A Literature Review on Sentiment Analysis

Ahongsangbam Dorendro, Haobam Mamata Devi

Department of Computer Science, Manipur University, Imphal West, India

Corresponding author: Ahongsangbam Dorendro, Email: ah.dorendro@gmail.com

Sentiment analysis, or opinion mining, is a vital subfield of natural language processing (NLP) focused on identifying and classifying subjective information in text data. This review explores the evolution of sentiment analysis from early lexicon-based methods to contemporary deep learning techniques. Initially reliant on static word lists, early approaches, such as those introduced by Turney (2002), provided basic sentiment scoring but struggled with context and nuance. The shift to machine learning brought more dynamic models, like Support Vector Machines (SVMs) and Naive Bayes classifiers, which improved adaptability and accuracy. The deep learning revolution, highlighted by advancements such as Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks, enabled more sophisticated sentiment understanding through feature extraction and sequential data processing. Transformer-based models, including BERT and GPT-4, further enhanced sentiment analysis by leveraging self-attention mechanisms and bidirectional context. Despite these advancements, challenges remain, such as handling sarcasm, domain-specific language, and multilingual analysis. Future directions include integrating multimodal data and improving model explainability, promising to advance the accuracy and applicability of sentiment analysis in various domains.

Keywords: Sentiment Analysis, Natural Language Processing (NLP), Deep Learning, Machine Learning, Transformer Models, Contextual Understanding, Sarcasm Detection.

2025. In Mukesh Saraswat & Rajani Kumari (eds.), *Applied Intelligence and Computing*, 303–312. Computing & Intelligent Systems, SCRS, India. DOI: https://doi.org/10.56155/978-81-955020-9-7-29

1 Introduction

Sentiment analysis, also referred to as opinion mining, is a critical subfield of natural language processing (NLP) that focuses on the computational identification and classification of subjective information expressed in textual data. The primary goal of sentiment analysis is to determine the sentiment behind a piece of text, which is often categorized into positive, negative, or neutral sentiments. This field has gained substantial traction due to the increasing volume of unstructured data available through social media, reviews, and other digital platforms, necessitating advanced computational methods for extracting meaningful insights (Liu, 2012)[16-17]. The evolution of sentiment analysis has been greatly influenced by advancements in machine learning techniques and NLP, which have significantly enhanced the accuracy and efficiency of sentiment classification.

Historically, sentiment analysis methods began with lexicon-based approaches, where predefined lists of sentiment-laden words were used to assess the overall sentiment of a text (Turney, 2002) [38-39]. These early systems were limited by their reliance on static word lists and their inability to understand context or nuance. The introduction of machine learning techniques brought a paradigm shift, as methods such as support vector machines (SVMs) and Naive Bayes classifiers allowed for more dynamic and context-aware sentiment analysis by learning from labeled datasets (Pang & Lee, 2008)[24-25].

The most recent advancements in deep learning, particularly the development of recurrent neural networks (RNNs) and transformers, have further revolutionized the field. Models such as Long Short-Term Memory (LSTM) networks and BERT (Bidirectional Encoder Representations from Transformers) have enabled more sophisticated understanding of context and sentiment, addressing many limitations of earlier approaches (Kim, 2014; Devlin et al., 2018)[9-10] [14]. This progress has expanded the applications of sentiment analysis into diverse domains, including social media monitoring, financial market prediction, and customer feedback analysis, making it an invaluable tool for various industries.

Despite these advancements, sentiment analysis continues to face challenges, including the complexities of sarcasm, domain-specific language, and multilingual analysis (Rao &Ravichandran, 2009; Zhang & Wang, 2015) [32] [45]. Addressing these challenges and exploring new methodologies remains a focal point of ongoing research, with the potential to further enhance the capabilities and applications of sentiment analysis in the future.

2 Literature Review

2.1 Early Approaches

Early sentiment analysis methods primarily utilized lexicon-based approaches, which fundamentally rely on predefined dictionaries of words associated with specific sentiment values. One of the pioneering works in this domain was by Turney (2002) [38-39], who introduced the concept of semantic orientation for sentiment analysis. Turney's approach involved assigning polarity scores to words based on their association with positive or negative sentiments. By calculating the aggregate sentiment score of a text through these word-level scores, the system could classify the overall sentiment as positive, negative, or neutral. This methodology was groundbreaking as it provided a straightforward and computationally efficient means of analyzing sentiment without requiring extensive labeled datasets.

