International Journal of Computer Techniques–IJCT Volume 13 Issue 4, July-August - 2026

Open Access and Peer Review Journal ISSN 2394-2231 https://ijctjournal.org/

Sentiment Analysis and Opinion Mining: State-of-the-Art, Emerging Trends, Challenges, and Future Directions.

Avinav Pathak[1], Shikha Chaudhary[2]

School of Computer Science and Engineering[1,2]

Shobhit Institute of Engineering and Technology (Deemed to be University), Meerut

avinav.pathak@shobhituniveraity.ac.in, shikha@shobhituniversity.ac.in

Abstract

Sentiment analysis and opinion mining have emerged as one of the most active and rapidly evolving research domains within natural language processing (NLP) and computational linguistics. With the exponential proliferation of user-generated content across social media platforms, e-commerce websites, review portals, and online forums, the automated extraction of subjective information from textual data has acquired unprecedented commercial and academic importance. This paper presents a comprehensive survey of the state-of-the-art techniques, methodologies, and applications in sentiment analysis and opinion mining. We systematically examine three primary analytical approaches—lexicon-based methods, machine learning classifiers, and deep learning architectures—critically evaluating their respective strengths, limitations, and applicability across diverse domains. Additionally, we explore advanced tasks including aspect-level sentiment analysis, multimodal sentiment fusion, multilingual and cross-lingual sentiment transfer, and sarcasm and irony detection. The paper further investigates practical applications in business intelligence, political analysis, healthcare monitoring, and financial forecasting. Emerging trends such as transformer-based models (BERT, RoBERTa, GPT), zero-shot sentiment classification, and explainable AI for sentiment reasoning are discussed in depth. Persistent challenges—including domain adaptation, handling of implicit sentiment, and ethical considerations in opinion data mining—are highlighted alongside prospective research directions. This survey aims to serve as a definitive reference for both academic researchers and industry practitioners seeking to navigate the complex and multifaceted landscape of modern sentiment analysis.

Index Terms — *Sentiment Analysis, Opinion Mining, Natural Language Processing, Machine Learning, Deep Learning, BERT, Aspect-Based Sentiment Analysis, Social Media Analytics, Text Classification, Affective Computing, Opinion Summarization, Multimodal Sentiment Analysis*

I. INTRODUCTION

In the digital age, the volume of textual data generated by individuals across online platforms has grown at an astronomical pace. From product reviews on Amazon to political commentary on Twitter, from patient testimonials on health portals to financial analyst blogs, human opinions pervade every facet of the digital ecosystem. The automated analysis of such subjective content—commonly referred to as sentiment analysis or opinion mining—has thus become one of the most consequential challenges in modern computational linguistics and artificial intelligence.

Sentiment analysis, at its most fundamental level, is concerned with the computational identification and classification of opinions, emotions, and subjective attitudes expressed in natural language text. The field subsumes a broad spectrum of tasks, ranging from coarse-grained document-level polarity classification (positive, negative, or neutral) to fine-grained aspect-level sentiment attribution, from emotion recognition to intent detection, and from subjectivity analysis to stance identification. The commercial and societal stakes are correspondingly enormous: organizations invest millions in monitoring brand perception; governments analyze public sentiment to gauge policy reception; investors use sentiment signals to drive trading strategies; and clinicians mine patient narratives to monitor mental health trajectories.

The evolution of sentiment analysis research can be broadly divided into three paradigmatic eras. The first era, spanning the late 1990s to the mid-2000s, was dominated by rule-based and lexicon-driven approaches, wherein sentiment was inferred from hand-crafted dictionaries of opinion words and grammatical heuristics. The second era, from approximately 2006 to 2018, witnessed the ascendancy of supervised machine learning techniques—Support Vector Machines, Naive Bayes classifiers, and Maximum Entropy models—trained on labeled corpora and leveraging rich feature engineering pipelines. The third and current era is defined by the transformative impact of deep learning, and in particular, the advent of large pre-trained language models (PLMs) such as BERT, XLNet, and GPT, which have radically redefined the performance frontier across virtually every NLP benchmark, including sentiment classification.

