



International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 8.699

Volume 14, Issue 2, February 2025

Comprehensive Study on Sentiment Analysis: From Rule-based to Modern LLM-based Systems

Shailja Gupta*, Rajesh Ranjan*, Surya Narayan Singh*

Carnegie Mellon University, USA

Carnegie Mellon University, USA

BIT Sindri, India

ABSTRACT: This paper presents a comprehensive survey of sentiment analysis, a pivotal area in natural language processing (NLP) and artificial intelligence (AI), highlighting its evolution from traditional rule-based techniques to cutting-edge approaches leveraging large language models (LLMs). The study systematically explores the historical development of sentiment analysis, emphasizing the transition from lexicon-based methods to advanced deep learning frameworks such as transformers and recurrent neural networks. Key challenges, including sarcasm detection, multilingual text processing, and bias mitigation, are analyzed in depth. Furthermore, the impact of LLMs, including pre-trained models like BERT and GPT, is examined, showcasing their ability to handle nuanced contexts and improve sentiment classification accuracy. Emerging trends, such as multimodal sentiment analysis and the integration of explainable AI, are discussed alongside future research directions to address current limitations. Through the in-depth analysis, the study identifies the gaps and limitations in current techniques and proposes a future research direction in this field. This study is a valuable resource for researchers and practitioners, offering insights into the advancements and opportunities in sentiment analysis driven by AI and LLMs.

KEYWORDS: Sentiment Analysis, Deep Learning, Explainable AI, Large Language Models, Artificial Intelligence.

I. INTRODUCTION

1.1. Background and Motivation

Sentiment analysis, also known as opinion mining, has become a pivotal area of research in natural language processing (NLP) and artificial intelligence (AI). This field focuses on identifying and extracting subjective information from textual data, often to determine the sentiment expressed by individuals or groups. The rapid growth of digital communication platforms has amplified the need for effective sentiment analysis tools to manage and interpret vast amounts of user-generated content. Traditional methods of sentiment analysis, based on manually crafted rules and lexicons, have been increasingly supplemented by advanced machine learning and deep learning approaches which offer more nuanced and scalable solutions (Liu et. al. 2012) (Zhang et. al. 2015)

1.2. Objectives of the Survey

The primary objective of this survey is to provide a comprehensive overview of sentiment analysis techniques, from traditional methods to cutting-edge deep learning models. This paper aims to:

- Review the evolution of sentiment analysis techniques and their impact on the field.
- Highlight key challenges and limitations associated with sentiment analysis, including multilingual texts and sarcasm detection issues.
- Explore the role of large language models (LLMs) in advancing sentiment analysis capabilities.
- Recommend future research directions to address existing challenges and enhance the effectiveness of sentiment analysis.

II. METHODOLOGY

This survey adopts a structured approach to review the evolution of sentiment analysis, focusing on techniques ranging from rule-based methods to advanced large language models (LLMs). Literature was identified from several databases including IEEE, arxiv.org, and others, using keywords like "Sentiment Analysis," "Deep Learning," and "Large Language

Models." Foundational studies work addressing key challenges, and papers on LLM applications were included as part of the survey.

Thematic Categorization: Literature was grouped into:

- Historical methods and advancements.
- Challenges in sentiment analysis (e.g., context, biases).
- Impact of LLMs and future trends.

2. Evolution of Sentiment Analysis

Sentiment analysis is a subfield of natural language processing (NLP) that focuses on determining and analyzing the emotional tone expressed in text. It involves classifying text data such as positive, negative, neutral, or more granular emotional states such as joy, anger, or sadness (Pang et. al. 2008). This process helps in understanding public opinion, customer satisfaction, and various emotional insights from textual data. The development of sentiment analysis can be traced back to the early 2000s when initial approaches were primarily rule-based. Early systems relied on handcrafted lexicons and pattern matching to detect sentiments (Turney et. al. 2003). With the advancement of machine learning, these methods evolved to include statistical models and feature-based classifiers. The introduction of deep learning and neural networks in the 2010s marked a significant shift, allowing for more sophisticated and accurate sentiment analysis by leveraging vast amounts of data and computational power (Kim et. al. 2014). It can be categorized into several types based on the granularity and focus of the analysis:

Document-Level Sentiment Analysis: Classifies the sentiment of entire documents or reviews, often used for summarizing overall opinion (Hu et. al. 2004).

Sentence-Level Sentiment Analysis: Focuses on individual sentences to determine sentiment, providing more detailed insights (Socher et. al. 2013).