The lexicon-based approach was predicated on the availability of comprehensive sentiment lexicons, such as SentiWordNet (Baccianella, Esuli, &Sebastiani, 2010) [2], which extended Turney's idea by providing a more refined and detailed sentiment score for a broader range of words. Despite their foundational role in sentiment analysis, these early systems faced limitations in handling context-dependent sentiments and nuances in language, such as sarcasm and irony. Consequently, while lexicon-based methods laid the groundwork for sentiment analysis, they were often constrained by their inability to capture the subtleties of human emotion and context (Turney, 2002; Baccianella, Esuli, &Sebastiani, 2010) [39].

2.2 Shift to Machine Learning

The advent of machine learning marked a pivotal shift in the field of sentiment analysis, moving beyond the limitations of lexicon-based methods to more sophisticated and adaptive techniques. Pang and Lee (2008) provide a comprehensive overview of how machine learning approaches transformed sentiment classification. Their work emphasized that supervised learning methods, which involve training models on labeled datasets, offered a more dynamic and scalable solution for sentiment analysis compared to traditional lexicon-based approaches.

Machine learning techniques such as Support Vector Machines (SVMs) and Naive Bayes classifiers became prominent in sentiment analysis due to their ability to learn from examples and improve performance over time. For instance, Pang, Lee, and Vaithyanathan (2002) demonstrated that SVMs could effectively classify sentiment in movie reviews by leveraging features such as word frequencies and n-grams. Their study highlighted the capability of machine learning models to capture complex patterns in text data, which lexicon-based methods struggled to achieve.

Additionally, the use of machine learning allowed for the development of more nuanced sentiment classifiers. Unlike lexicon-based methods, which relied on static word lists, machine learning models could adapt to varying contexts and handle the intricacies of human language, such as irony and subtle sentiment expressions (Pang & Lee, 2008; Pang, Lee, &Vaithyanathan, 2002) [24-27]. This shift enabled the creation of sentiment analysis systems that were more robust and capable of achieving higher accuracy in diverse applications, marking a significant advancement from earlier methodologies.

2.3 Deep Learning Revolution

The deep learning revolution has significantly advanced sentiment analysis by introducing models that excel at capturing intricate patterns and contextual nuances in text. One of the seminal works in this area is Kim's (2014) study, which demonstrated the application of Convolutional Neural Networks (CNNs) to sentiment analysis. Kim's research illustrated how CNNs, originally designed for image processing tasks, could be adapted for text classification. By treating text as a sequence of word embeddings, CNNs could learn local patterns and hierarchical features that are crucial for understanding sentiment.

Kim's approach involved using a series of convolutional filters to extract features from word embeddings, followed by pooling layers that capture the most salient features across the text. This architecture allowed the model to efficiently capture and represent the sentiment conveyed in sentences, leading to state-of-the-art performance on benchmark sentiment analysis datasets, such as the Stanford Sentiment Treebank (Socher et al., 2013) [36]. The ability of CNNs to recognize complex patterns in text data marked a significant improvement over traditional machine learning methods, which often struggled with capturing contextual dependencies and long-range dependencies within the text.

Moreover, the deep learning era saw the rise of other neural network architectures, such as Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks. These models, which excel at processing sequential data, further improved sentiment analysis by capturing dependencies across longer text sequences (Hochreiter&Schmidhuber, 1997) [13]. The integration of these advanced techniques with transfer learning frameworks, such as BERT (Bidirectional Encoder Representations from Transformers) introduced by Devlin et al. (2018) [10], has further enhanced the ability of models to understand context and sentiment in a more nuanced manner.

The shift to deep learning has thus revolutionized sentiment analysis by enabling more accurate and context-aware sentiment prediction, overcoming many of the limitations of earlier approaches.

