiment analysis, document and term-based approaches are used to address popular issues in this field as well as for topic classification; sample oversampling can be applied. A representative task on a larger scale domain is spam detection, some solutions using ensemble approach with high performance results were proposed and we list all of those which may make sense in real case scenario by their specific tasks (document or term level). This review: (i) presents a comprehensive discussion about existing neural networks, symbolic approaches and ensemble techniques that can be used for the task of text classification; in addition to proposing new metrics such as Multi-usefulness and irrelevancies(c); per class scores, (ii), 2D matrix correlation graphs(iii).

A Survey [4] on Deep Learning for Text Classification —The article discusses the advancements of text classification models from traditional machine learning to recent state-of-art deep learning (transformers)

methods. In this paper, we present a brief survey of several key developments including attention mechanisms for text classification (Sec 2), how they have been used to improve transfer learning in NLP implementations (Sec 3) and recently introduced pre-trained language models (PLM) that build on top of most recent ideas from both the previous points. Moreover, the work endeavours to address future avenues in this field: examples of which are Hybrid Models and Cross-Lingual Text Classification. In this paper [5], a full tutorial on how to build real-world text classification systems leveraging the state-of-the-art NLP model is provided. The pipeline from data pre-processing, through feature extraction to the training and evaluation of models is described in detail by Dozat et al. The paper elaborates the application of state-of-the-art NLP models as BERT, RoBERTa and GPT-2 for different text classification problems. The paper also addresses aspects of model deployment, like scalability, interpretability and computational efficiency. In this paper [6], we explore the application of hybrid deep learning models to classify medical texts. In this paper, the authors present such a model built with CNNs and using BiLSTMs to encode temporal sequence information at both word and character levels by feeding rich low-variance features extracted from medical documents into multitask classification; they also introduce three types of attention mechanisms: global self-attention, local context-aware multi-head attention (LCMHA), entity level latent component shared on multiple tasks. The work deals with challenges consistently present in medical text, such as the use of domain specific jargon and an inherently small amount of training data. Extensive experiments demonstrate the superiority of our hybrid model over both traditional methods and deep learning models in classification accuracy as well as robustness. In this study[7], the characteristic of article classification schema is dealt with using NLP and machine learning methods. The author compares the performance of different machine learning algorithms (e.g. Decision Trees, SVM, KNN) over text features using TF-IDF and word embeddings techniques. But as the research shows, both feature selection and model optimization will need to be carried out in order to optimize classification accuracy. The performance of hybrid modes were evaluated using combination various feature extraction techniques and algorithms, the results indicate that mixture approaches ensure better performance as a single method. In this paper[8], we propose a deep learning framework for disease prediction from clinical notes and medical documents by using biomedical Text Classification. The authors introduce an integrated architecture combining CNNs and LSTM layers with NLP techniques to pull out the necessary medical information for classification. Results suggest the application of deep learning could lead to more accurate health prediction and the elimination or significant reduction of a diagnostic category-manual burden in disease classification among medical institutions.

This paper [9] presents a new frame work for text classification that use evolutionary algorithms and swarm based optimization methods along with deep neural networks. The proposed model integrates stack of CNN layers and swarm optimized deep learning structure to improve feature learning as well the classification accuracy. It achieves state of the art results on multiple text classification tasks, with orders of magnitude fewer parameters — an order to two less than typical deep learning classifiers; which makes it particularly useful for smaller datasets. In this paper [10], we present an efficient and simple end-to-end deep learning architecture for text classification that sacrifices a little bit of accuracy in exchange for computational efficiency. A Lightweight Model Architecture Combining CNN and Attention Mechanisms Optimized for Real-Time Classification Task (Task 12) Takeaways The focus of this study is on lightweight model if we need to achieve high speed and low latency which are convenient for online applications like Content moderation, Spam detection etc. We reviewed and summarised how deep-learning technologies have been applied in text classification, within a social health care network setting.