Aspect-Based Sentiment Analysis: Analyzes sentiments related to specific aspects or features within the text, such as product reviews focusing on different attributes (Pontiki et. al. 2016)

Sentiment analysis, a crucial area in natural language processing (NLP), aims to detect and interpret emotions expressed in text. However, several challenges hinder its accuracy and effectiveness. As sentiment analysis has evolved, the complexity of linguistic nuances, cultural variations, and context-dependent meanings has posed significant obstacles. This section discusses three key challenges: ambiguity and context, sarcasm and irony, and multilingual or bilingual texts, all of which complicate the accurate interpretation of sentiment

Ambiguity and Context: Ambiguity in language is one of the foremost challenges in sentiment analysis. Words and phrases often carry multiple meanings depending on their context, which makes determining sentiment a complex task. For instance, consider the word "cold." In isolation, this word could indicate a physical temperature or describe someone's distant personality. In another context, it might imply negative emotions like rejection or loneliness. This contextual Dependency creates a significant hurdle for traditional rule-based models, which may struggle to accurately interpret nuanced meanings. The challenge of understanding context is particularly relevant in fields like product reviews, where a word like "light" might refer to the weight of a product in one review and to its brightness in another. This polysemy often results in inaccurate sentiment classification if the system fails to understand the broader context.

Sarcasm and Irony: Detecting sarcasm and irony is another formidable challenge in sentiment analysis. These linguistic constructs rely on the discrepancy between literal and intended meanings, which are often difficult to detect using traditional or even some modern machine learning methods. For instance, the phrase "Oh, great! Another delay!" is sarcastic, with an underlying negative sentiment, despite the use of a seemingly positive word, "great." Sarcasm and irony are particularly challenging because they require the system to not only understand the literal text but also infer the speaker's tone, context, and intent. This task demands a deeper understanding of pragmatics, which goes beyond the capabilities of simple keyword-based or lexicon-based approaches.

Multilingual and Bilingual Texts: Analyzing multilingual and bilingual texts introduces further complexity to sentiment analysis. Sentiment can vary dramatically between languages due to differences in grammar, idioms, and cultural expressions. For example, sentiment conveyed in English may not have a direct equivalent in languages like Hindi, Chinese, or Arabic, resulting in a potential loss of meaning during translation. Moreover, sentiment analysis systems trained in one language often fail to generalize to another, primarily due to the lack of annotated data in many languages. Additionally, cultural differences in expressing emotions can further complicate the analysis. For example, while direct

expressions of dissatisfaction might be common in Western languages, indirect or more formal expressions are preferred in some Asian languages. This disparity makes sentiment detection harder in multilingual contexts. Handling bilingual texts introduces the additional challenge of translation accuracy. Automated translation tools are often employed, but their limitations can distort the sentiment. A misinterpretation of a single word can flip the overall sentiment of a sentence, thus reducing the reliability of sentiment analysis in multilingual settings.

III. TRADITIONAL APPROACHES TO SENTIMENT ANALYSIS

3.1. Rule-Based Systems

Traditional sentiment analysis often relied on rule-based systems, which use predefined rules and linguistic resources to identify and classify sentiments.

Lexicon-Based Approaches: Lexicon-based methods use sentiment lexicons, which are lists of words associated with specific sentiments. Each word in the lexicon is assigned a sentiment score, and the overall sentiment of a text is determined by aggregating the scores of the words it contains. Examples include the AFINN and SentiWordNet lexicons (Baccianella et. al. 2010). These methods are straightforward but limited by the need for comprehensive lexicons and the inability to handle context-dependent meanings effectively.

Pattern-Based Approaches: Pattern-based methods involve creating rules based on syntactic and semantic patterns in the text. These rules can capture specific sentiment-bearing structures, such as phrases or negations, to determine sentiment (Pang et. al. 2008). While these methods can be more flexible than lexicon-based approaches, they often require extensive manual effort to develop and maintain.

3.2. Machine Learning Techniques

The advent of machine learning brought a shift from rule-based methods to approaches that learn from data.

Feature Engineering: In machine learning-based sentiment analysis, features such as term frequency-inverse document frequency (TF-IDF) or word embeddings are used to represent textual data. These features are then fed into classification algorithms to predict sentiment (Joachims et. al. 1998). Feature engineering involves selecting and transforming raw text data into meaningful features to improve model performance.