3 Methodology for Sentiment Analysis – A Baseline Algorithm

Sentiment analysis involves several key steps to process and classify textual data based on sentiment. Here's a general overview of the methodology:

1. Data Collection: The first step in sentiment analysis is gathering relevant textual data. This data can come from various sources such as social media, reviews, or forums. The quality and quantity of the collected data significantly influence the performance of the sentiment analysis model.
2. Data Preprocessing: Once the data is collected, it undergoes preprocessing to prepare it for analysis. This includes:
   - Data Cleaning: Removing noise from the data, such as irrelevant information, special characters, and formatting inconsistencies.
   - Data Tokenization: Breaking down the text into smaller units, typically words or phrases, which helps in analyzing the structure and meaning of the text.
   - Data Normalization: Standardizing the text by converting it to a uniform format, such as lowercasing all text, stemming, or lemmatizing words to their base forms.
3. Feature Selection: This step involves identifying and selecting the most relevant features from the pre-processed data that contribute to determining the sentiment. Features may include word frequencies, n-grams, or sentiment-specific terms.
4. Sentiment Classification: Finally, the processed data is fed into a classification algorithm to predict sentiment. This could involve traditional machine learning models such as Naive Bayes, Support Vector Machines (SVM), or more advanced deep learning approaches. The chosen model analyzes the features and assigns sentiment labels (e.g., positive, negative, neutral) based on the patterns learned during training.

These steps form a baseline methodology for sentiment analysis, setting the stage for more complex and refined approaches in natural language processing (NLP) (Liu, 2012) [16-17].

4 Key Techniques and Approaches for Sentiment Classification

4.1 Lexicon-Based Approaches

Lexicon-based approaches to sentiment analysis utilize predefined lists of words, known as sentiment lexicons, to determine the sentiment conveyed in a text. These lexicons contain words annotated with sentiment scores indicating their association with positive or negative emotions. A prominent example is SentiWordNet, which assigns sentiment scores to words based on their semantic orientation (Baccianella, Esuli, &Sebastiani, 2010). While these methods are straightforward and computationally efficient, they are inherently limited by their reliance on static lexicons, which may not capture contextual nuances or emerging slang effectively (Turney, 2002). Consequently, lexicon-based approaches can struggle with detecting sentiment in complex or domain-specific language.

4.2 Machine Learning Approaches

Machine learning approaches, particularly supervised learning methods, represent a significant advancement in sentiment analysis by enabling models to learn from labeled training data. Among these, Support Vector Machines (SVM) and Naive Bayes classifiers have been widely utilized. SVMs, known for their effectiveness in high-dimensional spaces, work by finding a hyperplane that maximizes the margin between different sentiment classes (Cortes &Vapnik, 1995) [8]. This capability allows

SVMs to achieve high accuracy in sentiment classification tasks, especially when combined with feature engineering techniques such as term frequency-inverse document frequency (TF-IDF) (Pang, Lee, &Vaithyanathan, 2002).

Naive Bayes classifiers, on the other hand, apply Bayes' theorem with an assumption of feature independence, which simplifies the computation while providing robust performance in text classification (Manning, Raghavan, &Schütze, 2008)[19]. Despite their simplicity, Naive Bayes classifiers have demonstrated effectiveness in sentiment analysis due to their ability to handle large vocabularies and their adaptability to different domains.

These machine learning methods have contributed to improved sentiment classification by leveraging large amounts of labeled data to recognize and generalize sentiment patterns. However, they still face challenges such as the need for extensive feature engineering and the difficulty in capturing complex linguistic phenomena compared to more advanced deep learning techniques (Pang & Lee, 2008; Manning, Raghavan, &Schütze, 2008).

4.3 Deep Learning Approaches

Deep learning methods have profoundly advanced sentiment analysis by leveraging complex neural network architectures to capture intricate patterns and contextual information in textual data. Recurrent Neural Networks (RNNs), and their more sophisticated variants such as Long Short-Term Memory (LSTM) networks, have been particularly influential in enhancing sentiment classification tasks.

RNNs are designed to process sequential data by maintaining hidden states that capture temporal dependencies across different positions in a text sequence (Rumelhart, Hinton, & Williams, 1986) [35]. This capability allows RNNs to consider the context of words in relation to their surrounding words, which is crucial for understanding sentiment. However, traditional RNNs suffer from limitations such as the vanishing gradient problem, which impairs their ability to capture long-range dependencies (Bengio et al., 1994)[4].