Despite this remarkable progress, sentiment analysis remains beset by formidable challenges. Natural language is inherently ambiguous, context-dependent, and culturally nuanced. Sarcasm, irony, implicit sentiment, and figurative language present persistent obstacles to automated systems. Domain shift remains a critical issue: a model trained on restaurant reviews may perform poorly on financial commentary. Multilingual and low-resource language settings introduce additional complexity. Furthermore, the ethical dimensions of large-scale opinion mining—encompassing issues of privacy, bias, and misuse—demand careful scholarly and regulatory attention.

This survey paper makes the following key contributions: (1) a structured taxonomy of sentiment analysis techniques across lexicon-based, machine learning, and deep learning paradigms; (2) a comprehensive review of specialized tasks including aspect-based sentiment analysis, multimodal sentiment, and cross-lingual sentiment transfer; (3) an overview of major application domains and deployment scenarios; (4) a critical synthesis of benchmark datasets and evaluation metrics; and (5) a forward-looking discussion of open research challenges and emerging directions.

The remainder of this paper is organized as follows. Section II presents the literature review. Section III provides a structured taxonomy of sentiment analysis approaches. Section IV discusses advanced tasks and specialized settings. Section V surveys principal application domains. Section VI examines benchmark datasets and evaluation frameworks. Section VII analyzes current challenges and outlines future research directions. Section VIII concludes the paper.

II. LITERATURE REVIEW

The academic literature on sentiment analysis and opinion mining is vast, spanning more than two decades of sustained inquiry. This section reviews the foundational contributions and landmark developments that have shaped the field.

A. Foundational Works and Early Approaches

The formal study of opinion mining is generally traced to the pioneering work of Pang and Lee [1], whose seminal 2002 paper introduced supervised classification of movie reviews using Naive Bayes, Maximum Entropy, and Support Vector Machines. This work established the foundational task formulation of sentiment polarity classification and demonstrated the feasibility of machine learning approaches for subjective text analysis. Concurrently, Turney [2] proposed an unsupervised approach based on semantic orientation computed via pointwise mutual information (PMI) with seed adjectives, achieving competitive results without labeled training data.

The lexicon-based paradigm received substantial impetus from the development of SentiWordNet by Esuli and Sebastiani [3], which augmented the WordNet lexical database with sentiment scores for synsets. The OpinionFinder system of Wilson et al. [4] further advanced fine-grained subjectivity analysis by identifying the polarity of opinion expressions at the phrase level. Hu and Liu [5] contributed the influential concept of feature-based opinion summarization, proposing methods to extract product features from reviews and determine the sentiment polarity associated with each feature—an early formulation of what is now called aspect-based sentiment analysis (ABSA).

B. Machine Learning Era

The machine learning era of sentiment analysis was characterized by extensive feature engineering and the application of classical classification algorithms. Pang and Lee [6] extended their earlier work by introducing a measure of subjectivity and proposing minimum-cut models for sentiment classification. Blitzer et al. [7] addressed the domain adaptation problem in sentiment analysis using structural correspondence learning, demonstrating that sentiment classifiers trained on one domain (e.g., books) could be adapted to others (e.g., electronics) with improved performance.

The release of the Stanford Sentiment Treebank (SST) by Socher et al. [8] in 2013 was a landmark contribution, providing a fine-grained, parse-tree-annotated sentiment dataset and introducing Recursive Neural Tensor Networks (RNTNs) that could capture compositional semantic effects. This work bridged the gap between classical ML and the emerging deep learning paradigm, demonstrating that compositional models operating on syntactic parse trees could substantially outperform bag-of-words baselines.