Classification Algorithms: Various classification algorithms, such as Naive Bayes, Support Vector Machines (SVM), and Logistic Regression, are employed to classify text into sentiment categories (McCallum et. al. 1998). These algorithms learn patterns from labeled training data and apply them to new, unseen text. While effective, these methods often require careful tuning and feature selection to achieve optimal performance.

3.3. Limitations of Traditional Methods

Traditional approaches to sentiment analysis, while pioneering, have several limitations:

Handling complex sentiments: Traditional methods often struggle with complex sentiments, such as mixed or nuanced emotions, as they typically rely on binary or simple categorical classifications. These methods may not capture the full range of emotional expressions in the text.

Generalizability Issues: Rule-based and machine-learning approaches may have difficulty generalizing across different domains or languages. Rule-based systems depend heavily on domain-specific rules, while machine learning models may require extensive retraining to adapt to new contexts or languages.

IV. STATE-OF-THE-ART APPROACHES TO SENTIMENT ANALYSIS

4.1. Deep Learning Approaches in Sentiment Analysis Deep learning has revolutionized sentiment analysis by enabling models to capture complex patterns and contextual information from text data. Neural networks, inspired by the human brain, consist of layers of interconnected nodes or neurons that process input data and learn features through training. Deep learning models, a subset of neural networks with many layers, have significantly advanced sentiment analysis by improving the ability to understand and generate text (LeCun et. al. 2015). **Recurrent Neural Networks (RNNs)** are designed to handle sequential data by maintaining a hidden state that captures information from previous steps in the sequence. This makes them suitable for text analysis, where the order of words and context are crucial. However, standard RNNs suffer from issues such as vanishing and exploding gradients, which can hinder their performance on long sequences (Bengio et. al. 1994). **Long Short-Term Memory (LSTM) Networks** are a type of RNN designed to address the limitations of standard RNNs. They incorporate memory cells and gating mechanisms to regulate the flow of information, allowing the model to retain long-term dependencies and handle longer sequences more

effectively (Hochreiter et. al. 1997). LSTMs have been widely used in sentiment analysis for their improved performance in capturing contextual information and handling complex sentence structures.

4.2. Convolutional Neural Networks (CNNs) Originally developed for image processing, CNNs have been adapted for text analysis. CNNs use convolutional layers to automatically extract features from text data by applying filters that capture local patterns, such as phrases or n-grams (Kim et. al. 2014). This approach is particularly effective for identifying key sentiment-bearing phrases and patterns in text. CNNs are often combined with other models, such as RNNs or LSTMs, to enhance their performance in sentiment classification tasks.

4.3. Transformer Models and Large Language Models (LLMs)

Initially, sentiment analysis was performed using simpler models such as bag-of-words and n-grams. Still, these approaches had limitations in capturing contextual information and semantic nuances (Manning et. al. 2008)]. The advent of neural network-based models brought significant improvements, with embeddings like Word2Vec and GloVe offering richer representations of words by capturing semantic relationships (Mikolov et. al. 2013). However, the real breakthrough came with the introduction of transformer-based models, which leverage self-attention mechanisms to process and understand the text more nuancedly (Vaswani et. al. 2017).

BERT (Bidirectional Encoder Representations from Transformers): BERT is a transformer-based model significantly impacting sentiment analysis. It uses bidirectional attention to consider each word's context in both directions (left-to-right and right-to-left), enabling it to capture nuanced meanings and relationships between words (Devlin et. al. 2019). BERT is pre-trained on a large corpus and can be fine-tuned for specific sentiment analysis tasks, leading to substantial improvements in accuracy and performance. **GPT (Generative Pre-trained Transformer),** developed by OpenAI, is another influential transformer-based model. Unlike BERT, GPT uses a unidirectional approach (left-to-right) for text generation and understanding. The model is pre-trained on diverse internet text and can be fine-tuned for various NLP tasks, including sentiment analysis (Radford et. al. 2018). GPT's ability to generate coherent and contextually relevant text has made it a valuable tool for understanding sentiment and generating responses. Pre-trained LLMs have set new benchmarks in various NLP tasks, including sentiment analysis. These models are trained on vast amounts of text data and learn to generate contextual representations of words and phrases. Recent advancements in **Multimodal LLMs** have also made multimodal sentiment analysis possible, where textual data is integrated with other modalities such as images or audio. This approach aims to capture sentiment more comprehensively by considering multiple sources of information (Radford et. al. 2021). Models like CLIP (Contrastive Language-Image Pretraining) combine text and image data to improve sentiment analysis by leveraging visual context along with textual information. This integration helps in understanding sentiment in a more nuanced way, especially in contexts where visual elements contribute significantly to the overall sentiment (Ramesh et. al. 2021). These recent advancements in technology, particularly transformers, offer significant advantages in understanding the context and nuances of sentiment expression. They can capture subtle emotional tones and complex sentence structures that traditional methods might miss. However, despite their advantages, deep learning models, especially transformers, require substantial computational resources and memory. Training these models involves processing large datasets and using powerful hardware, which can be a limitation for some applications (Vaswani et. al. 2017).

