

A Study of Sentiment Analysis of Drug Reviews using Transformer-Based Embeddings and Hybrid Classifiers

Nayan Sarkar ¹, Dr. Balveer Singh ²

^{1,2} Department of Computer Science & Engineering, P.K. University, Shivpuri, M.P., India.

ABSTRACT

This study presents a sentiment analysis model for classifying user-generated drug reviews into positive, neutral, and negative categories using transformer-based embeddings and hybrid classifiers. A dataset of 5,170 drug reviews was collected from WebMD using the Beautiful Soup Python library and manually labeled for accuracy. After preprocessing and encoding, embeddings were generated using BERT, SciBERT, BioBERT, and SBERT models to capture contextual meaning. These embeddings were then used as input features for machine learning classifiers such as Decision Tree, Support Vector Machine (SVC), Random Forest, and Recurrent Neural Network (RNN). Experiments were conducted using Python (v3.9.11) and the SMOTE technique to address class imbalance. Performance evaluation using precision, recall, and F1-score revealed that the Random Forest and RNN models achieved the highest accuracy. The results demonstrate that transformer-based embeddings combined with hybrid classifiers offer an effective framework for analyzing public sentiment on drug efficacy and safety.

Keywords: *Sentiment Analysis, Drug Reviews, Recurrent Neural Network (RNN), Machine Learning.*

1. Introduction

There has been an explosion of user-generated information, especially evaluations and experiences from patients about different medications and medical treatments, due to the proliferation of digital health platforms. When it comes to medicine efficacy, safety, and patient happiness, these evaluations are a goldmine of practical information. Nevertheless, there are substantial obstacles to successful analysis of textual feedback due to its unstructured character. One answer is sentiment analysis, a branch of NLP that automates the process of extracting and interpreting feelings and opinions from text. Aiming to enhance healthcare decision-making and patient-centered medication development, this research use machine learning and deep learning methods to assess user-generated drug evaluations. The goal is to turn unstructured patient input into relevant pharmaceutical insights.

This project aims to construct an all-encompassing sentiment analysis framework that can reliably categorize evaluations of drugs as either favorable, negative, or neutral. The suggested system would decipher the subtleties of language, medical jargon, and emotional undertones in patient accounts by using sophisticated natural language processing models and machine learning techniques. Because it sheds light on medication effectiveness, patient adherence, and safety issues from an evidence-based perspective, understanding these attitudes is vital for pharmaceutical firms, healthcare providers, and lawmakers. Systematically analyzing patient input for improved clinical and commercial decision-making is vital due to the growing confidence patients have in online health communities.

Data gathering and preparation is the first stage of this procedure. Verified medication review sites and internet health discussion groups will be the sources of user-generated content. We will scrape the web using Python's BeautifulSoup package to efficiently retrieve data such review content, drug name, rating, and date. After extraction, the raw text is preprocessed using natural language processing tools such as NLTK and spaCy to lowercase, tokenize, remove stop words, and lemmatize. In order to prepare the text

for feature extraction and training machine learning models, this step is necessary to clean and standardize it. Improving model performance relies heavily on proper preprocessing, which involves removing noise and highlighting pertinent textual patterns that match to sentiment expressions.

Utilizing word embeddings and contextual language models, the analyzed data is subsequently numerically represented. In many cases, the contextual links between words cannot be captured by traditional approaches such as TF-IDF and Bag-of-Words (BoW). The study gets around these restrictions by using contextualized embeddings generated by pre-trained transformer-based models. As a starting point, one of these models is BERT, which stands for Bidirectional Encoder Representations from Transformers. By taking into account both the left and right word dependencies in a phrase, BERT is able to represent bidirectional context. Because of this, the model can comprehend nuanced language, like denial and sarcasm, which are common in user evaluations. "This drug worked well initially, but caused severe headaches later" is an example of a mixed-sentiment statement that BERT's contextual awareness aids in accurately classifying, as opposed to previous models.