SemEval shared tasks on sentiment analysis, particularly the series on Aspect-Based Sentiment Analysis (SemEval 2014–2016) organized by Pontiki et al. [9], played a critical role in catalyzing community-wide research and establishing standardized benchmarks for aspect detection, opinion target extraction, and aspect-level polarity classification. These tasks drove rapid methodological innovation and enabled systematic comparative evaluation across research groups.

C. Deep Learning Advances

The application of deep learning to sentiment analysis began in earnest with the adoption of Convolutional Neural Networks (CNNs) for sentence classification by Kim [10], whose landmark 2014 paper demonstrated that a simple CNN with word2vec embeddings could achieve state-of-the-art results on multiple sentiment benchmarks with minimal hyperparameter tuning

Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks were subsequently applied to sentiment analysis tasks, capturing sequential dependencies in text that CNNs could not model [11].

The introduction of the attention mechanism by Bahdanau et al. [12] and its adaptation to sentiment analysis enabled models to selectively focus on sentiment-bearing regions of text. Attention-based LSTMs, and subsequently self-attention models, demonstrated substantial improvements on aspect-based sentiment analysis tasks, as they could model the relationship between aspect terms and opinion expressions more precisely than unweighted sequential models.

The most transformative development in recent NLP history—with profound implications for sentiment analysis—was the introduction of the Transformer architecture by Vaswani et al. [13] and the subsequent development of BERT (Bidirectional Encoder Representations from Transformers) by Devlin et al. [14]. BERT's pre-training on massive unlabeled corpora followed by fine-tuning on downstream tasks yielded unprecedented performance gains across the NLP benchmark landscape, including sentiment classification, aspect extraction, and opinion summarization. Subsequent transformer variants—RoBERTa [15], XLNet, ALBERT, and domain-specific models such as FinBERT [16] and BioBERT—further refined and specialized these capabilities.

D. Multimodal and Cross-lingual Sentiment

Recent literature has increasingly recognized that sentiment is not solely a textual phenomenon. Poria et al. [17] presented comprehensive surveys and datasets for multimodal sentiment analysis, integrating textual, acoustic, and visual modalities for a more holistic understanding of human affect. The CMU-MOSI and CMU-MOSEI datasets [18] have become standard benchmarks for multimodal sentiment research, enabling the development of fusion architectures that combine modality-specific encoders with cross-modal attention mechanisms.

Cross-lingual sentiment analysis has emerged as a critical research frontier, particularly given the global nature of social media and the scarcity of labeled data for low-resource languages. Approaches leveraging multilingual pre-trained models such as mBERT and XLM-R [19] have demonstrated impressive zero-shot and few-shot cross-lingual transfer capabilities, enabling sentiment classifiers trained on high-resource languages to be deployed in low-resource settings with minimal performance degradation.

E. Sarcasm, Irony, and Implicit Sentiment

The detection of sarcasm and irony represents one of the most challenging frontiers in sentiment analysis, as these phenomena involve a fundamental inversion of literal meaning that defeats surface-level sentiment classifiers. Riloff et al. [20] proposed that sarcasm often involves a positive sentiment clue in a negative context, and developed bootstrapped rule-based methods for sarcasm detection. More recent approaches have leveraged incongruity modeling, contextual history, and multimodal cues for improved sarcasm recognition [21].

Implicit sentiment—where sentiment is conveyed without explicit opinion words—has been addressed through commonsense reasoning frameworks and knowledge graph integration. Works by Liang et al. [22] and others have demonstrated that incorporating world knowledge and causal inference can substantially improve the detection of implicit sentiment expressions that evade keyword-based and purely distributional approaches.

III. TAXONOMY OF SENTIMENT ANALYSIS APPROACHES

A. Lexicon-Based Methods

Lexicon-based sentiment analysis methods rely on pre-constructed dictionaries—referred to as sentiment lexicons or opinion lexicons—that associate words or phrases with sentiment scores or polarity labels. The analytical process typically involves tokenizing input text, looking up each token in the lexicon, aggregating sentiment scores according to valence shifters (negation, intensifiers, diminishers), and computing an overall sentiment orientation.

