ISSN: 2582-7219

# International Journal of Multidisciplinary Research in Science, Engineering and Technology

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

Impact Factor: 9.864

Volume 9, Issue 8, August 2026

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 9.864| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

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

# From Lexicons to Large Language Models: A Comprehensive Survey of Modern Sentiment Analysis Ecosystems

**Dr. Shrinath Pai**

Principal, Shree Guru Sudhindra College, Bhatkal, Karnataka, India

**ABSTRACT:** Sentiment Analysis (SA) and Opinion Mining have evolved rapidly from rigid rule-based lexicons to highly contextualized Large Language Models (LLMs). This survey presents a comprehensive review of the sentiment analysis ecosystem, tracing its trajectory across five distinct technological epochs: Lexicon-Based approaches, Traditional Machine Learning, Deep Learning Architectures, Pre-trained Transformer models, and Generative Large Language Models. We evaluate the core methodologies, architectural shifts, and benchmark datasets defining each phase. Furthermore, we examine critical sub-fields including Aspect-Based Sentiment Analysis (ABSA), Multimodal Sentiment Analysis (MSA), and Multilingual/Cross-lingual domain transfers. The review provides a synthesis of current research challenges, such as handling sarcasm, domain adaptation, bias, interpretability, and computational overhead, while outlining strategic directions for future research in low-resource adaptation and zero-shot reasoning.

**KEYWORDS:** Sentiment Analysis, Natural Language Processing, Large Language Models, Deep Learning, Aspect-Based Sentiment Analysis.

## I. INTRODUCTION

### Background

Sentiment Analysis (SA)—the computational study of opinions, sentiments, emotions, and attitudes expressed in written text—has emerged as one of the most active research domains within Natural Language Processing (NLP). With the exponential growth of user-generated content on social media, e-commerce platforms, blogs, and interactive forums, the capacity to automatically extract and structure human affect has become vital for business intelligence, political forecasting, public health monitoring, and social science research. [1]

Historically, sentiment analysis relied heavily on hand-crafted lexical resources (e.g., SentiWordNet, VADER) where predefined polarity scores were aggregated over tokenized text. While computationally efficient, these early systems struggled with linguistic nuances, contextual shift, and negation. The emergence of statistical machine learning algorithms in the early 2000s introduced data-driven sentiment classification using features like n-grams and term frequency-inverse document frequency (TF-IDF). [2]

The past decade has witnessed two profound architectural shifts: first, the deep learning revolution characterized by Recurrent Neural Networks (RNNs) and Convolutional Neural Networks (CNNs) capable of learning dense representation vectors; second, the transformer revolution powered by self-attention mechanisms, leading to models such as BERT, RoBERTa, and eventually generative Large Language Models (LLMs) such as GPT-4, LLaMA, and PaLM. [3]

© 2026 IJMRSET | Volume 9, Issue 8, August 2026| DOI:10.15680/IJMRSET.2026.0908022

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 9.864| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