V. LIMITATION OF CURRENT TECHNOLOGY IN SENTIMENT ANALYSIS

Despite the advancements in sentiment analysis, several challenges persist that affect the accuracy and effectiveness of sentiment analysis systems. This section discusses key challenges and potential solutions in sentiment analysis.

5.1. Sentiment Detection in Bilingual and Multilingual Texts

Cross-Lingual Models: Sentiment analysis across different languages presents unique challenges due to variations in linguistic structures, cultural contexts, and sentiment expressions. Cross-lingual models aim to address these issues by leveraging shared semantic spaces or aligning embeddings across languages. Techniques such as multilingual embeddings and transfer learning have been employed to build models that can understand and analyze sentiments in multiple languages.

Domain-Specific Adaptations: Sentiment expression can vary significantly across domains. For example, sentiment related to financial markets may be expressed differently than sentiment in social media. Domain adaptation techniques

involve fine-tuning models on domain-specific data to improve their performance in particular contexts. This adaptation is crucial for achieving accurate sentiment analysis in specialized fields (Bolukbasi et. al. 2016).

5.2. Detection of Sarcasm and Irony

Sarcasm and irony pose significant challenges for sentiment analysis because they often invert the literal meaning of words. Detecting these forms of sentiment requires models to understand context, tone, and sometimes cultural nuances. Approaches such as using context-aware models (e.g., transformers) and incorporating external knowledge sources have shown promise in improving sarcasm detection. Additionally, annotated datasets with sarcastic and ironic examples can help train more robust models.

5.3. Bias and Fairness in Sentiment Analysis

Identifying and Mitigating Biases: Sentiment analysis models can inherit and perpetuate biases present in training data. These biases may relate to demographic factors such as gender, race, or socioeconomic status, leading to unfair or inaccurate sentiment classification. Techniques to identify and mitigate biases include diverse data collection, fairness-aware training algorithms, and bias detection frameworks. Ensuring that models are trained on balanced and representative datasets is crucial for minimizing bias (Gupta et. al. 2024).

Ethical Considerations: The ethical implications of sentiment analysis extend beyond bias. Issues such as privacy, data security, and the potential misuse of sentiment data must be considered. Implementing ethical guidelines and ensuring transparency in how sentiment data is used can help address these concerns. Moreover, user consent and data anonymization practices should be prioritized to protect individual privacy.

VI. FUTURE RESEARCH DIRECTIONS

The field of sentiment analysis is continuously evolving, driven by advancements in technology and the growing demand for sophisticated text understanding. This section explores the emerging trends and future research directions that are likely to shape the landscape of sentiment analysis in the coming years.

6.1. Advances in Deep Learning Architectures

New Architectures and Techniques: The development of novel deep learning architectures promises to enhance sentiment analysis capabilities. Techniques such as graph neural networks (GNNs) and attention mechanisms are being explored for their potential to capture complex relationships and contextual information in text. These advancements could lead to more accurate and nuanced sentiment analysis by better understanding the interplay between different elements of a text (Wu et. al. 2021).

Enhanced Transfer Learning: Transfer learning continues to be a significant area of research, with a focus on improving the effectiveness of pre-trained models for sentiment analysis. Advances in few-shot learning and zero-shot learning techniques could enable models to perform sentiment analysis with minimal labeled data, making them more adaptable to new domains and languages (Brown et. al. 2020).

6.2. Enhanced Handling of Context and Nuance

Contextualized Sentiment Analysis: Future research is likely to focus on improving the ability of sentiment analysis models to handle nuanced contexts. This includes a better understanding of sarcasm, irony, and subtle emotional cues. Innovations in context-aware models and multi-turn dialogue systems may provide deeper insights into the sentiment expressed in complex or conversational texts (Zhang et. al. 2020).