The study also makes use of two specialized versions of the BERT model, SciBERT and BioBERT. SciBERT excels in comprehending domain-specific terms and technical jargon because it is educated on academic and scientific literature. By contrast, BioBERT is very effective at processing healthcare-related terms and medical terminology since it is trained on biomedical corpora like PubMed and clinical notes. By improving the semantic comprehension of drug review medical language, these models increase sentiment classification accuracy. Phrases such as "caused mild nausea" or "reduced systolic pressure" have substantial contextual meaning in biological language, which models like BioBERT can accurately decipher, even if they do not convey express emotional tone.

Also, for semantic analysis at the sentence level, Sentence-BERT (SBERT) is used. Using a Siamese network design, SBERT creates sentence embeddings of predetermined sizes, in contrast to regular BERT that gives embeddings at the token level. Subject similarity analysis, emotion grouping, and paraphrasing identification are all made easier with this capability. Research shows that SBERT is useful for categorizing reviews with comparable sentiments or for spotting small differences in sentiment intensity. A more sophisticated understanding of patient feedback is made possible by this sentence-level depiction, which aids pharmaceutical stakeholders in identifying trends, such as recurrent adverse effects or overall satisfaction patterns across several medications.

After the feature extraction step is over, sentiment classification is done using a variety of machine learning classifiers. To set the stage for baseline performance measurements, classic models include Decision Tree, Support Vector Classification (SVC), and Random Forest. To better comprehend which textual characteristics contribute most to sentiment prediction, Decision Tree classifiers create hierarchical divides based on features that best separate sentiment classes. This makes them clearly interpretable. Conversely, SVC finds the best hyperplane to divide sentiment categories in a space with many dimensions. Complex text data is well-suited to SVC because, with the help of kernel functions, it can effectively manage non-linear connections.

As an ensemble learning technique, Random Forest uses a combination of Decision Trees to make predictions more stable and less prone to overfitting. To make sure the categorization findings are solid, it uses majority voting to combine the output from many trees. Important for dealing with varied and unexpected user-generated evaluations, this ensemble method enables the model to generalize well to unseen data. To further understand the sequential relationships in text data, Recurrent Neural Networks (RNNs) are used. Reasoning neural networks (RNNs) are useful for simulating the temporal sentiment

changes that occur in drug evaluations, which often include numerous words with changing sentiments over time (e.g., early pleasure followed by unfavorable experiences). The network is able to learn context from word to word thanks to its recurrent connections, which enable it to remember what words have come before.

The research uses stringent assessment standards to guarantee the findings are reliable. To test the success of the classification, the models are evaluated using metrics like as accuracy, precision, recall, F1-score, and ROC-AUC. Also, to make sure the model is consistent and to reduce bias, k-fold cross-validation is run. One way to see how well the models do at classifying attitudes as positive, negative, or neutral is via a confusion matrix. To find the best mix of algorithms and embedding models for sentiment classification, it is necessary to compare them.

Various open-source libraries and frameworks are essential to the execution of this study. Because of its great support for NLP and machine learning applications, Python is the main language used for programming. For traditional methods, you may use scikit-learn or another library, and for deploying deep learning models, you can use TensorFlow or PyTorch. Hugging Face's Transformers library streamlines the process of fine-tuning models for sentiment analysis by providing access to pre-trained models such as BERT, SciBERT, BioBERT, and SBERT. Matplotlib and Seaborn are used for data visualization and performance analysis. These tools make it easier to understand results and display conclusions visually.