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

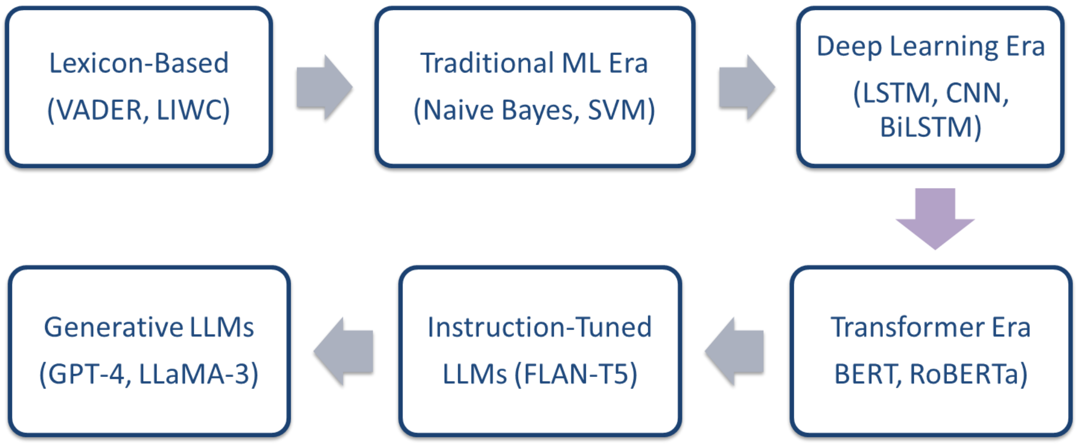


**Statement of the Problem**

Despite significant advancements, modern sentiment analysis ecosystems present complex trade-offs. While pre-trained transformers and generative LLMs demonstrate remarkable zero-shot and few-shot classification capabilities, they introduce major challenges regarding computational resource requirements, opacity ("black-box" decision making), susceptibility to prompt variations, and algorithmic bias. Conversely, light-weight traditional systems remain efficient but fail in complex, multi-lingual, or highly domain-specific context settings. Moreover, fine-grained tasks such as Aspect-Based Sentiment Analysis (ABSA) and Multimodal Sentiment Analysis (MSA) require nuanced representations that basic classification pipelines struggle to provide. [4]

**Purpose of the Review**

This survey aims to provide a structured evaluation of the paradigm shifts in sentiment analysis. Specifically, this review:

1. Systematically traces the technological trajectory from lexical dictionaries to generative LLMs.
2. Evaluates the methodological paradigms underlying sentence-level, document-level, and aspect-level SA.
3. Examines specialized sub-domains including Multimodal SA, Cross-lingual SA, and Emotion Recognition.
4. Critically analyzes open challenges, current performance benchmarks, and future directions for computational sentiment research. [5]

## II. METHODOLOGY

**Search Strategy**

A systematic search was conducted across major academic databases: IEEE Xplore, ACM Digital Library, PubMed, Google Scholar, Scopus, and arXiv repositories. Queries were constructed using boolean operators combining primary and secondary keywords: [6]

- ("Sentiment Analysis" OR "Opinion Mining") AND ("Lexicon" OR "Machine Learning" OR "Deep Learning" OR "BERT" OR "Large Language Model" OR "LLM")
- ("Aspect-Based Sentiment Analysis" OR "ABSA") AND ("Transformers" OR "Instruction Tuning")
- ("Multimodal Sentiment Analysis") AND ("Fusion" OR "Cross-Attention")

**Inclusion and Exclusion Criteria**

To maintain rigor and relevance, explicit criteria were established for literature filtering: [7]

**Inclusion Criteria:**

1. Peer-reviewed journal articles, top-tier conference papers (ACL, EMNLP, NAACL, AAAI, KDD, IEEE TPAMI), and seminal arXiv preprints.
2. Studies introducing novel computational architectures, benchmark datasets, or empirical evaluations in Sentiment Analysis.
3. Literature published between 2000 and 2026, with an emphasis on foundational deep learning and transformer papers post-2017. [8]

© 2026 IJMRSET | Volume 9, Issue 8, August 2026| DOI:10.15680/IJMRSET.2026.0908022

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 9.864| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

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

**Exclusion Criteria:**

1. Studies lacking rigorous baseline evaluations or validation on standard benchmarks.
2. Non-English papers without peer-reviewed English translations.
3. Short abstract-only reports, non-peer-reviewed blog posts, or redundant domain-specific applications lacking methodological novelty. [9]

**Data Extraction**

From the qualified set of papers, key operational data were extracted: model classification family, core architectural components, target dataset(s), task granularity (document, sentence, aspect), performance metrics (Accuracy, F1-Score), reported limitations, and computational hardware demands. [10]

**Quality Assessment**

Each selected paper was assessed for methodological robustness using a standard checklist assessing:

- Clarity of research objectives and architectural definitions.
- Reproducibility of experiments (code availability, hyperparameter specification).
- Significance of empirical baseline comparisons. [11]

## III. REVIEW OF LITERATURE

**Organizational Structure**

The literature is categorized into five primary architectural paradigms: (1) Lexicon-Based Approaches, (2) Machine Learning Models, (3) Deep Neural Architectures, (4) Pre-trained Transformer Networks, and (5) Generative Large Language Models. Additionally, we analyze specialized domains: Aspect-Based Sentiment Analysis (ABSA) and Multimodal Sentiment Analysis (MSA). [12]

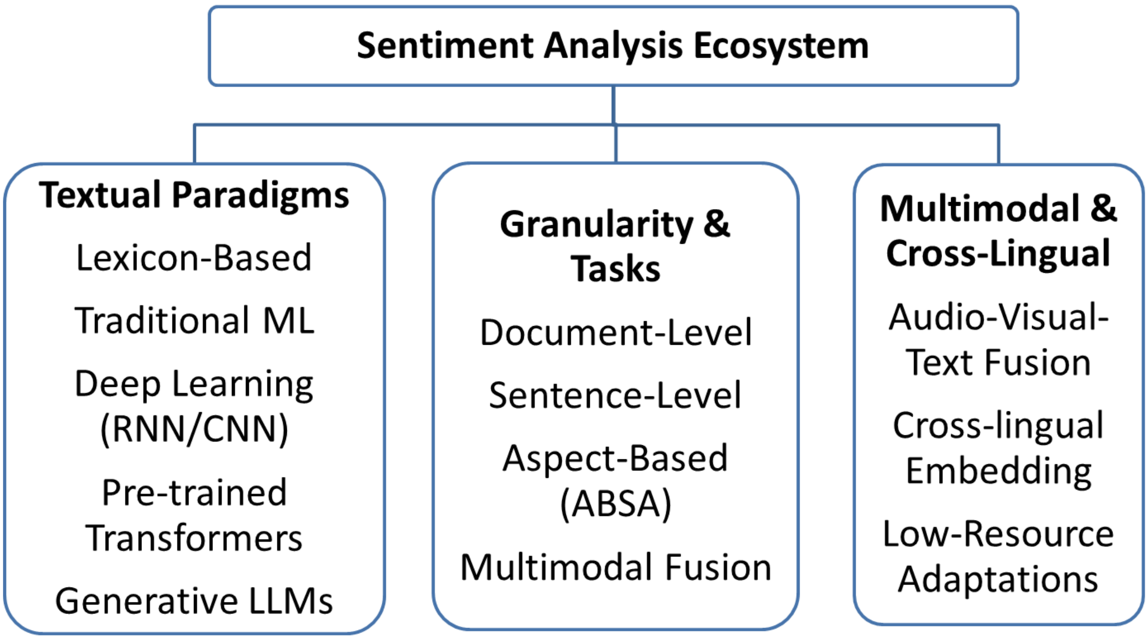


**Thematic Subsections**

**1. Lexicon-Based and Rule-Based Paradigms**

Early sentiment analysis depended on curated dictionaries containing word polarity scores.

- **Manual and Corpus-Based Lexicons:** Resources such as LIWC (Linguistic Inquiry and Word Count) and General Inquirer established initial baseline representations.

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 9.864| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

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

- **Dictionary-Based Computational Expansion:** SentiWordNet leveraged WordNet relationships to propagate sentiment scores dynamically.
- **Heuristic Frameworks:** VADER (Valence Aware Dictionary and sEntiment Reasoner) incorporated rule-based heuristics to handle punctuation, capitalization, degree modifiers, and adversarial conjunctions (e.g., "but"), setting a high benchmark for unannotated social media text processing. [13]

*Limitation:* Lexical methods fail when encountering out-of-vocabulary (OOV) slang, sarcasm, or context-dependent polarity shifts (e.g., "cold beer" vs. "cold soup"). [13]