The authors [11] investigate multiple NLP models transpiring from RNNs to CNNs and transformers for healthcare-based applications like disease surveillance, mental health assessment, patient symptom feedback categorisation etc. The paper also notes the challenges of using deep learning in healthcare including data privacy and model interpretability. This paper [12] presents a study on the problem of text classification for medical domain - particularly when applied to highly skewed unstructured clinical notes. In particular the authors propose measures for reducing imbalanced classes (e.g., by using weighted loss functions and data augmentation) which they evaluate based on classification performance. These results demonstrate that class imbalance is one of the key challenges for generalizing accurate models in medical applications, especially when dealing with rare but critical diagnoses. In this paper[13], we propose a global model for weakly supervised deep learning representations of clinical text and develop it into a novel method to predict whether patients have certain diseases based on their document type. Sparse training data is generated through weak labelling, and meaningful features are learned from clinical texts with the help of pre-trained deep learning models. Results show that this method is very effective in transferring labels without taking recourse to heavy and time-consuming manual labelings.

This research paper [14] identifies how good machine learning and deep learning models work for the assignment of text classification tasks. To this end, the authors compare deep learning models like LSTM and CNNs to traditional classifiers such as SVM and Naive Bayes on multiple text datasets. It puts the advantages in perspective when deep learning is concerned with complex language pattern recognition while also discussing some of those instances where simple traditional machine learning models seem to outperform due to their simplicity and less computational needs. The integration of NLP with machine learning algorithms for text classification has been explored in this paper [15]. In this regard, the authors suggest a hybrid model that uses word embeddings as features after selecting important ones and reinforces such words with machine learning algorithms like Random Forest or SVM.

In this paper, a study on text categorization in mobile communication scenarios is described to solve these problems by limited computing resources and lightweight models. These results show that the hybrid system can provide good classification accuracy without losing efficiency.

III. COMPARATIVE ANALYSIS OF DEEP LEARNING MODELS FOR TEXT CLASSIFICATION

This section briefly covers the salient points of various text-based models with their underlying mechanisms and training strategies, specifically for deep learning methods, as they have changed the game for text classification by learning hierarchical and contextual features from text data. Finally, the last section does a comparative analysis of the different types of deep learning models, such as CNNs, RNNs and LSTMs and Transformers.

Convolutional Neural Networks

CNNs are commonly used for text classification because they can recognize local patterns (n-grams) via convolutional filters. They are very efficient with short texts such as sentiment analysis and spam detection. CNNs can also efficiently run in parallel which allows them to handle large datasets. However, their filters with a predetermined size constrain them from capturing long-term dependencies, an essential characteristic of textual information.

A. Recurrent Neural Networks

RNNs work on sequential data, which is why it keeps a hidden state that contains temporal information. These models excel at context understanding tasks such as document classification and sequence labeling. Although capable of processing sequence data, RNNs enjoy vanishing gradients, preventing the proper calibration for long texts or deep architectures, leading to unprepared long-term sequence learning tasks.

B. Long Short-Term Memory Networks

LSTMs are used to overcome the limitations of the RNN by adding gates to manage the flow of information. They perform well at memorizing long-range dependencies and are extensively used for applications such as topic classification and sentiment analysis. Yes, but LSTMs are neck-brace intensive and don't scale very well to large data or more complex tasks.

C. Transformer Models

The architecture of Transformers has led to transformative advancements in text classification given their self-attention mechanisms that can capture global dependencies and contextual nuances. Pre-trained models such as BERT and GPT are significantly ahead of the curve on various benchmarks.

Transformers are the most powerful model for documents of arbitrary length and for multi-lingual tasks. Nevertheless, their large computational and memory consumption hinder their deployment on resource-limited devices.

D. Hybrid Models

Hybrid approaches that combine CNNs and LSTMs or incorporate attention mechanisms utilize the strengths of multiple architectures to improve performance. It works especially well in scenarios where both local feature extraction and context are necessary. But hybrid models also come with added complexity, which can result in a higher computational cost and increased risk of overfitting.

E. Limitations Across Models

Some common limitations across models involve sensitivity to the data quality and amount, challenges in interpretability of results, and high computational resource requirements. Moreover, intermediate balanced datasets are still a significant problem to address and it usually has to be done by using data generation techniques or other re-sampling methods.

This comparative study demonstrates that the optimal architecture certainly depends on the application, its dataset size, domain characteristics, and computational requirements, indicating that no single model is optimal in every circumstance. Hence, future work must emphasize the need for lightweight, interpretable, and scalable models which can mitigate some of these weaknesses, thereby improving the practicality of using deep learning to address text classification problems, as well as a range of datasets pertaining to it.