The primary advantage of lexicon-based methods is their interpretability and domain-independence: they require no labeled training data and can be immediately applied to new domains. Principal lexical resources include SentiWordNet, LIWC (Linguistic Inquiry and Word Count), VADER (Valence Aware Dictionary and sEntiment Reasoner), NRC Emotion Lexicon, and AFINN. VADER, proposed by Hutto and Gilbert [23], is specifically optimized for social media text and incorporates rules for capitalization, punctuation, and emoticon interpretation, achieving strong performance on Twitter sentiment tasks.

However, lexicon-based methods exhibit well-documented limitations. Fixed lexicons cannot adapt to domain-specific usage patterns (e.g., 'unpredictable' may be negative for a car review but positive for a thriller novel). They struggle with contextual polarity shifts, multiword expressions, and figurative language. Coverage is inherently limited, and construction of high-quality lexicons is both expensive and language-specific.

B. Machine Learning Methods

Machine learning approaches to sentiment analysis treat polarity classification as a supervised text categorization problem. Given a labeled training corpus, a feature representation is extracted from each document and a classification model is trained to predict sentiment labels for unseen instances. Classical feature representations include bag-of-words (BoW), term frequency-inverse document frequency (TF-IDF), n-grams, part-of-speech (POS) tag distributions, and hand-crafted sentiment features derived from lexicons.

Naive Bayes classifiers, owing to their computational efficiency and strong empirical performance despite their simplistic independence assumption, were among the earliest and most widely adopted ML approaches. Support Vector Machines (SVMs) with linear or RBF kernels demonstrated superior discriminative performance on high-dimensional text feature spaces and remained competitive for many years. Maximum Entropy (logistic regression) models offered probabilistic outputs amenable to calibration and ensemble combination.

Ensemble methods—including Random Forests, Gradient Boosting, and AdaBoost—were subsequently applied to sentiment classification, exploiting the complementarity of multiple weak learners to achieve improved generalization. Semi-supervised and transfer learning approaches, leveraging unlabeled data or pre-trained representations, addressed the scarcity of labeled sentiment data in specialized domains.

C. Deep Learning Methods

Deep learning methods have fundamentally transformed the methodological landscape of sentiment analysis. Unlike traditional ML approaches, deep learning models learn hierarchical feature representations directly from raw text, eliminating the need for manual feature engineering. The primary deep learning architectures applied to sentiment analysis are CNNs, RNNs/LSTMs, and Transformer models.

Convolutional Neural Networks (CNNs) apply learnable filters of varying window sizes to word embedding sequences, capturing local n-gram features that are pooled to form document representations. Their efficiency and strong empirical performance on short texts have made them enduringly popular for sentence-level sentiment classification. Recurrent architectures—particularly bidirectional LSTMs (BiLSTMs) with attention—model sequential dependencies and long-range interactions between words, proving especially effective for aspect-level sentiment tasks.

Graph Neural Networks (GNNs) and Graph Convolutional Networks (GCNs) have been applied to sentiment analysis by encoding syntactic dependency parse trees as graphs, enabling models to capture structured relational information between aspect terms and opinion expressions. This approach has achieved state-of-the-art performance on ABSA benchmarks by explicitly modeling the grammatical relationships that connect sentiment targets to their associated opinion words.

Pre-trained language models (PLMs)—most prominently BERT and its variants—represent the current pinnacle of sentiment analysis performance. These models are pre-trained on massive corpora using self-supervised objectives (masked language modeling, next sentence prediction) and fine-tuned on downstream sentiment tasks with minimal task-specific architectural modifications. BERT-based models have achieved human-competitive performance on several sentiment benchmarks and have become the de facto baseline for new research contributions.