Emotion Detection: Beyond simple sentiment classification, there is growing interest in detecting specific emotions such as joy, anger, or sadness. This requires models that can distinguish between different emotional states and understand their implications in various contexts. Advances in emotion-aware sentiment analysis could provide more detailed insights into user sentiments and reactions (Chutia et. al. 2024).

Cross-Linguistic and Cross-Cultural Analysis: Future research should explore cross-linguistic and cross-cultural aspects of sentiment analysis to improve models' effectiveness across diverse languages and cultures. This includes developing techniques for handling multilingual data and understanding cultural variations in sentiment expressions (Miah et. al., 2024).

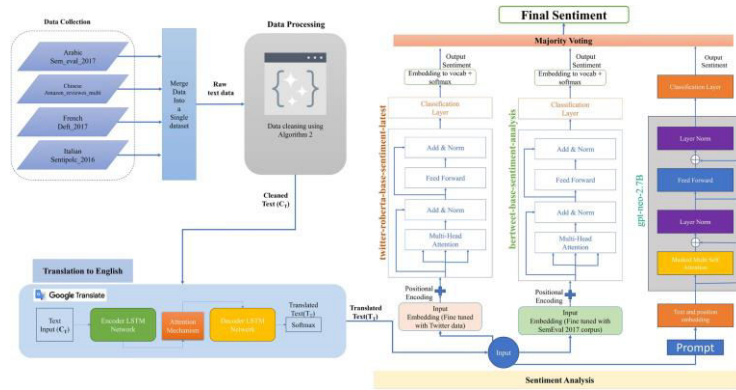


Figure 1: Overview of a multimodal approach to cross-lingual sentiment analysis (Miah et. al. 2024)

Explainable AI: Future research should explore ways to enhance the interpretability and transparency of sentiment analysis models, which is an important research direction. Explainable AI techniques (Uddin et. al. 2024) can help users understand how models arrive at their conclusions, increasing trust and accountability in sentiment analysis applications.

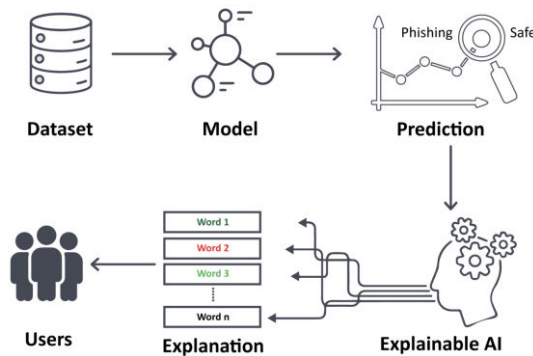


Figure 2: Overview of an Explainable Transformer-based Model (Uddin et. al., 2024)

Human-AI Collaboration: The integration of human expertise with AI-driven sentiment analysis can lead to more accurate and contextually relevant results. Future research should focus on developing systems that facilitate effective collaboration between human analysts and AI models, combining the strengths of both to improve sentiment analysis outcomes.

VII. CONCLUSION

This survey comprehensively examines the evolution of sentiment analysis, charting its development from rule-based systems to the transformative impact of large language models (LLMs). By analyzing historical milestones, state-of-the-art approaches, and key challenges such as ambiguity, sarcasm detection, and multilingual text processing, this study highlights the pivotal role of advanced techniques like transformers in enhancing sentiment analysis capabilities. The integration of multimodal data and explainable AI methodologies underscores the potential for more robust, interpretable, and context-aware sentiment analysis systems. However, significant challenges persist, including addressing cultural nuances, mitigating biases, and improving adaptability across domains and languages. Future research must prioritize these challenges while exploring innovative directions, such as multimodal sentiment analysis and emotion-specific classification. Addressing these gaps will further empower sentiment analysis as a critical tool in natural language processing and artificial intelligence, enabling its applications to expand across industries and use cases. This work serves as a valuable resource for researchers and practitioners, offering a structured understanding of sentiment analysis and a

roadmap for advancing the field. By leveraging the insights and state-of-the-art techniques presented, the community can continue to push the boundaries of what is possible in sentiment analysis.