2. Review of Literature

Alkhnbashi, Omer et al., (2024) In recent years, internet medical forums have grown into a vital, low-cost resource for medical service management input, where patients may share their stories and get answers from others. This data provides insight into the efficacy of pharmaceutical treatments, pain management tactics, and alternative therapies, in addition to helping to gauge patient happiness and enhance healthcare quality. This study thoroughly documents key aspects of patients' experiences, showcasing the positive and negative emotions expressed in their narratives. Our dataset includes around 15,000 items extracted from various sections of the well-known patient.info medical forum. Using a unique combination of content analysis, deep learning techniques, aspect-based sentiment analysis, and LLM, we conduct an examination of this data. The goal of our method is to unearth the many facets conveyed by patients' comments. The experiment verified that deep learning models could properly predict the sentiment values linked with seven distinct aspect categories found in the feedback. The LLM with few-shot learning stood out from the rest of the models. Our study provides valuable insights into patients' experiences in online forums and highlights the usefulness of advanced analytical methods for extracting meaning from patients' unstructured input. This knowledge might be very beneficial for healthcare professionals and medical service management.

Tanasescu, Laura et al., (2024) It is critical for most companies to increase staff performance and productivity because of the various changes that have occurred in people and processes in many organizations in the last few years. The fact that this can keep the business growing and competitive makes it even more important. Creating a machine learning system to predict the performance ratings of people in a particular company entails data pre-processing, variable selection, algorithm building, and hyperparameter shaping. Various methods are compared and contrasted in this research. In order to increase overall productivity and decrease reliance on subjective human judgment in performance assessments, we set out to determine the most effective strategies for predicting the provided factors.

Raju, Nidadavolu et al., (2024) Healthcare, banking, and e-commerce decision-makers are seeking relevant insights from big data, and predictive analytics is a vital tool for them due to the ever-increasing complexity and volume of data. When working with large amounts of data, it may be challenging to choose the optimal machine learning algorithm due to differences in processing requirements, scalability, and algorithmic performance. This paper provides a comprehensive review of popular machine learning algorithms for predictive analytics, focusing on their performance and practicality in big data environments. Research sorts algorithms into three categories based on their learning type: supervised, unsupervised, and reinforcement learning. Algorithms are then evaluated based on many factors, such as the accuracy of their predictions, the efficiency of their computations, their scalability, and their suitability for real-time analytics. In order to determine how well they manage massive datasets, a number of techniques are tested and evaluated, including decision trees, linear regression, neural networks, support vector machines, and clustering methods. The study offers a set of metrics, including training time, F1-score, and accuracy, to further assess the algorithms' computational feasibility and predictive potential. Through the application of these algorithms to a sample of a huge data collection, this case study sheds light on their real-world performance in several predictive analytics scenarios. Comparison tables and performance graphs are good visual representations of the data that assist to comprehend the pros and cons of various algorithm possibilities. Algorithms like random forests and neural networks perform better at generating predictions, but they may be too resource-intensive for real-time processing in big data applications, according to the study. In this work's later half, we provide suggestions for selecting machine learning algorithms based on specific predictive analytics objectives, data attributes, and processing requirements. The article continues by discussing the challenges of implementing these algorithms in big data environments, but it also offers solutions, such as incorporating deep learning and employing distributed computing, that might enhance predictive analytics going ahead.

Udeh, Chioma et al., (2024) Big Data Analytics is revolutionizing the way businesses get insights from diverse and expansive datasets in the dynamic realm of business intelligence. After a lengthy analysis of Big Data Analytics and its impact on modern BI, this paper provides a concise summary of the key issues. Big Data Analytics represents a sea shift in how organizations make decisions by allowing them to fully use their data assets. The study delves into the many uses of Big Data Analytics, showcasing its significance in customer-centric endeavors, operational effectiveness, risk mitigation, and long-term strategy formulation. Companies may greatly enhance their strategic planning with the help of predictive analytics when they utilize it to anticipate market shifts. Incorporating analytics-derived insights into decision-making drives future-oriented strategic efforts that are better informed. When it comes to identifying and preventing fraud, risk management becomes much more proactive with the integration of Big Data Analytics. The ability to manage enormous volumes of data in real-time raises awareness while decreasing financial risks associated with fraudulent activities. Through the use of scenario modeling, organizations may improve their risk mitigation and anticipation capabilities. The focus shifts to operational efficiency as analytics expose inefficiencies in supply chains, retail operations, and industrial processes. Through the use of data analytics, which enable real-time decision-making, retailers can maintain a responsiveness to changing customer preferences and market circumstances. Increased focus on customers is a direct result of personalized marketing and predictive analytics used by customer care departments. In an effort to boost customer satisfaction and retention, this analysis looks at how Big Data Analytics enables companies to provide personalized experiences. The paper provides a concise overview of Big Data Analytics' game-changing trajectory in modern BI, focusing on its role in simplifying strategic complexity, decreasing risk, enhancing operations, and making customer-centric choices. In today's data-driven world, businesses may benefit from a comprehensive review that tries to show them how to use Big Data Analytics to their advantage.