**2. Traditional Machine Learning Paradigms**

The shift toward supervised statistical models addressed lexicon limitations by learning from annotated training corpora. [14]

- **Feature Engineering:** Bag-of-Words (BoW), n-grams, and TF-IDF vectors formed standard text inputs, often paired with hand-crafted syntactic features (POS tags).
- **Algorithms:** Naive Bayes (NB), Support Vector Machines (SVM), and Maximum Entropy classifiers proved highly effective. Pang et al. (2002) established that SVMs trained on unigrams outpaced lexical methods in movie review classification.
- **Ensemble Methods:** Random Forests and Gradient Boosting Decision Trees (GBDTs) were introduced to handle high-dimensional feature spaces. [14]

*Limitation:* Traditional ML approaches treat text as an unordered collection of token vectors, ignoring long-range context, structural syntax, and semantic similarity. [15]

**3. Deep Learning Architectures**

The introduction of word embedding models—Word2Vec (Mikolov et al., 2013) and GloVe (Pennington et al., 2014)—revolutionized text representation by mapping words to dense vector spaces preserving semantic relationships. [15]

- **Convolutional Neural Networks (CNNs):** Kim (2014) demonstrated that 1D-CNNs with varied filter sizes could capture local dynamic n-gram features efficiently.
- **Recurrent Neural Networks (RNNs/LSTMs):** Long Short-Term Memory (LSTM) networks and Gated Recurrent Units (GRUs) solved the vanishing gradient problem, allowing sequential contextual modeling. Bidirectional LSTMs (BiLSTMs) captured both preceding and succeeding context.
- **Attention Mechanisms:** The integration of attention modules (Bahdanau et al., 2014) allowed deep models to assign differential weights to words, improving focus on sentiment-bearing tokens. [16]

**4. Pre-trained Transformer Networks**

The introduction of the Transformer architecture (Vaswani et al., 2017) based entirely on self-attention eliminated sequential processing dependencies, paving the way for pre-trained language models.

- **Bi-directional Encoders (BERT):** Devlin et al. (2018) introduced BERT, utilizing Masked Language Modeling (MLM) to construct deeply contextualized representations. BERT achieved state-of-the-art results across text classification datasets (SST-2, IMDB).
- **Architectural Refinements:**
  - *RoBERTa (Liu et al., 2019):* Optimized BERT training by removing the Next Sentence Prediction (NSP) loss and training on larger mini-batches.
  - *ALBERT (Lan et al., 2019):* Reduced parameter footprint via factorized embedding parameterization and cross-layer parameter sharing.
  - *XLNet (Yang et al., 2019):* Integrated autoregressive permutation training to overcome MLM pre-training discrepancies.

**5. Generative Large Language Models (LLMs)**

The transition from task-specific fine-tuning (encoder models) to task-agnostic prompt engineering (decoder models) marks the current frontier of sentiment research.

- **Zero-Shot and Few-Shot Prompting:** Models such as GPT-3.5, GPT-4, LLaMA-3, and PaLM 2 process sentiment tasks directly via natural language prompts without parameter updating.

© 2026 IJMRSET | Volume 9, Issue 8, August 2026| DOI:10.15680/IJMRSET.2026.0908022

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 9.864| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

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

- **Instruction-Tuning and Alignment:** Models fine-tuned using Reinforcement Learning from Human Feedback (RLHF) and Direct Preference Optimization (DPO) demonstrate an ability to execute complex reasoning, trace multi-step sentiment logic, and generate structured analytical outputs.
- **Parameter-Efficient Fine-Tuning (PEFT):** Methods like LoRA (Low-Rank Adaptation) and QLoRA allow massive generative models to be fine-tuned for specialized sentiment applications on edge hardware.