REFERENCES

1. 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 2010).
2. 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.
3. Bolukbasi, T., Chang, K., Zou, J., Saligrama, V., & Kalai, A. (2016). Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings. ArXiv. <https://arxiv.org/abs/1607.06520>
4. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., and Sutskever, I.,(2020). Language Models are Few-Shot Learners. Proceedings of the Conference on Neural Information Processing Systems (NeurIPS 2020).
5. Chutia, T., Baruah, N. A review on emotion detection by using deep learning techniques. Artif Intell Rev 57, 203 (2024)
6. Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, 1 (Long and Short Papers), 4171–4186.
7. Gupta, S., & Ranjan, R. (2024). Evaluation of LLMs Biases Towards Elite Universities: A Persona-Based Exploration. ArXiv. <https://arxiv.org/abs/2407.12801>
8. Hochreiter, S., & Schmidhuber, J. (1997). Long Short-Term Memory. Neural Computation, 9(8), 1735–1780.
9. Hu, M., & Liu, B. (2004). Mining and Summarizing Customer Reviews. Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2004).
10. Joachims, T. (1998). Text Classification with Support Vector Machines: Learning with Many Relevant Features. Proceedings of the 10th European Conference on Machine Learning (ECML 1998).
11. Kim, Y. (2014). Convolutional Neural Networks for Sentence Classification. Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP 2014).
12. Kim, Y. (2014). Convolutional Neural Networks for Sentence Classification. Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP 2014).
13. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep Learning. Nature, 521(7553), 436–444.
14. Liu, B. (2012). Sentiment Analysis and Opinion Mining. Synthesis Lectures on Human Language Technologies, 5(1), 1–167.
15. Manning, C. D., Raghavan, P., & Schütze, H. (2008). Introduction to Information Retrieval. Cambridge University Press.
16. McCallum, A., & Nigam, K. (1998). A Comparison of Event Models for Naive Bayes Text Classification. AAAI Conference on Artificial Intelligence.
17. Miah, M. S., Kabir, M. M., Sarwar, T. B., Safran, M., Alfarhood, S., & Mridha, M. F. (2024). A multimodal approach to cross-lingual sentiment analysis with ensemble of transformer and LLM. Scientific Reports, 14(1), 1-18.
18. Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed Representations of Words and Phrases and Their Compositionality. Proceedings of the 26th Annual Conference on Neural Information Processing Systems (NeurIPS 2013).
19. Pang, B., & Lee, L. (2008). Opinion Mining and Sentiment Analysis. Foundations and Trends® in Information Retrieval, 2, 1–135.
20. Pontiki, M., Galanis, D., Papageorgiou, H., Androutsopoulos, I., Manandhar, S., Al-Smadi, M., Al-Ayyoub, M., Zhao, Y., Qin, B., De Clercq, O., Hoste, V., Apidianaki, M., Tannier, X., Loukachevitch, N., Kotelnikov, E., Bel, N., Jiménez-Zafra, S. M., & Eryiğit, G. (2016). Semeval-2016 Task 5: Aspect Based Sentiment Analysis. Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval 2016).
21. Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., and Sutskever, I. (2021). Learning Transferable Visual Models From Natural Language Supervision. Proceedings of the 2021 Conference on Computer Vision and Pattern Recognition (CVPR 2021).

22. Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving Language Understanding by Generative Pre-Training. OpenAI Blog.
23. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., and Sutskever, I. (2021). Zero-Shot Text-to-Image Generation. Proceedings of the 2021 Conference on Computer Vision and Pattern Recognition (CVPR 2021).
24. Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C. D., Ng, A., & 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 2013).
25. Turney, P. D., & Littman, M. L. (2003). Measuring Praise and Criticism: Inference of Semantic Orientation from Association. ACM Transactions on Information Systems (TOIS), 21(4), 315–346.
26. Uddin, M. A., & Sarker, I. H. (2024). An Explainable Transformer-based Model for Phishing Email Detection: A Large Language Model Approach. ArXiv. /abs/2402.13871
27. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is All You Need. Proceedings of the 31st International Conference on Neural Information Processing Systems (NeurIPS 2017).
28. Wu, L., Chen, Y., Shen, K., Guo, X., Gao, H., Li, S., Pei, J., & Long, B. (2021). Graph Neural Networks for Natural Language Processing: A Survey. ArXiv. <https://arxiv.org/abs/2106.06090>
29. Zhang, X., Zhao, J., & LeCun, Y. (2015). Character-level Convolutional Networks for Text Classification. Proceedings of the 28th International Conference on Neural Information Processing Systems (NIPS 2015).
30. Zhang, Z., Zhao, H., & Wang, R. (2020). Machine Reading Comprehension: The Role of Contextualized Language Models and Beyond. ArXiv. <https://arxiv.org/abs/2005.06249>



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  ijirset@gmail.com



www.ijirset.com

Scan to save the contact details