Hanif, Atikah et al., (2023) Technological advancements have caused an explosion of data in several industries, including aviation. With the use of big data analytics, it is possible to extract valuable insights from these enormous datasets. Relevant papers were selected from twenty academic articles that were indexed in internet databases such as ScienceDirect, Web of Science, Google Scholar, and IEEE Xplore. The words "data mining," "machine learning," "big data analytics," "predictive analytics," and the "airlines industry" were all part of the search. Based on the findings, Big Data Analytics (BDA) offers several applications in the aviation industry, including optimizing airline operations, increasing customer pleasure and service quality, strengthening safety and risk management, and aircraft maintenance. The investigation reveals that there are many significant challenges to the adoption of BDA in the aviation industry, including concerns over data integration, privacy and security, and compliance with regulations. When it comes down to it, this review article is a gold mine of information on the pros and cons of BDA in the aviation industry. In doing so, it increases our understanding of the subject and prepares the way for future research and practical applications.

3. Proposed Approach

Libraries and Models Used

The libraries that were used and the many pre-trained models are listed below: -

- BeautifulSoup
- Bidirectional Encoder Representations from Transformers
- SciBERT
- BioBERT
- SBERT
- Decision Tree
- Support Vector Classification
- Random Forest
- Recurrent Neural Network

Dataset

The dataset was constructed by manually classifying 5,170 medication reviews into three separate categories according to their sentiment using the BeautifulSoup Python package. The reviews were sourced from the publically available website WebMD. The dataset included two columns, "Reviews" and "Classification," since the categories were "Positive," "Neutral," and "Negative."

Preprocessing

The first and most important stage in data mining and machine learning is data preprocessing, which entails cleaning and organizing raw data so it can be better analyzed and modelled. Data accuracy, consistency, and completeness are ensured by a sequence of procedures that deal with discrepancies, missing numbers, outliers, and mistakes. In order to improve the data's quality and make it more usable for different algorithms, data preparation also include data reduction, data transformation, and feature engineering.

Generation of Embeddings

We need to convert the pre-processed data, which includes reviews and labels, into numerical numbers. This led to the creation of embeddings of the reviews and the ability to encode labels using LabelEncoder. A low-dimensional representation of a high-dimensional data set, such text, pictures, or sounds, is called an embedding. Automated systems are able to comprehend and handle complicated data because they record the

significance and connections between the data elements. Embeddings are used in several applications, such as recommender systems, image recognition, anomaly detection, and natural language processing, and they are trained using massive volumes of data. A number of natural language processing models, such as S-BERT, BioBERT, SciBERT, and BERT, are fed the reviews. By taking the review's context into account, the embeddings acquired from these models capture the text's meaning as vectors representing the reviews.

Proposed Models of Machine Learning

The categorization procedure starts when the embeddings are generated and their labels are identified. This is done by using different machine learning models to process the embeddings along with the labels that go with them. Decision tree, logistic regression, support vector, and random forest classifiers are some of the most popular machine learning methods used for classification problems. There are advantages and disadvantages to each of these algorithms; picking the right one depends on the details of the data and the result you're after.

Model Comparison

We compare and assess the models' performance. Since RNN produces better results in terms of accuracy, it is selected for use in model construction based on evaluation criteria.

