Sentiment Analysis of Machine Learning Algorithms: A Transformer-Based Approach

Saad Iqbalbhai Mandli

Leelaben Dashrathbhai Ramdas Patel Institute of Technology and Research, Gandhinagar

[1] *Date of Receiving: 12/02/2025;*
Date of Acceptance: 28/03/2025;
Date of Publication: 05/05/2025

Abstract

Sentiment analysis is one of the most widely used applications of Natural Language Processing because it helps convert unstructured opinions into measurable insights. Traditional machine learning algorithms such as Naive Bayes, Logistic Regression, Support Vector Machine, Random Forest and Gradient Boosting have been used extensively for classifying reviews, tweets, feedback and news articles into positive, negative or neutral categories. However, these methods depend heavily on manual feature engineering, bag-of-words, TF-IDF vectors and domain-specific pre-processing. Deep learning methods improved performance by learning distributed representations, but recurrent and convolutional architectures still struggle with long-range context, sarcasm, negation and aspect-level sentiment. Transformer-based models, especially BERT and its variants, address many of these limitations through self-attention, contextual embeddings and transfer learning. This paper studies sentiment analysis from the perspective of machine learning algorithms and proposes a transformer-based approach for improving classification accuracy, contextual understanding and adaptability across domains. A comparative analysis is presented between traditional machine learning, deep learning and transformer-based models. The paper concludes that transformer-based models provide superior contextual representation and generalization, although they require higher computational resources and careful fine-tuning.

1. Introduction

Sentiment analysis, also known as opinion mining, aims to identify the emotional polarity expressed in text. The polarity may be binary, such as positive or negative, or multi-class, such as positive, negative and neutral. In business, sentiment analysis is used to examine customer reviews, social media reactions, brand perception and service feedback. In politics, it helps understand public opinion, while in education and healthcare it can support feedback evaluation and patient experience analysis.

Early sentiment analysis systems relied on rule-based lexicons, where sentiment was identified by matching words with positive or negative dictionaries. Later, supervised machine learning methods became popular because they could learn from labeled text. Algorithms such as Naive Bayes, Support Vector Machine and Logistic Regression performed well when text was represented using TF-IDF or bag-of-words. These models remain useful because they are simple, fast and interpretable. However, they treat words mainly as independent features and often fail to understand context. For example, the sentence “The movie is not bad” may be wrongly classified if the model focuses only on the word “bad”.

Deep learning introduced word embeddings and neural networks that could represent semantic relationships more effectively. CNN, RNN, LSTM and GRU models improved sentiment classification by learning patterns from word sequences. Yet, these models still process text sequentially or through local filters, making it difficult to capture long-

[1] ***How to cite the article**:* Mandli S.I. (2025); Sentiment Analysis of Machine Learning Algorithms: A Transformer-Based Approach; *International Journal of Innovations in Applied Sciences and Engineering*; Vol 11, 76-81

range dependencies efficiently. The Transformer architecture changed this direction by using self-attention, allowing the model to assign importance to different words in a sentence simultaneously. BERT further improved performance by learning bidirectional contextual representations, making it highly suitable for sentiment analysis.

2. Literature Review

The development of sentiment analysis has moved from lexicon-based methods to machine learning, deep learning and transformer-based systems. Classical algorithms such as Naive Bayes and SVM have been widely studied because they are efficient for high-dimensional sparse text features. SVM is especially strong in text classification because it can separate classes using high-dimensional decision boundaries. Han et al. showed that SVM-based approaches can be improved using probabilistic latent semantic analysis and kernel methods for Twitter sentiment analysis.

Deep learning surveys show that neural networks can reduce manual feature engineering and learn semantic representations directly from data. Zhang, Wang and Liu reviewed deep learning for sentiment analysis and highlighted the role of CNNs, RNNs and attention mechanisms. Wankhade, Rao and Kulkarni later reviewed sentiment analysis methods, applications and challenges, emphasizing issues such as domain dependency, sarcasm and multilingual processing.