### 6. Aspect-Based Sentiment Analysis (ABSA)

Document- or sentence-level sentiment scores can obscure fine-grained nuances when multiple entities or attributes are mentioned within a single text instance. ABSA decomposes sentiment analysis into sub-tasks:

1. Aspect Term Extraction (ATE)
2. Aspect Category Detection (ACD)
3. Aspect Sentiment Classification (ASC)

Graph Convolutional Networks (GCNs) operating over dependency trees (e.g., Zhang et al., 2019) aligned syntactic relations with aspect targets. Modern approaches recast ABSA as a generative sequence-to-sequence problem using models like T5 or fine-tuned LLMs, extracting aspect-sentiment pairs in structured formats (e.g., JSON). [17]

Sentence: "The food was incredible, but the service was terribly slow."
Aspect 1: "food" Polarity: POSITIVE
Aspect 2: "service" Polarity: NEGATIVE

### 7. Multimodal Sentiment Analysis (MSA)

The growth of modern content platforms necessitates processing diverse modalities simultaneously: Text, Audio, and Video.

- **Fusion Strategies:** Early fusion (concatenating feature vectors at input), Late fusion (aggregating individual model decision probabilities), and Hybrid/Cross-modal Attention fusion.
- **Architectures:** Models like the Multimodal Transformer (MulT) use cross-modal attention to align asynchronous signals across visual frames, audio pitch/tone, and textual transcripts. [18]

**Summary of Major Studies and Benchmarks**

The following table synthesizes the landmark contributions across paradigms, illustrating key methodologies, datasets, and benchmark performances.

| Paradigm | Landmark Study | Primary Methodology | Key Datasets | Performance / Metrics | Primary Limitations |
|---|---|---|---|---|---|
| **Lexicon** | Hutto & Gilbert (2014) | Rule-based lexical valence scoring (VADER) | SST-2, Twitter Corpus | ~79.0% Accuracy (Twitter) | Ignores complex context & irony. |
| **Traditional ML** | Pang et al. (2002) | Bag-of-Words + SVM / Naive Bayes | Movie Review Corpus | 82.9% Accuracy (SVM) | High dimensionality; loses syntax. |
| **Deep Learning** | Kim (2014) | 1D CNN over static/non-static Word2Vec | SST-2, TREC, IMDB | 88.1% Accuracy (SST-2) | Limited long-range dependency modeling. |
| **Deep Learning** | Tai et al. (2015) | Tree-Structured LSTM | SST-5 (Fine-grained) | 51.0% Accuracy (SST-5) | Requires explicit dependency parsing. |
| **Transformer** | Devlin et al. (2018) | Masked Language Modeling + | SST-2, CoLA, | 93.5% Accuracy | High computational |

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 9.864| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

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

| Paradigm | Landmark Study | Primary Methodology | Key Datasets | Performance / Metrics | Primary Limitations |
|---|---|---|---|---|---|
| | | Multi-head Attention (BERT) | GLUE | (SST-2) | costs during fine-tuning. |
| **Transformer** | Liu et al. (2019) | Optimized Pre-training strategy (RoBERTa) | SST-2, IMDB | 96.4% Accuracy (SST-2) | Opaque decision-making mechanisms. |
| **Generative LLM** | Wei et al. (2022) | Chain-of-Thought (CoT) Prompting | StrategyQA, Custom SA | ~90%+ Zero-shot reasoning | High latency, potential hallucinations. |
| **ABSA** | Zhang et al. (2019) | Syntactic Dependency Tree + GCN | Laptop14, Rest14 | 81.6% F1-Score (Rest14) | Sensitive to parsing error propagation. |
| **Multimodal** | Tsai et al. (2019) | Multimodal Transformer (MulT) | CMU-MOSI, CMU-MOSEI | 82.5% Accuracy (MOSI) | Alignment errors across asynchronous modes. |

**Critical Analysis**

A comparative evaluation of these paradigms yields distinct trade-offs across five operational dimensions:

1. **Contextual Scope:** Lexical models analyze tokens in isolation. Traditional ML considers limited sliding n-gram windows. Deep Learning architectures extend context across sequences, while Transformers and LLMs process thousands of tokens simultaneously, capturing structural and discourse-level dynamics.
2. **Data Efficiency:** Lexicon-based frameworks require zero training data. Deep learning and transformer models demand extensive annotated data or compute-intensive pre-training. Generative LLMs bridge this gap via zero-shot/few-shot in-context learning, though optimal performance often requires targeted instruction-tuning.
3. **Domain Adaptability:** Rule-based and basic ML models deteriorate rapidly when applied outside their target training domain due to vocabulary shifts. Pre-trained language models transfer rich language representations across diverse domains with minimal fine-tuning.
4. **Computational Complexity:** Operational demands have grown exponentially. Lexicons run efficiently on basic CPU hardware. Deep neural networks require consumer GPUs. Conversely, serving generative LLMs requires distributed GPU clusters, raising deployment costs and energy footprints.
5. **Interpretability:** Lexicons offer full transparency via deterministic rule paths. Linear ML models can be audited using feature weights. Deep neural networks, transformers, and LLMs operate as complex mathematical "black boxes," requiring auxiliary explainability frameworks (e.g., LIME, SHAP, integrated gradients, or explicit reasoning chains). [19]

## IV. DISCUSSION

**Synthesis of Findings**

The evolution of sentiment analysis highlights a fundamental shift from human-defined rules to autonomous, representations-based learning. While accuracy across standard benchmarks (such as SST-2 and IMDB) has reached saturated levels (>96%), modern sentiment analysis focus has shifted toward complex linguistic conditions: multi-party conversations, cross-cultural context, multimodal streams, and high-granularity ABSA tasks.

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 9.864| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

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

**Identification of Trends**

- **Convergence to Generative Paradigms:** The traditional boundary between extraction, classification, and text generation has blurred. SA tasks are increasingly formulated as text-to-text transformation problems.
- **Parameter-Efficient Adaptation:** Methods such as Low-Rank Adaptation (LoRA), Prefix Tuning, and Adapter Modules have become essential for deploying large models within constrained environments.
- **Explainable Sentiment Analysis (XSA):** Shifting focus from raw label output (Positive/Negative) toward generating clear rationale explanations for classified sentiment scores.

**Theoretical Framework**

We conceptualize the modern Sentiment Analysis ecosystem through a 4-Layer Architectural Stack model:
Layer 4: Reasoning & Decoding Layer (Zero-shot / CoT / PEFT)
Layer 3: Contextual & Structural Layer (Self-Attention / GCN)
Layer 2: Multimodal & Cross-Lingual Alignment Layer
Layer 1: Dense Representation Layer (Token / Positional / Modal)

**Implications**

- **Industry & Enterprise:** Organizations can leverage hybrid models: deployment of lightweight, parameter-efficient distilled transformers for high-volume streaming data (e.g., customer support tickets), coupled with generative LLM auditing for low-volume, high-value strategic analytics.
- **Academic Research:** Focus must move away from benchmark optimization on clean, balanced datasets toward evaluating robustness against real-world domain noise, slang, sarcasms, and low-resource languages. [20]

## V. CONCLUSION

**Summary**

Sentiment Analysis has transitioned from simplistic lexical lookup techniques to sophisticated generative architectures capable of understanding nuanced human emotion and context. While pre-trained transformers and modern Large Language Models have pushed performance metrics to unprecedented heights, key operational challenges remain regarding computational efficiency, interpretability, domain adaptation, and multimodal alignment. [21]

**Limitations**

This review is bounded by its structural scope:

1. It focuses predominantly on text-centric and primary multimodal formats (audio/video), excluding sensor-based sentiment/affective computing.
2. The rapid pace of LLM development means empirical evaluations of proprietary commercial models reflect specific snapshot iterations. [22]