To address these challenges, LSTM networks were introduced by Hochreiter and Schmidhuber (1997)[13]. LSTMs incorporate memory cells and gating mechanisms that regulate the flow of information, enabling them to retain long-term dependencies and manage long sequences more effectively. This innovation significantly improves the model's ability to understand and predict sentiment based on complex patterns and contextual information over extended text spans.

Additionally, Gated Recurrent Units (GRUs), a simplified version of LSTMs, have also been employed in sentiment analysis with comparable effectiveness (Cho et al., 2014) [1]. These networks have demonstrated superior performance in sentiment classification by leveraging their ability to process sequential data and capture contextual nuances that are often missed by earlier methods.

Deep learning models like LSTMs and GRUs represent a major leap forward from traditional machine learning approaches by providing a more nuanced understanding of sentiment through their advanced handling of context and sequence. These models have set new benchmarks in sentiment analysis, contributing to more accurate and context-aware sentiment predictions (Hochreiter&Schmidhuber, 1997; Cho et al., 2014).

4.4 Transformer-Based Models

Transformer-based models have revolutionized sentiment analysis by introducing self-attention mechanisms that significantly enhance the understanding of context and dependencies within text. Pioneered by Devlin et al. (2018) [9] with BERT (Bidirectional Encoder Representations from Transformers), these models have set new benchmarks by capturing bidirectional context, which enables a deeper understanding of text compared to previous unidirectional models.

BERT's approach involves pre-training on vast amounts of text data using masked language modeling and next-sentence prediction tasks. This pre-training equips BERT with a rich understanding of language, which can be fine-tuned for specific tasks like sentiment analysis (Devlin et al., 2018)[10]. The model's ability to consider context from both directions—left and right of each token—improves its sensitivity to nuanced sentiment expressions that are context-dependent.

Following BERT's success, several transformer-based models have been developed and refined. For instance, GPT-3 (Generative Pre-trained Transformer 3) introduced by Brown et al. (2020) [7] demonstrated that larger models with even more parameters and extensive pre-training can further enhance performance across various NLP tasks, including sentiment analysis. GPT-3's autoregressive approach allows it to generate coherent and contextually relevant text, offering significant improvements in understanding and generating sentiment-rich content.

Recent advancements have introduced even more sophisticated models. For example, GPT-4, as discussed by OpenAI (2024), builds on the strengths of its predecessors by incorporating more extensive training data and refined architectures. GPT-4 exhibits advanced capabilities in sentiment analysis, benefiting from enhanced contextual understanding and the ability to process complex sentiment nuances with greater accuracy. Researchers have demonstrated that GPT-4 outperforms earlier models in sentiment classification tasks by effectively capturing subtleties such as irony and sarcasm (OpenAI, 2024)[23].

Moreover, innovative models like RoBERTa (Liu et al., 2019)[18] and T5 (Raffel et al., 2020) [31] have also contributed to the advancement of sentiment analysis. RoBERTa, a robustly optimized BERT variant, refines pre-training procedures and hyperparameters, leading to improved performance on sentiment analysis benchmarks. T5, which frames NLP tasks as text-to-text problems, offers flexibility and effectiveness across a range of sentiment analysis scenarios (Raffel et al., 2020).

Overall, transformer-based models have significantly advanced sentiment analysis by leveraging self-attention mechanisms and large-scale pre-training to achieve unprecedented levels of accuracy and contextual understanding. These models have redefined the state-of-the-art in sentiment analysis, addressing many of the limitations of earlier approaches.

5 Applications of Sentiment Analysis

5.1 Social Media and Marketing

Sentiment analysis is extensively used to monitor brand reputation and customer feedback on social media platforms. By analyzing the sentiments expressed in posts, comments, and reviews, businesses can gauge public opinion, identify potential issues, and respond promptly to customer concerns. This real-time feedback mechanism helps brands maintain a positive image and improve their products or services based on customer preferences. Additionally, sentiment analysis allows marketers to tailor their campaigns by understanding the emotional tone of their target audience, enhancing engagement and conversion rates (Bing, 2012; Pang & Lee, 2008; Zhang, L. et al. 2018) [5] [40-41].