The Transformer architecture, introduced by Vaswani et al., replaced recurrence with self-attention. BERT, introduced by Devlin et al., became a major shift because it could be fine-tuned with a simple task-specific layer for multiple NLP tasks. Sun, Huang and Qiu applied BERT to aspect-based sentiment analysis by converting aspect sentiment into a sentence-pair task. Xu et al. further analysed BERT representations for aspect-based sentiment analysis and found that BERT captures aspect and opinion information through learned contextual features. Recent studies have also investigated Explainability in transformer sentiment models because black-box predictions are difficult to trust in business and academic applications.

3. Research Gap and Objectives

Although many studies have compared machine learning algorithms for sentiment analysis, there remains a gap in connecting traditional ML evaluation with transformer-based contextual modelling. Many classical studies compare Naive Bayes, SVM, Random Forest and Logistic Regression using accuracy, precision, recall and F1-score. However, they often use sparse text features and do not deeply examine semantic complexity. On the other hand, transformer-based research often reports high performance but gives less practical comparison with lightweight ML models that may be more suitable for small businesses or low-resource environments.

The central research gap is the need for a balanced framework that evaluates sentiment analysis algorithms not only by accuracy but also by context handling, scalability, interpretability, computational cost and domain adaptability. A transformer-based approach can be strong, but it should not automatically replace classical ML in all cases. Small datasets, limited hardware or simple binary classification may still favor SVM or Logistic Regression. Therefore, this paper aims to propose a comparative framework and a transformer-based sentiment analysis pipeline.

The objectives are: to study classical ML algorithms for sentiment classification; to examine the contribution of transformer-based models; to compare algorithms on practical dimensions; and to propose a fine-tuning approach using BERT-style architecture for sentiment classification.

	text	label	label_name
0	i didnt feel humiliated	0	sadness
1	i can go from feeling so hopeless to so damned...	0	sadness
2	im grabbing a minute to post i feel greedy wrong	3	anger
3	i am ever feeling nostalgic about the fireplac...	2	love
4	i am feeling grouchy	3	anger

Figure 1: Samples of Dataset

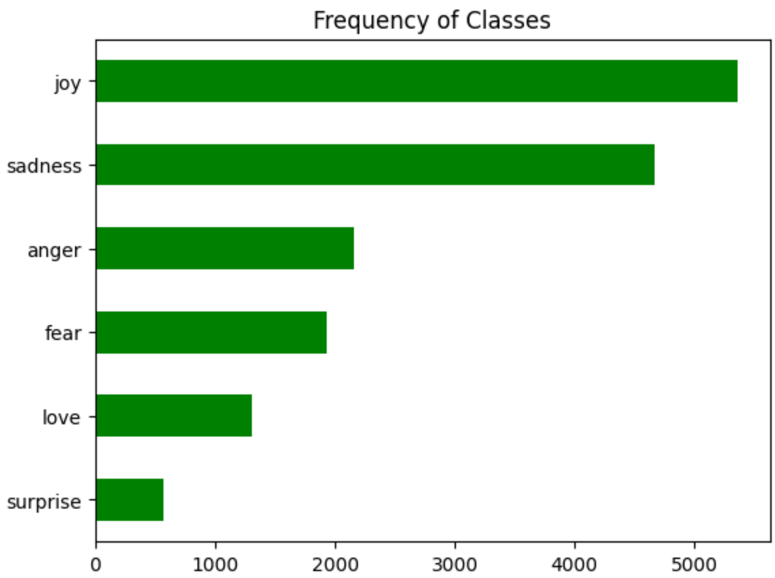

Figure 2: Frequency of Classes

4. Proposed Methodology

The proposed methodology follows a supervised sentiment classification pipeline. The input dataset may contain customer reviews, tweets, product feedback or educational comments. Each text record is labeled as positive, negative or neutral. The pipeline begins with data cleaning, including removal of duplicate records, irrelevant symbols and noisy formatting. Unlike traditional ML models, transformer models do not require aggressive pre-processing because punctuation, casing and sentence structure may carry meaning. However, normalization of URLs, emojis and repeated characters is useful for social media data.

For baseline models, text is converted into TF-IDF vectors. These vectors are used to train Naive Bayes, Logistic Regression, SVM and Random Forest classifiers. For the transformer-based model, text is tokenized using a WordPiece or subword tokenizer. Special tokens are added, and the sequence is passed into a pre-trained transformer encoder. The final hidden representation of the classification token is connected to a dense layer with softmax activation. The model is fine-tuned using labeled sentiment data.