IV. EMERGING TRENDS, CHALLENGES, AND OPPORTUNITIES IN DEEP LEARNING FOR TEXT CLASSIFICATION

The field of deep learning for text classification is rapidly evolving, driven by the need for models that can handle the increasing complexity and volume of textual data. Emerging trends reflect the innovative strategies researchers are employing to overcome existing limitations and improve the performance and scalability of text classification systems.

One significant trend is the integration of attention mechanisms and Transformer architectures, such as BERT, GPT, and RoBERTa. These models have redefined text classification by introducing context-aware embeddings and enabling bidirectional understanding of text. Transformers excel in capturing long-range dependencies, making them particularly effective for complex tasks like multi-label classification and hierarchical categorization.

Another promising development is the rise of hybrid models that combine the strengths of different architectures, such as CNNs and RNNs, or integrate domain-specific features with deep learning methods. This hybridization addresses challenges such as overfitting and computational inefficiency while enhancing the robustness of predictions.

Despite these advancements, several challenges persist. Data imbalance, a common issue in real-world datasets, can significantly affect model performance, particularly when underrepresented classes are crucial. Techniques like data augmentation, synthetic oversampling, and class-weight adjustments are being explored to mitigate this. Additionally, interpretability and explain ability remain critical challenges, as the black-box nature of deep learning models often limits their acceptance in high-stakes applications such as healthcare and legal systems.

Opportunities lie in multimodal learning, which integrates textual data with other data types such as images, videos, or graphs to enhance classification accuracy. This is particularly relevant for domains like social media analysis and healthcare, where textual data is often accompanied by other contextual information. Furthermore, advances in federated learning and edge computing present new possibilities for deploying deep learning models in decentralized and resource-constrained environments, enabling privacy-preserving text classification at scale.

Lastly, the exploration of low-resource settings, including transfer learning and few-shot learning techniques, is expanding the accessibility of deep learning for text classification to languages and domains with limited labelled data. This approach has the potential to democratize AI and bridge the gap between high-resource and underrepresented languages.

By addressing these challenges and leveraging emerging trends, researchers can unlock new opportunities, pushing the boundaries of what deep learning models can achieve in text classification.

V. PRACTICAL APPLICATIONS OF DEEP LEARNING IN TEXT CLASSIFICATION ACROSS DIVERSE DOMAINS

Deep learning-based approaches have greatly changed the field of text classification by offering highly effective and scalable solutions for various applications. These developments have been particularly prominent across a number of fields — natural language text extraction and classification have often been key to decision-making and operational effectiveness.

A. Sentiment Analysis

One popular example of text classification is sentiment analysis, which involves using deep learning models to read customer opinions, feedback, and reviews to classify the sentiment expressed as positive, negative, or neutral. Businesses use this data to optimize products, refine customer experience, and tailor marketing efforts. Traditional lexicon-based approaches map positive scores to positive words, leading to potential misreading of nuanced sentiment in text, like sarcasm or mixed feelings [16].

B. Spam Detection and Email Categorization

The deep learning text classifiers are extensively used in to identify a spam email and sort genuine emails in folders such as promotions, updates, primary inbox, etc. Neural networks are used in these models to detect more subtle patterns and language constructs used in a spam, greatly increasing accuracy compared to rule-based systems [17].

C. Healthcare

Text classification has a significant role in the healthcare sector, such as medical record analysis, symptoms classification, and disease prediction. For instance, patient records and clinical notes may be tagged to aid in diagnostics or to flag potential risk factors. Inspired by the effectiveness of deep learning in drug discovery, it has also been applied in research papers and drug reviews for treatment efficacy [18].

D. Legal and Financial Sectors

Another important use case is legal document classification, the automated categorization of contracts, case summaries, and legal filings for easy retrieval. In finance, end-to-end deep learning models are used to classify customer feedback, financial reports and transactional data for applications like fraud detection, and score credit [19].

E. Social Media Analysis

In recent years, the rise of social media platforms has put deep learning to use in the areas of content moderation, topic classification and hate speech detection.

These applications include monitoring posts, comments, and tweets for potential violations of community guidelines, as well as detecting emerging trends or sentiments [20].

F. Education

Text classification is used in the education sector for essay grading, organizing educational resources in the right headings, and for plagiarism checking. Analysing student feedback and engagement data to provide personalized learning experience [21].