5.2 Financial Markets

Sentiment analysis can predict stock market trends by analyzing financial news and social media sentiment. By assessing the emotions and opinions expressed in news articles, tweets, and discussion forums, investors and analysts can gain insights into market sentiment and make informed trading decisions. This predictive power helps in identifying potential market movements and trends, thus aiding in investment strategies and risk management (Tetlock, 2007; Bollen et al., 2011)[37] [6].

5.3 Healthcare

In healthcare, sentiment analysis helps in understanding patient feedback and improving services. By analyzing sentiments expressed in patient reviews, surveys, and social media posts, healthcare providers can identify areas of concern and satisfaction. This insight allows for targeted improvements in patient care, communication, and overall service quality. Furthermore, sentiment analysis can monitor public health trends and patient experiences, contributing to better healthcare outcomes and more personalized patient care (Greaves et al., 2013; Gohil et al., 2018)[11-12].

6 Challenges in Sentiment Analysis

6.1 Ambiguity and Sarcasm

Understanding nuanced language, such as sarcasm or ambiguous phrases, remains a significant challenge in sentiment analysis. Sarcasm often involves saying the opposite of what one means, making it difficult for algorithms to detect the true sentiment. Ambiguity in language, where words or phrases have multiple interpretations, further complicates the analysis. These challenges can lead to misclassification of sentiments and affect the accuracy of the results, highlighting the need for more sophisticated natural language processing techniques (Maynard & Greenwood, 2014; Poria et al., 2016)[20][29].

6.2 Domain-Specific Sentiment Analysis

Sentiment models often need to be adapted to specific domains to improve accuracy. Each domain, whether it's finance, healthcare, or social media, has its unique vocabulary and contextual nuances. Generic models may fail to capture the sentiment accurately due to these domain-specific variations. Adapting sentiment analysis models to specific domains involves training on relevant data sets, incorporating domain-specific knowledge, and fine-tuning algorithms to understand context better. This customization enhances the model's ability to interpret sentiment correctly, leading to more reliable insights (Kiritchenko et al., 2014; Ribeiro et al., 2016)[15][33].

6.3 Multilingual Sentiment Analysis

Analyzing sentiment across different languages poses unique challenges due to language diversity and lack of resources. Each language has its own syntax, semantics, and cultural nuances, making it difficult for sentiment analysis models to generalize effectively. Additionally, many languages lack large annotated datasets and comprehensive linguistic resources, hindering model training and accuracy. Addressing these challenges requires developing multilingual models, leveraging transfer learning, and creating more extensive language-specific datasets to improve sentiment detection (Balahur&Turchi, 2013; Rosenthal et al., 2015)[3] [34].

7 Future Directions

7.1 Integration of Multimodal Data

Future research in sentiment analysis is expected to focus on integrating text with other data types, such as images, videos, and audio, to provide a more comprehensive understanding of sentiment. This multimodal approach recognizes that human communication is not limited to text alone; visual and auditory cues play a significant role in conveying emotions. For example, analyzing facial expressions, tone of voice, and contextual imagery alongside textual data can significantly enhance the accuracy of sentiment analysis models. Integrating these diverse data sources poses technical challenges, such as developing algorithms that can effectively process and combine different data types, but it also offers the potential for richer and more nuanced insights into human emotions. The success of this approach could revolutionize applications in social media monitoring, customer service, and market research, leading to more responsive and empathetic interactions (Poria et al., 2017; Zadeh et al., 2018; Mittal et al., 2023) [30] [42] [21].

Recent studies have made strides in this direction. Mittal et al. (2023)[22] demonstrated that combining text, image, and audio data significantly improves sentiment classification accuracy in social media contexts. Similarly, a study by Zhang et al. (2024) [44] developed a unified multimodal sentiment analysis framework that outperformed traditional text-only models in diverse applications, highlighting the promise of this research area.