Evaluation is performed using accuracy, precision, recall, F1-score and confusion matrix. F1-score is especially important when class distribution is imbalanced. For deployment, the trained model can be exposed through an API that receives text input and returns sentiment probability scores.

5. Transformer-Based Architecture

The transformer-based approach relies on self-attention. Self-attention allows each token to attend to other tokens in the sentence. This is useful in sentiment analysis because polarity often depends on relationships between distant words. In "The camera quality is good, but the battery is terrible," the model must understand that different aspects have different sentiments. Traditional models may classify the whole sentence as mixed or rely on dominant keywords, while transformers can represent the relationship between aspect terms and opinion words more effectively.

BERT is an encoder-only transformer model trained using masked language modeling and next sentence prediction. During fine-tuning, the pre-trained model is adapted to the sentiment task with labelled examples. This transfer learning process allows the model to use linguistic knowledge learned from large corpora while adjusting to domain-specific sentiment data. For aspect-based sentiment analysis, the input can be structured as a sentence pair: the original review and the target aspect. This helps the model classify sentiment toward a specific feature rather than the entire sentence.

A transformer-based sentiment model typically contains token embedding, positional embedding, multi-head self-attention, feed-forward layers, normalization, dropout and classification head. Hyperparameters include learning rate, batch size, maximum sequence length, dropout rate and number of epochs.

Architecture Component	Role in Sentiment Analysis
Tokenizer	Converts text into subword tokens
Positional Encoding	Preserves word order information
Self-Attention	Captures contextual relationships
Transformer Encoder	Produces deep contextual embeddings
Classification Head	Predicts sentiment class
Softmax Layer	Generates class probabilities

6. Comparative Analysis

Traditional machine learning algorithms perform well when text is clean, domain-specific and represented with strong features. Naive Bayes is fast and works well with small datasets, but it assumes feature independence. Logistic Regression is simple, interpretable and strong for linear decision boundaries. SVM is effective for sparse high-dimensional data and often performs better than Naive Bayes in text classification. Random Forest handles nonlinear relationships but may not always perform well with sparse TF-IDF features.

Deep learning models such as CNN and LSTM reduce manual feature engineering. CNN captures local n-gram patterns, while LSTM captures sequential dependencies. However, LSTM models are slower to train and may struggle with very long sequences. Transformer-based models outperform these approaches in contextual understanding because they process all tokens through attention. BERT and similar models are especially effective for complex sentiment expressions, sarcasm, negation and aspect-level classification.

The practical comparison shows that no single model is best in every situation. For small projects with limited hardware, Logistic Regression or SVM may be more appropriate. For high-stakes applications, large datasets and domain-specific analysis, transformers are more suitable. For real-time systems, distilled or smaller transformer models may provide a balance between accuracy and speed.

Model	Strength	Weakness	Best Use Case
Naive Bayes	Very fast, simple	Weak context handling	Small baseline systems
Logistic Regression	Interpretable, efficient	Linear assumption	Business review classification
SVM	Strong for sparse text	Slower with large data	High-dimensional TF-IDF tasks
Random Forest	Handles nonlinear patterns	Less effective on sparse text	Structured feature combinations
CNN	Captures local patterns	Limited long context	Short review classification
LSTM	Handles sequence order	Slow training	Moderate-length text
BERT/Transformer	Strong contextual understanding	Computationally expensive	Complex sentiment and ABSA

7. Discussion

The comparative analysis indicates that transformer-based models are not merely another classifier but a different representation strategy. Classical ML depends on feature extraction before learning. Transformers learn representations and classification jointly during fine-tuning. This is a major advantage for sentiment analysis because the emotional meaning of a word depends on context. For example, “unpredictable” may be negative for a washing machine but positive for a thriller movie. A TF-IDF model may struggle with this distinction, while a transformer can use surrounding words and domain context.

Another benefit is transfer learning. A pre-trained transformer can be fine-tuned with a smaller labeled dataset compared with training a deep neural network from scratch. This is valuable when labeled sentiment data is limited. However, transformers also introduce challenges. They require GPU resources for efficient training, careful hyperparameter tuning and attention to bias in pre-training data. They may also produce confident but wrong predictions, which makes Explainability important.