Experimental Procedure

The machine running Windows 11 and powered by an Intel Core i5 CPU was used to complete this job. Visual Studio Code in conjunction with the Python 3.9.11 environment was used to run the code. Since gathering data is the first order of business, web scraping should be executed before a plethora of other responsibilities. To get over the issue of dataset imbalance, the Synthetic Minority Oversampling Technique (SMOTE) is used. A dataset with ratings and labels is now available. We used LabelEncoder to encode the labels. The reviews were embedded using several pre-trained models such as Bert, SciBERT, BioBERT, and SBERT. The Python Transformers module was used for this purpose. Training and testing sets are subsequently created from the dataset. A model's generalizability and prediction power on novel and unseen data may be better evaluated with correct data splitting. Two variables were used: one for storing independent characteristics and another for saving goal values. The dataset was analyzed using a number of classification models, such as Logistic Regression, Support Vector Machine, Decision Tree, and Random Forest. We need to choose one of them for our model based on this. Among the classifiers, Random Forest has the highest accuracy. To improve accuracy and decrease variation, an ensemble learning approach called a random forest classifier uses many decision trees to produce predictions.

4. Results and Discussion

In this part, we provide the results of the experiment that used the suggested method for sentiment analysis on UGC on drugs. In this part, we provide the comparison findings and then go into depth about what we found.

Table 1: Performance Metrics for Various Classifiers using BERT with Recurrent Neural Network

Sentiment	Precision	Recall	F1-Score	Support
Negative	0.63	0.70	0.66	510
Neutral	0.40	0.78	0.16	205
Positive	0.47	0.60	0.53	325
Accuracy			0.55	1040
Macro Avg	0.50	0.46	0.45	1040
Weighted Avg	0.54	0.55	0.52	1040

Table 1 presents the performance metrics of a Recurrent Neural Network (RNN) model using BERT embeddings for sentiment analysis of drug reviews. The model demonstrates the highest recall for neutral sentiments (0.78), indicating that it is effective at identifying most neutral reviews, although the corresponding F1-score is low (0.16), suggesting imbalanced prediction performance for this class. Negative sentiment shows a balanced performance, with a precision of 0.63, recall of 0.70, and an F1-score of 0.66, reflecting the model's relatively strong ability to correctly classify negative reviews. Positive sentiment is moderately predicted, with precision at 0.47, recall at 0.60, and F1-score at 0.53, showing that some positive reviews are misclassified as other sentiments.

Table 2: Performance Metrics for Various Classifiers using SBERT with Recurrent Neural Network

Sentiment	Precision	Recall	F1-Score	Support
Negative	0.64	0.74	0.68	510
Neutral	0.29	0.23	0.25	205
Positive	0.52	0.46	0.48	325
Accuracy			0.55	1040
Macro Avg	0.48	0.47	0.47	1040
Weighted Avg	0.53	0.55	0.53	1040

Table 2 displays the performance metrics of a Recurrent Neural Network (RNN) model using SBERT embeddings for sentiment analysis of drug reviews. The model shows its strongest performance in detecting negative sentiments, with a precision of 0.64, recall of 0.74, and F1-score of 0.68, indicating that negative reviews are identified accurately and consistently. Positive sentiment classification is moderate, with precision at 0.52, recall at 0.46, and an F1-score of 0.48, suggesting that some positive reviews are misclassified as neutral or negative. The neutral class demonstrates weaker performance, with a precision of 0.29, recall of 0.23, and F1-score of 0.25, indicating difficulty in correctly identifying neutral reviews.