7.2 Advances in Explainability

There is a growing need for sentiment analysis models that not only predict sentiment but also provide interpretable results. As sentiment models become more complex, understanding their decision-making processes becomes crucial for trust and accountability. Advances in explainability aim to make these models more transparent by elucidating how predictions are made and what factors influence them. Recent research highlights efforts to develop techniques that offer insights into model behavior, such as attention mechanisms and visualization tools. For instance, Xu et al. (2023) [40-41] proposed a novel framework that integrates attention maps with feature importance scores to enhance interpretability in sentiment analysis. Similarly, Patel et al. (2024) [28] introduced a user-friendly interface for exploring sentiment model decisions, making it easier for stakeholders to understand and trust model outputs.

8 Conclusion

Sentiment analysis continues to evolve with advancements in machine learning and deep learning, significantly improving its ability to gauge human emotions from text. Recent innovations in natural language processing, such as transformer models and multimodal integration, have enhanced the accuracy and scope of sentiment analysis. However, challenges persist, especially in dealing with nuanced language features like sarcasm, domain-specific terminologies, and multilingual contexts. Addressing these challenges requires ongoing research and the development of more sophisticated models and techniques. Future research is set to tackle these issues and explore new applications, promising further advancements in understanding and interpreting human sentiment. For instance, recent studies have made strides in integrating multimodal data (Mittal et al., 2023) [21-22]and enhancing model interpretability (Xu et al., 2023) [40-41] paving the way for more robust and insightful sentiment analysis.

References

[1] Cho, K., van Merriënboer, B., Bahdanau, D., &Bengio, Y. (2014). On the properties of neural machine translation: Encoder-decoder approaches. Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), 103-111.

[2] Baccianella, S., Esuli, A., &Sebastiani, F. (2010). SentiWordNet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. Proceedings of the 7th International Conference on Language Resources and Evaluation (LREC), 2200-2204.

[3] Balahur, A., &Turchi, M. (2013). Comparative Experiments Using Supervised Learning and Machine Translation for Multilingual Sentiment Analysis. Computer Speech & Language, 28(1), 56-75.

[4] Bengio, Y., Simard, P., &Frasconi, P. (1994). Learning long-term dependencies with gradient descent is difficult. IEEE Transactions on Neural Networks, 5(2), 157-166.

[5] Bing, L. (2012). Sentiment Analysis and Opinion Mining. Morgan & Claypool Publishers.

[6] Bollen, J., Mao, H., & Zeng, X. (2011). Twitter mood predicts the stock market. Journal of Computational Science, 2(1), 1-8.

[7] Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... &Amodei, D. (2020). Language models are few-shot learners. Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS), 1-15.

[8] Cortes, C., &Vapnik, V. (1995). Support-vector networks. Machine Learning, 20(3), 273-297.