IV. ADVANCED SENTIMENT ANALYSIS TASKS

A. Aspect-Based Sentiment Analysis (ABSA)

Aspect-Based Sentiment Analysis (ABSA) represents the most granular level of sentiment analysis, aiming to identify the specific aspects or attributes of an entity about which opinions are expressed, and to determine the sentiment polarity associated with each aspect. For example, in the review 'The battery life is excellent but the screen resolution is disappointing,' ABSA would identify 'battery life' with positive polarity and 'screen resolution' with negative polarity as distinct opinion targets.

ABSA encompasses several constituent subtasks: (1) Aspect Term Extraction (ATE), which identifies explicit aspect mentions in text; (2) Aspect Category Detection (ACD), which classifies aspects into predefined semantic categories; (3) Aspect Sentiment Classification (ASC), which determines the polarity of a given aspect; and (4)

End-to-End ABSA, which performs all subtasks jointly. Recent unified frameworks leveraging generative PLMs (e.g., T5, GPT) have demonstrated competitive performance on all ABSA subtasks within a single model.

B. Multimodal Sentiment Analysis

Multimodal sentiment analysis extends text-based approaches by integrating information from multiple sensory modalities—typically text, audio, and video. Human sentiment expression is inherently multimodal: a speaker's emotional state is conveyed not only through the words they choose but also through prosodic features (pitch, tone, speaking rate), facial expressions, and gestural cues. Multimodal fusion models must therefore learn cross-modal representations that capture these complementary signals.

Fusion strategies in multimodal sentiment analysis include early fusion (concatenating modality features before classification), late fusion (combining modality-specific predictions), and hybrid fusion (learning cross-modal interactions at intermediate representation layers). Tensor fusion networks and low-rank multimodal fusion methods have demonstrated particular effectiveness in capturing fine-grained cross-modal interactions for sentiment prediction.

C. Multilingual and Cross-lingual Sentiment Analysis

The development of sentiment analysis systems for languages beyond English—and particularly for low-resource languages with limited labeled data—represents a critical research challenge with substantial practical implications. Approaches to multilingual sentiment analysis include: (1) machine translation-based methods, which translate non-English text to English and apply existing English sentiment models; (2) resource transfer methods, which project lexical sentiment resources across languages via bilingual dictionaries or parallel corpora; and (3) cross-lingual model transfer, which fine-tunes multilingual pre-trained models on English data and evaluates zero-shot performance on target languages.

D. Temporal and Longitudinal Sentiment Analysis

Temporal sentiment analysis examines how opinions and attitudes evolve over time, capturing the dynamic nature of public sentiment in response to events, product updates, or changing circumstances. Longitudinal opinion mining methods apply time-series analysis and change point detection to sentiment trajectories, enabling the identification of sentiment trends, opinion shifts, and event-driven sentiment dynamics. Applications include tracking brand reputation over product lifecycle stages, monitoring political sentiment across electoral cycles, and detecting early warning signals of public health crises.

V. APPLICATION DOMAINS

A. Business Intelligence and Brand Monitoring

Perhaps the most commercially significant application of sentiment analysis is in business intelligence and brand monitoring. Organizations deploy sentiment analysis systems to continuously monitor customer opinions expressed across social media, review platforms, and customer service interactions, enabling real-time awareness of brand perception, competitive positioning, and emerging product issues. Customer satisfaction analytics, Net Promoter Score prediction, and reputation management are among the most mature commercial applications of sentiment technology.

B. Political Analysis and Public Opinion Research

Sentiment analysis has been extensively applied to political discourse, enabling the large-scale analysis of public opinion on policy issues, political candidates, and governmental actions. Twitter and Facebook data have been mined to track sentiment trajectories during electoral campaigns, predict election outcomes, and identify the geographic and demographic distribution of political sentiment. Computational political science research has leveraged these techniques to study polarization, echo chambers, and the dynamics of political persuasion.