G. Customer Support and Chatbots

For example, text classification is the foundation behind the tunnelling of AI chatbots and also virtual assistants, whereby these agents can understand the queries from the users and direct them toward possible answers or services. These systems utilize deep learning models to gauge intent and context for smooth interactions [22].

H. News and Document Classification

For instance, text classification is used in the media field to label news articles with a specific topic, i.e sports, politics, entertainment, etc.

In document management systems, it helps in categorizing considerable repositories of text for quick access and retrieval [23].

I. Cybersecurity

Text classification in Cyberspace is used for detecting phishing, identifying malicious URLs, and categorization of alerts. This makes these applications capable of identifying and neutralizing potential threats instantly [24].

These applications showcase the wide-ranging capabilities of deep learning in addressing complex text classification tasks, revolutionizing industries and enhancing their operational efficiency and effectiveness. With the evolution of deep learning models, their application scope can broaden even more with endless possibilities.

VI. CONCLUSION

This paper provided an in-depth overview of various deep learning methodologies for text classification and their transformative applications in multiple domains. Often when wording their instructions and with more complex text entering flows, they had to learn new patterns to accurately classify into their existing end categories—and therefore, the study explored a suite of models from CNNs, RNNs, and LSTMs to Transformer-based such as BERT and GTPs.

The gap of traditional machine learning methods on contextuality and scalability can be addressed significantly with DIW, as you can deep dive into data capturing complex patterns and long-range dependencies in text. It highlighted that hybrid methods and attention mechanisms are key in improving the classification accuracy and robustness.

Additionally, applications ranging from sentiment analysis to spam detection to healthcare text categorization to cybersecurity showcased the practicality and adaptability of these models in real-world situations. These successes are still marred with severe obstacles such as data immensity, huge computability resources required and lack of explain ability.

Recent studies proposed addressing some of the existing deficiencies, such as lightweight approaches, interpretability and low-resource environments by developing efficient models. Multimodal learning, federated learning, and few-shot learning innovations can further extend the potential usage of text classification systems. Moreover, future model development should focus on integrating ethical considerations like reducing bias and preserving privacy. By addressing these challenges and harnessing emerging trends, deep learning is well-positioned to further advance text classification, creating new opportunities for research and practical implementation across domains.