Table 3: Performance Metrics for Various Classifiers using SciBERT with Recurrent Neural Network

Sentiment	Precision	Recall	F1-Score	Support
Negative	0.58	0.77	0.66	510
Neutral	0.35	0.10	0.16	205
Positive	0.49	0.48	0.48	325
Accuracy			0.56	1040
Macro Avg	0.47	0.45	0.43	1040
Weighted Avg	0.51	0.55	0.50	1040

Table 3 presents the performance of a Recurrent Neural Network (RNN) model using SciBERT embeddings for sentiment analysis of drug reviews. The model exhibits strong performance in detecting negative sentiments, achieving a precision of 0.58, recall of 0.77, and F1-score of 0.66, indicating effective identification of negative reviews. Positive sentiment classification is moderate, with precision at 0.49, recall at 0.48, and an F1-score of 0.48, showing balanced but not optimal performance. Neutral sentiment detection is relatively weak, with a precision of 0.35, recall of 0.10, and F1-score of 0.16, highlighting the model's difficulty in correctly classifying neutral reviews.

Table 4: Performance Metrics for Various Classifiers using BioBERT with Recurrent Neural Network

Sentiment	Precision	Recall	F1-Score	Support
Negative	0.62	0.73	0.67	510
Neutral	0.32	0.12	0.18	205
Positive	0.50	0.55	0.52	325
Accuracy			0.56	1040
Macro Avg	0.48	0.47	0.46	1040
Weighted Avg	0.52	0.56	0.53	1040

Table 4 shows the performance of a Recurrent Neural Network (RNN) model using BioBERT embeddings for sentiment analysis of drug reviews. The model demonstrates strong performance in detecting negative sentiments, with a precision of 0.62, recall of 0.73, and F1-score of 0.67, indicating reliable classification of negative reviews. Positive sentiment is moderately well predicted, with precision at 0.50, recall at 0.55, and F1-score at 0.52, suggesting that the model can capture positive sentiments reasonably accurately. Neutral sentiment detection remains challenging, with a precision of 0.32, recall of 0.12, and F1-score of 0.18, highlighting difficulty in correctly identifying neutral reviews.

5. Conclusion

Artificial intelligence has the ability to revolutionize the pharmaceutical and healthcare industries, as shown by research that used machine learning methods to analyze user-generated medication evaluations sentiment. The study shows that patient views may be systematically analyzed to find out about treatment effectiveness, side effects, and overall satisfaction using sophisticated natural language processing (NLP) models including BERT, SciBERT, BioBERT, and SBERT. By combining several machine learning methods such as Support Vector Classification, Random Forest, and Recurrent Neural Networks, sentiment classification becomes more accurate and dependable. This guarantees that positive, negative, and neutral feedback are properly distinguished.

By taking this approach, lawmakers, healthcare providers, and pharmaceutical firms may all make evidence-based choices that put patients' needs first. Product enhancement, patient education, and post-market monitoring may all benefit from the results of these types of research. Furthermore, stakeholders may enhance the healthcare ecosystem's responsiveness, openness, and trust by comprehending public mood patterns. Finally, the study proves that sentiment analysis is a must-have for future patient-centered healthcare by helping to close the information gap between patients and the pharmaceutical sector.

References

1. Hanif and A. Jaafar, "A Review: Application of Big Data Analytics in Airlines Industry," *Open International Journal of Informatics*, vol. 11, no. 2, pp. 196-209, 2023, doi: 10.11113/oiji2023.11n2.254.
2. C. Udeh, O. Orieno, O. Daraojimba, N. Ndubuisi, and O. Oriekhoe, "Big Data Analytics: A Review of Its Transformative Role in Modern Business Intelligence," *Computer Science & IT Research Journal*, vol. 5, no. 1, pp. 219-236, 2024, doi: 10.51594/csitrj. v5i1.718.
3. N. Raju and P. Devi, "A Comparative Analysis of Machine Learning Algorithms for Big Data Applications in Predictive Analytics," *International Journal of Scientific Research and Management (IJSRM)*, vol. 12, no. 10, pp. 1608-1630, 2024, doi: 10.18535/ijsrn/v12i10.ec09.
4. L. Tanasescu, A. Vines, A.-R. Bologa, and O. Virgolici, "Data Analytics for Optimizing and Predicting Employee Performance," *Applied Sciences*, vol. 14, no. 1, pp. 3254, 2024, doi: 10.3390/app14083254.