Explainable sentiment analysis is becoming increasingly important. Techniques such as attention visualization, LIME, SHAP and integrated gradients can help identify which words influenced the model output. This is useful for business users who need to understand why a review was classified as negative or why a customer complaint was prioritized.

8. Conclusion

Sentiment analysis has evolved from lexicon-based techniques to machine learning, deep learning and transformer-based architectures. Traditional machine learning algorithms remain useful because they are simple, fast and interpretable. SVM and Logistic Regression are strong baselines for many text classification tasks. Deep learning methods improve representation learning but still face limitations in long-range contextual understanding. Transformer-based models, especially BERT-style architectures, provide a more powerful approach by using self-attention and pre-trained contextual embeddings.

The proposed transformer-based approach improves sentiment classification by capturing semantic relationships, negation, aspect-level meaning and domain-specific context. Comparative analysis shows that transformers are best suited for complex sentiment tasks, while classical ML remains suitable for lightweight applications. Future work should focus on efficient transformer models, multilingual sentiment analysis, explainable AI and domain-specific fine-tuning. A practical sentiment analysis system should therefore not choose an algorithm only by accuracy, but also by cost, interpretability, scalability and deployment needs.

References

1. Ali, H., Hashmi, E., Yayilgan Yildirim, S., & Shaikh, S. (2024). Analyzing Amazon products sentiment: A comparative study of machine learning, deep learning, and transformer-based techniques. *Electronics, 13*(7), 1305. https://doi.org/10.3390/electronics13071305
2. Barbieri, F., Camacho-Collados, J., Espinosa Anke, L., & Neves, L. (2020). TweetEval: Unified benchmark and comparative evaluation for tweet classification. In *Findings of the Association for Computational Linguistics: EMNLP 2020* (pp. 1644–1650). https://doi.org/10.18653/v1/2020.findings-emnlp.148
3. Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT 2019)* (pp. 4171–4186). https://doi.org/10.18653/v1/N19-1423
4. Han, K.-X., Chien, W., Chiu, C.-C., & Cheng, Y.-T. (2020). Application of support vector machine in the sentiment analysis of Twitter dataset. *Applied Sciences, 10*(3), 1125. https://doi.org/10.3390/app10031125
5. Perikos, I., & Diamantopoulos, A. (2024). Explainable aspect-based sentiment analysis using transformer models. *Big Data and Cognitive Computing, 8*(11), 141. https://doi.org/10.3390/bdcc8110141
6. Pilicita Garrido, A., & Barra, E. (2025). Sentiment analysis with transformers applied to education: Systematic review. *International Journal of Interactive Multimedia and Artificial Intelligence, 9*(2), 72–83. https://doi.org/10.9781/ijimai.2025.02.008
7. Sun, C., Huang, L., & Qiu, X. (2019). Utilizing BERT for aspect-based sentiment analysis via constructing auxiliary sentence. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT 2019)* (pp. 380–385). https://doi.org/10.18653/v1/N19-1035
8. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems* (Vol. 30). Curran Associates, Inc.
9. Wang, J., Wei, J., Tian, F., et al. (2024). A comparative study of machine learning models for sentiment analysis of transboundary rivers news media articles. *Soft Computing, 28*, 13331–13347. https://doi.org/10.1007/s00500-024-10357-2
10. Wankhade, M., Rao, A. C. S., & Kulkarni, C. (2022). A survey on sentiment analysis methods, applications, and challenges. *Artificial Intelligence Review, 55*, 5731–5780. https://doi.org/10.1007/s10462-022-10144-1
11. Xu, H., Shu, L., Yu, P., & Liu, B. (2020). Understanding pre-trained BERT for aspect-based sentiment analysis. In *Proceedings of the 28th International Conference on Computational Linguistics (COLING 2020)* (pp. 244–250). https://doi.org/10.18653/v1/2020.coling-main.21
12. Zhang, L., Wang, S., & Liu, B. (2018). Deep learning for sentiment analysis: A survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 8*(4), Article e1253. https://doi.org/10.1002/widm.1253