## VI. FUTURE RESEARCH DIRECTIONS

1. **Low-Resource and Cross-Lingual Transfer:** Developing efficient cross-lingual representations for low-resource languages without relying on costly translation pipelines.
2. **Robust Handling of Pragmatic Nuances:** Enhancing model detection of irony, sarcasm, hyperbole, and implicit sentiment through structured knowledge-graph integration.
3. **Green NLP & Model Compression:** Refining quantization, pruning, and knowledge distillation methods to bring state-of-the-art sentiment analysis to edge devices without degrading accuracy. [21]
4. **Unified Multimodal-Reasoning Frameworks:** Architecting natively multimodal LLMs capable of real-time joint processing of speech intonation, facial micro-expressions, and textual context.

## REFERENCES

[1] Bahdanau, D., Cho, K., & Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.

[2] Cambria, E., Schuller, B., Xia, Y., & Havasi, C. (2017). New avenues in multimodal sentiment analysis. *IEEE Intelligent Systems*, 32(5), 15-21.

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 9.864| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

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

[3] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
[4] Hu, M., & Liu, B. (2004). Mining and summarizing customer reviews. *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, 168-177.
[5] Hutto, C., & Gilbert, E. (2014). VADER: A parsimonious rule-based model for sentiment analysis of social media text. *Proceedings of the International AAAI Conference on Web and Social Media*, 8(1), 216-225.
[6] Kim, Y. (2014). Convolutional neural networks for sentence classification. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1746-1751.
[7] Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., & Soricut, R. (2019). ALBERT: A lite BERT for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
[8] Liu, B. (2012). *Sentiment analysis and opinion mining*. Morgan & Claypool Publishers.
[9] Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., & Stoyanov, V. (2019). RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*.
[10] Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
[11] Pang, B., Lee, L., & Vaithyanathan, S. (2002). Thumbs up? Sentiment classification using machine learning techniques. *Proceedings of the ACL Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 79-86.
[12] Pennington, J., Socher, R., & Manning, C. D. (2014). GloVe: Global vectors for word representation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1532-1543.
[13] Pontiki, M., Galanis, D., Papageorgiou, H., Androutsopoulos, I., Manandhar, S., AL-Smadi, M., ... & Ziakis, Z. (2016). SemEval-2016 task 5: Aspect based sentiment analysis. *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, 19-30.
[14] 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*, 1631-1642.
[15] Tai, K. S., Socher, R., & Manning, C. D. (2015). Improved semantic representations from tree-structured long short-term memory networks. *arXiv preprint arXiv:1503.00075*.
[16] Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M. A., Lacroix, T., ... & Lample, G. (2023). LLaMA: Open and efficient base language models. *arXiv preprint arXiv:2302.13971*.
[17] Tsai, Y. H. H., Bai, S., Liang, P. P., Kolter, J. Z., Morency, L. P., & Salakhutdinov, R. (2019). Multimodal transformer for unaligned multimodal language sequences. *Proceedings of the ACL*, 6558-6569.
[18] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998-6008.
[19] Wei, J., Wang, X., Schuurmans, D., Bosma, M., Fei, F., Xia, F., Chi, E., Le, Q. V., & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35, 24824-24837.
[20] Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R. R., & Le, Q. V. (2019). XLNet: Generalized autoregressive pretraining for language understanding. *Advances in Neural Information Processing Systems*, 32, 5753-5763.
[21] Zhang, C., Li, Q., & Song, D. (2019). Aspect-based sentiment classification with aspect-specific graph convolutional networks. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 4568-4578.

INNO SPACE
SJIF Scientific Journal Impact Factor

INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA

निस्केयर
NISCAIR


# INTERNATIONAL JOURNAL OF MULTIDISCIPLINARY RESEARCH

## IN SCIENCE, ENGINEERING AND TECHNOLOGY



| Mobile No: +91-6381907438 | Whatsapp: +91-6381907438 | ijmrset@gmail.com |



www.ijmrset.com