REFERENCES

- [1] J. Gu, K. Ren, and B. Gao, "Deep learning-based text knowledge classification for whole-process engineering consulting standards," *Journal of Engineering Research*, vol. 12, pp. 61–71, 2024.
- [2] R. Li, M. Liu, D. Xu, J. Gao, F. Wu, and L. Zhu, "A review of machine learning algorithms for text classification," in W. Lu et al. (Eds.): *CNCERT 2021, CCIS 1506*, pp. 226–234, 2022. doi: https://doi.org/10.1007/978-981-16-9229-1_14.
- [3] S. Minaee, N. Kalchbrenner, E. Cambria, N. Nikzad, M. Chenaghlu, and J. Gao, "Deep learning based text classification: A comprehensive review," *arXiv preprint arXiv:2004.03705v3 [cs.CL]*, Jan. 2021.
- [4] Z. Wan, "Text classification: A perspective of deep learning methods," *arXiv preprint arXiv:2309.13761v1 [cs.CL]*, Sep. 2023.
- [5] Dogra, Varun, Verma, Sahil, Kavita, , Chatterjee, Pushpita, Shafi, Jana, Choi, Jaeyoung, Ijaz, Muhammad Fazal, A Complete Process of Text Classification System Using State-of-the-Art NLP Models, *Computational Intelligence and Neuroscience*, 2022, 1883698, 26 pages, 2022. <https://doi.org/10.1155/2022/1883698>
- [6] Prabhakar, Sunil Kumar, Won, Dong-Ok, Medical Text Classification Using Hybrid Deep Learning Models with Multihead Attention, *Computational Intelligence and Neuroscience*, 2021, 9425655, 16 pages, 2021. <https://doi.org/10.1155/2021/9425655>
- [7] T. T. Dien, B. H. Loc and N. Thai-Nghe, "Article Classification using Natural Language Processing and Machine Learning," 2019 International Conference on Advanced Computing and Applications (ACOMP), Nha Trang, Vietnam, 2019, pp. 78-84, doi: 10.1109/ACOMP.2019.00019.
- [8] G. S. Greeshma, B. Ahuja, H. Samota, K. Vanitha, A. Dubey, S. Mujoo, and S. P. Saneesh, "A novel approach for biomedical text classification using deep learning and NLP for disease prediction," *International Journal of Intelligent Systems and Applications in Engineering*, vol. 12, no. 20s, pp. 657–666, 2024.
- [9] Prabhakar SK, Rajaguru H, So K and Won D-O (2022) A Framework for Text Classification Using Evolutionary Contiguous Convolutional Neural Network and Swarm Based Deep Neural Network. *Front. Comput. Neurosci.* 16:900885. doi: 10.3389/fncom.2022.900885
- [10] Alqahtani A, Ullah Khan H, Alsubai S, Sha M, Almadhor A, Iqbal T and Abbas S (2022) An ecient approach for textual data classification using deep learning. *Front. Comput. Neurosci.* 16:992296. doi: 10.3389/fncom.2022.992296
- [11] P. Lavanya and E. Sasikala, "Deep Learning Techniques on Text Classification Using Natural Language Processing (NLP) In Social Healthcare Network: A Comprehensive Survey," 2021 3rd International Conference on Signal Processing and Communication (ICPSC), Coimbatore, India, 2021, pp. 603-609, doi: 10.1109/ICSPSC51351.2021.9451752.
- [12] ext classification of unstructured medical notes with various levels of class imbalance. *BMC Med Res Methodol* 22, 181 (2022). <https://doi.org/10.1186/s12874-022-01665-y>
- [13] Wang, Y., Sohn, S., Liu, S. et al. A clinical text classification paradigm using weak supervision and deep representation. *BMC Med Inform Decis Mak* 19, 1 (2019). <https://doi.org/10.1186/s12911-018-0723-6>
- [14] Kolluri, Johnson and Razia, Shaik and Nayak, Soumya Ranjan, Text Classification Using Machine Learning and Deep Learning Models (June 4, 2020). International Conference on Artificial Intelligence in Manufacturing & Renewable Energy (ICAIMRE) 2019, Available at SSRN: <https://ssrn.com/abstract=3618895> or <http://dx.doi.org/10.2139/ssrn.3618895>
- [15] Hui Li, Zeming Li, Text Classification Based on Machine Learning and Natural Language Processing Algorithms, *Wireless Communications and Mobile Computing*, 2023, 9796719, 1 page, 2023. <https://doi.org/10.1155/2023/9796719>
- [16] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *NAACL-HLT*.
- [17] Almeida, T. A., Hidalgo, J. M. G., & Yamakami, A. (2011). Contributions to the Study of SMS Spam Filtering: New Collection and Results. *Proceedings of the ACM Symposium on Document Engineering*.
- [18] Lee, J., Yoon, W., Kim, S., et al. (2020). BioBERT: A Pre-trained Biomedical Language Representation Model for Biomedical Text Mining. *Bioinformatics*, 36(4), 1234-1240.
- [19] Chalkidis, I., Fergadiotis, M., Malakasiotis, P., et al. (2019). Legal-BERT: The Muppets Straight Outta Law School. *arXiv preprint arXiv:2010.02559*.
- [20] Mozafari, M., Farahbakhsh, R., & Crespi, N. (2020). Hate Speech Detection and Racial Bias Mitigation in Social Media Based on BERT Model. *PLOS ONE*, 15(8), e0237861.

- [21] Ullmann, T., & Wild, F. (2020). Automatic Detection of Plagiarism in Student Writings Using Deep Learning Models. *IEEE Transactions on Learning Technologies*, 13(3), 456-463.
- [22] Jain, M., Kumar, P., Bhansali, S., et al. (2018). Evaluating and Informing the Design of Chatbots. *Proceedings of the 2018 Conference on Human Factors in Computing Systems (CHI)*.
- [23] Lewis, D. D., Yang, Y., Rose, T. G., & Li, F. (2004). RCV1: A New Benchmark Collection for Text Categorization Research. *Journal of Machine Learning Research*, 5, 361-397.
- [24] Vinayakumar, R., Alazab, M., Soman, K. P., et al. (2019). Deep Learning Approach for Intelligent Intrusion Detection System. *IEEE Access*, 7, 41525-41550.