5. O. Alkhnbashi, R. Nassr, and M. Hammoudeh, "Aspect-Based Sentiment Analysis of Patient Feedback Using Large Language Models," *Big Data and Cognitive Computing*, vol. 8, no. 12, pp. 2-29, 2024, doi: 10.3390/bdcc8120167.
6. T. V. Iyelolu, E. E. Agu, C. Idemudia, and T. I. Ijomah, "Leveraging artificial intelligence for personalized marketing campaigns to improve conversion rates," *International Journal of Engineering Research and Development*, vol. 20, no. 8, pp. 253–270, 2024.
7. M. R. Ayyagari, "A framework for analytical CRM assessments challenges and recommendations," *International Journal of Information, Business and Management*, vol. 13, no. 2, pp. 108–121, 2021.
8. K. W. Walker, "Exploring adaptive boosting (AdaBoost) as a platform for the predictive modeling of tangible collection usage," *The Journal of Academic Librarianship*, vol. 47, no. 6, p. 102450, 2021.
9. Sumit, "Drug review sentiment analysis using boosting algorithms," *International Journal of Trend in Scientific Research and Development (IJTSRD)*, vol. 5, no. 4, pp. 937–941, 2021.
10. V. Djurisic, L. Kascelan, and S. Rogic, "Bank CRM optimization using predictive classification based on the support vector machine method," *Applied Artificial Intelligence*, vol. 34, no. 1, pp. 1–22, 2020.
11. B. Zorić, "Benefits of educational data mining," *Journal of International Business Research and Marketing*, vol. 6, no. 1, pp. 12–16, 2020.
12. Y. Rohani, Z. Torabi, and S. Kianian, "A novel hybrid genetic algorithm to predict students' academic performance," *Journal of Electrical and Computer Engineering Innovations*, vol. 8, no. 2, pp. 219–232, 2020.
13. F. Ofori, E. Maina, and R. Gitonga, "Using machine learning algorithms to predict students' performance and improve learning outcomes: A literature-based review," *Journal of Information and Technology*, vol. 4, no. 1, pp. 33–55, 2020.
14. D. Buenaño-Fernández, D. Gil, and S. Luján-Mora, "Application of machine learning in predicting performance for computer engineering students: A case study," *Sustainability*, vol. 11, no. 10, pp. 2833, 2019.
15. R. P. Bunker and F. Thabtah, "A machine learning framework for sport result prediction," *Applied Computing and Informatics*, vol. 15, no. 1, pp. 27–33, 2019.
16. N. Chen, C. Xiong, W. Du, C. Wang, X. Lin, and Z. Chen, "An improved genetic algorithm coupling a back-propagation neural network model (IGA-BPNN) for water-level predictions," *Water*, vol. 11, no. 9, p. 1795, 2019.
17. Sahoo et al., "Deep learning-based health recommender system using collaborative filtering," *Computation*, vol. 7, no. 2, pp. 1283–1299, 2019.
18. V. Deshmukh et al., "Comparative study of different approaches of opinion mining," *International Journal of Research*, vol. 8, no. 1, pp. 644–651, 2019.
19. M. Abdullah and A. Fahad, "Enhancing performance of educational data using big data and Hadoop," *International Journal of Applied Engineering Research*, vol. 14, no. 19, pp. 3814–3819, 2019.
20. Al-Homery, H. Asharai, and A. Ahmad, "The core components and types of CRM," *Pakistan Journal of Humanities and Social Sciences*, vol. 7, no. 1, pp. 121–145, 2019.
21. E. T. Lau, L. Sun, and Q. Yang, "Modeling, prediction, and classification of student academic performance using artificial neural networks," *SN Applied Sciences*, vol. 1, no. 9, pp. 1–10, 2019.