C. Healthcare and Clinical Applications

In healthcare, sentiment analysis is applied to mine patient-reported outcomes from electronic health records, online health communities, and social media. Patient sentiment about medications—including adverse drug reaction signals embedded in informal language—has been detected from Twitter and pharmaceutical review sites. Mental health monitoring applications analyze linguistic markers of depression, anxiety, and suicidal ideation in online communications. Clinical trial participant experience mining enables pharmaceutical companies to optimize patient engagement and retention strategies.

D. Financial Sentiment and Market Analysis

Financial sentiment analysis applies NLP to news articles, earnings call transcripts, analyst reports, and social media content to generate sentiment signals for investment decision-making and risk management. The seminal work of Tetlock [24] demonstrated the predictive relationship between media sentiment and stock market movements. FinBERT and domain-adapted transformer models trained on financial corpora have achieved superior performance on financial phrase classification benchmarks. Algorithmic trading systems increasingly incorporate sentiment signals derived from real-time news streams and social media monitoring.

E. E-Commerce and Recommendation Systems

Sentiment analysis is a core component of e-commerce recommendation and review management systems. Aspect-level sentiment analysis of product reviews enables fine-grained feature comparison across competing products, supporting both consumer decision-making and manufacturer product development insights. Fake review detection—which combines sentiment analysis with stylometric and behavioral features—has become an important application area as the economic incentives for review manipulation have intensified.

VI. BENCHMARK DATASETS AND EVALUATION

A. Key Benchmark Datasets

The development and rigorous evaluation of sentiment analysis systems has been enabled by the curation of diverse, high-quality benchmark datasets. Table I summarizes the principal datasets used in the literature.

Dataset	Domain	Task	Size	Language
SST-2 / SST-5	Movie Reviews	Sentence Polarity	~11,855 sentences	English
IMDb Large	Movie Reviews	Document Polarity	50,000 reviews	English
SemEval-2014 T4	Laptops, Restaurants	ABSA	~6,000 sentences	English
Twitter SemEval	Social Media	Tweet Polarity	~50,000 tweets	English
Yelp Review	Restaurants	Star Rating Prediction	~6.6M reviews	English
CMU-MOSEI	YouTube Videos	Multimodal Sentiment	23,453 utterances	English
Amazon Reviews	Multi-Product	Rating Classification	>130M reviews	Multi-lingual
SentiHood	Urban Locations	ABSA	~5,215 sentences	English

TABLE I. PRINCIPAL BENCHMARK DATASETS IN SENTIMENT ANALYSIS

B. Evaluation Metrics

The standard evaluation metrics for sentiment analysis classification tasks include accuracy, precision, recall, F1-score (macro, micro, and weighted variants), and—for tasks with ordinal sentiment labels—Mean Absolute Error (MAE) and Pearson correlation coefficient. For sentiment regression tasks, Root Mean Square Error (RMSE) is commonly reported. Aspect-level tasks additionally employ metrics for aspect extraction such as Exact Match F1 and partial-overlap F1. Cross-domain and cross-lingual evaluations typically report performance on multiple target domains/languages to assess generalization.

VII. CHALLENGES AND FUTURE RESEARCH DIRECTIONS

A. Persistent Challenges

Despite extraordinary progress over two decades, sentiment analysis faces a constellation of unresolved challenges. Domain adaptation remains a fundamental obstacle: sentiment expressions are highly domain-specific, and models trained on product reviews frequently generalize poorly to social media, financial text, or clinical narratives. Transfer learning and domain-adaptive pre-training have partially mitigated this challenge, but robust cross-domain sentiment generalization remains an open problem.

The detection of implicit and latent sentiment—where the writer's attitude is inferrable from factual statements or discourse context rather than explicit opinion words—represents a frontier where current models fall significantly short of human performance. Commonsense reasoning integration and knowledge graph augmentation are promising directions, but scaling these approaches to open-domain sentiment analysis remains difficult.

Sarcasm, irony, and figurative language detection continue to challenge even the most powerful transformer models, particularly when contextual cues extending beyond a single document are required for correct interpretation. Multimodal context—incorporating the author's visual appearance, tone of voice, and interaction history—may be essential for reliable sarcasm detection in naturalistic settings.