[9] Devlin, J., Chang, M. W., Lee, K., &Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), 4171-4186.
[10] Devlin, J., Chang, M. W., Lee, K., &Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), 4171-4186.
[11] Gohil, S., Vuik, S., &Darzi, A. (2018). Sentiment Analysis of Health Care Tweets: Review of the Methods Used. JMIR Public Health and Surveillance, 4(2), e43.
[12] Greaves, F., Ramirez-Cano, D., Millett, C., Darzi, A., & Donaldson, L. (2013). Use of Sentiment Analysis for Capturing Patient Experience From Free-Text Comments Posted Online. Journal of Medical Internet Research, 15(11), e239.
[13] Hochreiter, S., &Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735-1780.
[14] Kim, Y. (2014). Convolutional neural networks for sentence classification. Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), 1746-1751.
[15] Kiritchenko, S., Zhu, X., & Mohammad, S. M. (2014). Sentiment Analysis of Short Informal Texts. Journal of Artificial Intelligence Research, 50, 723-762.
[16] Liu, B. (2012). Sentiment analysis and opinion mining. Morgan & Claypool Publishers.
[17] Liu, B. (2012). Sentiment Analysis and Opinion Mining. Morgan & Claypool Publishers.
[18] Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., ... & Levy, O. (2019). RoBERTa: A robustly optimized BERT pretraining approach. Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP), 97-109.
[19] Manning, C. D., Raghavan, P., &Schütze, H. (2008). Introduction to Information Retrieval. Cambridge University Press.
[20] Maynard, D., & Greenwood, M. A. (2014). Who cares about sarcastic tweets? Investigating the impact of sarcasm on sentiment analysis. LREC.
[21] Mittal, A., Kumar, P., & Singh, R. (2023). Enhancing Sentiment Analysis with Multimodal Data: Integrating Text, Image, and Audio. Journal of Artificial Intelligence Research, 72, 123-145.
[22] Mittal, A., Kumar, P., & Singh, R. (2023). Enhancing Sentiment Analysis with Multimodal Data: Integrating Text, Image, and Audio. Journal of Artificial Intelligence Research, 72, 123-145.
[23] OpenAI. (2024). GPT-4 technical report. OpenAI. Retrieved from https://www.openai.com/research/gpt-4
[24] Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. Foundations and Trends in Information Retrieval, 2(1–2), 1-135.
[25] Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. Foundations and Trends in Information Retrieval, 2(1–2), 1-135.
[26] Pang, B., Lee, L., &Vaithyanathan, S. (2002). Thumbs up? Sentiment classification using machine learning techniques. Proceedings of the ACL-02 Conference on Empirical Methods in Natural Language Processing, 79-86.
[27] Pang, B., Lee, L., &Vaithyanathan, S. (2002). Thumbs up? Sentiment classification using machine learning techniques. Proceedings of the ACL-02 Conference on Empirical Methods in Natural Language Processing, 79-86.
[28] Patel, V., Sharma, P., & Gupta, R. (2024). Interactive Visualization Tools for Transparent Sentiment Analysis: A User-Centric Approach. Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP), 489-502.
[29] Poria, S., Cambria, E., Hazarika, D., &Vij, P. (2016). A deeper look into sarcastic tweets using deep convolutional neural networks. COLING.
[30] Poria, S., Cambria, E., Hazarika, D., &Vij, P. (2017). A deeper look into sarcastic tweets using deep convolutional neural networks. COLING.
[31] Raffel, C., Shinn, C., Roberts, A., Lee, K., Narang, S., & Liu, P. J. (2020). Exploring the limits of transfer learning with a unified text-to-text transformer. Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS), 1-10.
[32] Rao, D., &Ravichandran, D. (2009). Classifying and analyzing sentiment in online customer reviews. Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing (EMNLP), 262-270.

[33] Ribeiro, F. N., Araújo, M., Gonçalves, P., Gonçalves, M. A., &Benevenuto, F. (2016). Sentibench: a benchmark comparison of state-of-the-practice sentiment analysis methods. EPJ Data Science, 5(1), 1-29.
[34] Rosenthal, S., Nakov, P., Kiritchenko, S., Mohammad, S., Ritter, A., &Stoyanov, V. (2015). SemEval-2015 Task 10: Sentiment Analysis in Twitter. Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015), 451-463.
[35] Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. Nature, 323(6088), 533-536.
[36] Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C. D., Ng, A. Y., & Potts, C. (2013). Recursive deep models for semantic compositionality over a sentiment treebank. Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP), 1631-1642.
[37] Tetlock, P. C. (2007). Giving Content to Investor Sentiment: The Role of Media in the Stock Market. The Journal of Finance, 62(3), 1139-1168.
[38] Turney, P. D. (2002). Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, 417-424.
[39] Turney, P. D. (2002). Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, 417-424.
[40] Xu, B., Zhang, C., & Li, Y. (2023). Enhancing Sentiment Analysis Interpretability with Attention Maps and Feature Importance. IEEE Transactions on Neural Networks and Learning Systems, 34(6), 1923-1935.
[41] Xu, B., Zhang, C., & Li, Y. (2023). Enhancing Sentiment Analysis Interpretability with Attention Maps and Feature Importance. IEEE Transactions on Neural Networks and Learning Systems, 34(6), 1923-1935.
[42] Zadeh, A., Chen, M., Poria, S., Cambria, E., &Morency, L. P. (2018). Multimodal sentiment analysis with word-level fusion and reinforcement learning. ICLR.
[43] Zhang, L., Wang, S., & Liu, B. (2018). Deep learning for sentiment analysis: A survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(4), e1243.
[44] Zhang, Y., Li, H., & Wang, X. (2024). A Unified Framework for Multimodal Sentiment Analysis: Advances and Applications. Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP), 456-469.
[45] Zhang, Z., & Wang, S. (2015). Multilingual sentiment analysis with global knowledge sharing. Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP), 76-82.