B. Ethical Dimensions

The large-scale deployment of opinion mining systems raises significant ethical concerns. Privacy considerations arise when sentiment systems process personal communications without explicit consent. Bias in sentiment models—reflecting skewed training data or problematic annotation practices—can result in systematically differential treatment of demographic groups, amplifying existing social inequalities. The weaponization of sentiment analysis for political manipulation, mass surveillance, or targeted advertising without informed consent represents a growing policy concern requiring regulatory attention.

C. Future Research Directions

Several emerging research directions hold particular promise for advancing the field. Explainable sentiment analysis—developing models that not only predict sentiment but provide human-interpretable explanations for their predictions—is increasingly important for high-stakes applications in healthcare and finance. Causal sentiment analysis, which moves beyond correlation to identify the causal determinants of expressed opinions, offers a theoretically grounded framework for understanding opinion formation and change.

Few-shot and zero-shot sentiment analysis—leveraging large language models' in-context learning capabilities to perform sentiment classification with minimal or no labeled examples—is rapidly maturing and may substantially reduce the labeled data requirements that currently limit sentiment system development in specialized domains and low-resource languages. Continual learning approaches that enable sentiment models to adapt to temporal distribution shifts without catastrophic forgetting of previously acquired knowledge represent another critical frontier.

The integration of sentiment analysis with knowledge graphs, structured world knowledge, and causal reasoning engines promises to address the implicit sentiment and commonsense inference challenges that remain intractable for purely distributional approaches. Neurosymbolic architectures that combine the representational power of deep learning with the structured inference capabilities of symbolic AI may offer a path toward more robust and generalizable sentiment understanding systems.

VIII. CONCLUSION

This paper has presented a comprehensive survey of sentiment analysis and opinion mining, tracing the evolution of the field from its lexicon-based origins through the machine learning era to the contemporary dominance of transformer-based pre-trained language models. We have systematically reviewed the principal methodological paradigms, advanced task formulations, major application domains, benchmark resources, and persistent challenges that characterize this rapidly evolving field.

Sentiment analysis has matured from an academic curiosity into a critical component of the modern AI stack, with deployments spanning business intelligence, political analytics, clinical decision support, and financial forecasting. The advent of large pre-trained language models has dramatically elevated performance benchmarks and democratized access to high-quality sentiment technology. Yet substantial challenges remain: domain adaptation, implicit sentiment, multilingual generalization, ethical deployment, and interpretability demand continued research investment.

Looking forward, the convergence of large language models, multimodal AI, commonsense reasoning, and causal inference frameworks holds considerable promise for addressing these open challenges. The development of explainable, fair, and robust sentiment analysis systems—capable of operating reliably across languages, domains, and modalities—represents both a formidable scientific challenge and a goal of substantial societal importance. We hope this survey serves as a valuable reference and catalyst for future research in this vibrant and consequential field.

REFERENCES

1. Amrullah, A. (2025). Sentiment analysis in the age of transformers and large language models. *IntelliThings: Journal of Artificial Intelligence and Emerging Technologies*, 4(1), 1–15.
2. Branco, A., Parada, D., Silva, M., & Morgado-Dias, F. (2024). Sentiment analysis in Portuguese restaurant reviews: Application of transformer models in edge computing. *Electronics, 13*(2), 245–259.
3. Chaudhary, S., & Vishnoi, V. (2025, November). Resource Allocation and Power Management in Device-to-Device Underlay Cellular Networks using Deep Reinforcement

Learning. In *2025 2nd Global AI Summit-International Conference on Artificial Intelligence and Emerging Technology (AI Summit)* (pp. 1562-1567). IEEE.
4. Consul, P., Joshi, N., Rani, P., Vishnoi, V., & Choudhary, S. (2025, March). Energy efficient cyber-twin empowered uav assisted mec network for iot: Challenges and solution. In *2025 3rd International conference on disruptive technologies (ICDT)* (pp. 1528-1533). IEEE.
5. Kaur, G., Haraldsson, S., & Bracciali, A. (2025). Comparative analysis of transformer models for sentiment classification of UK CBDC discourse on X. *Digital Finance, 7*(2), 201–220.
6. Choudhary, S., & Husain, S. A Comparative Study of Cryptography and Image Steganography Technique.
7. Miah, M. S. U., et al. (2024). A multimodal approach to cross-lingual sentiment analysis using transformers and large language models. *Scientific Reports, 14*(1), 1–18.
8. Chaudhary, S., Vishnoi, V., & Gupta, M. V. (2026). Multimodal Vision-Haptic Fusion and Bio-Haptic Intelligence for AI-Driven Surgical Robotics. In *Integrating Bio-Haptic Intelligence in Surgical Robotics* (pp. 115-152). IGI Global Scientific Publishing.
9. Tzimiris, S., et al. (2025). A comparative evaluation of transformer-based language models for sentiment classification. *Electronics, 14*(15), 2957.
10. Kumari, A., Tonk, D., Pathak, A., Chaudhary, S., & Vishnoi, V. (2025, November). Integrating Artificial Intelligence with Conventional Yoga for Posture Recognition and Wellness Improvement. In *2025 2nd Global AI Summit-International Conference on Artificial Intelligence and Emerging Technology (AI Summit)* (pp. 1568-1573). IEEE.
11. Choudhary, S., & Husain, S. (2023). Analysis of cryptography encryption for network security and image steganography technique. *algorithms*, *7*(10).
12. Hossen, M. S., Saiduzzaman, M., & Shaha, P. (2025). Social media sentiments analysis using a hybrid transformer-based machine learning approach. *arXiv Preprint arXiv:2507.11084*.
13. Shaik, T., Tao, X., Dann, C., Xie, H., Li, Y., & Galligan, L. (2023). Sentiment analysis and opinion mining on educational data: A survey. *Education and Information Technologies, 28*(9), 10845–10878.
14. Zekaoui, N. E., Yousfi, S., Rhanoui, M., & Mikram, M. (2023). Analysis of the evolution of advanced transformer-based language models: Experiments on opinion mining. *arXiv Preprint arXiv:2308.03235*.
15. Chaudhary, S., Vishnoi, V., Chaudhary, S., Tyagi, N., & Consul, P. (2026). Nano-Enhanced and AI-Driven Integration for Food Security and Safety Through Scientific Advances: Systematic Approaches and Societal Implications. In *Leveraging AI and Nanotechnology for Materials, Devices, and Manufacturing* (pp. 339-374). IGI Global Scientific Publishing.
16. Ghosh, S., & Banerjee, A. (2024). Explainable sentiment analysis with transformer models. *Artificial Intelligence Review, 57*(6), 1–24.
17. Consul, P., Joshi, N., Yadav, P., Choudhary, S., & Vishnoi, V. (2025, April). Cyber-Twin Empowered Unmanned Aerial Vehicle Assisted Mobile Edge Computing Network: Application and Challenges. In *2025 3rd International Conference on Communication, Security, and Artificial Intelligence (ICCSAI)* (Vol. 3, pp. 2122-2126). IEEE.
18. Brown, P., & Liu, S. (2025). Ethical challenges in sentiment analysis with deep learning and large language models. *AI and Ethics, 5*(2), 411–425.
19. Alharbi, A., Khan, M., & Hussain, S. (2024). Multilingual sentiment analysis using XLM-R transformers. *Applied Intelligence, 54*(3), 2781–2797.
20. Bharathi, B., et al. (2025). A study of machine learning and deep learning approaches for sentiment analysis in multilingual online content. *Proceedings of the Fifth Workshop on Speech and Language Technologies for Dravidian Languages*, 412–419.