

Review

From Lexicons to Transformers: An AI View of Sentiment Analysis

Maikel Leon 

Department of Business Technology, Miami Herbert Business School, University of Miami, Coral Gables, FL 33146, USA; mleon@miami.edu

Received: 2 July 2025; **Revised:** 13 August 2025; **Accepted:** 17 August 2025; **Published:** 30 August 2025

Abstract: Understanding public opinion at scale is both a scientific challenge and a practical necessity in the digital era, as the proliferation of online communication platforms has created unprecedented opportunities to monitor attitudes in near real time. Early work in subjectivity detection and semantic orientation laid the methodological foundations for automated sentiment extraction, focusing on distinguishing objective from subjective content and determining polarity. Contemporary applications, however, face far more complex requirements, demanding systems capable of processing massive, noisy, and dynamic data streams while integrating multimodal signals from text, images, audio, and video. This paper presents a historical review of sentiment analysis and opinion monitoring through the lens of artificial intelligence, tracing developments from the early 1990s to the present and classifying approaches from lexicon-based heuristics to classical machine learning, deep neural architectures, transfer learning, and multimodal fusion, with an emphasis on both technical and conceptual advances. Extensive tables summarize algorithms, datasets, and case studies across various domains, including politics, finance, and entertainment, highlighting practical lessons and performance trends. The review also addresses pressing ethical concerns, including bias, fairness, and transparency, and considers the implications of rapidly evolving AI capabilities. We conclude by outlining future directions that emphasize adaptability, context awareness, and the seamless integration of emerging technologies into scalable and reliable opinion analysis systems.

Keywords: Sentiment Analysis; Public Opinion Monitoring; Lexicon-Based Techniques; Deep Learning; Multimodal Sentiment Integration

1. Introduction

Public opinion shapes elections, markets, and policymaking. Traditionally, surveys and focus groups have served as the primary tools for gauging sentiment; however, limited sample sizes and temporal sparsity have constrained the effectiveness of these methods. The rise of blogs, microblogs, and online reviews in the late 1990s and early 2000s created an unprecedented reservoir of opinionated data. Researchers quickly recognized the need for automated techniques to transform raw text into actionable insights. The terms sentiment analysis and opinion mining were formalized in seminal papers by Nasukawa and Yi [1] and Dave et al. [2], though earlier work on lexical semantics, subjectivity, and metaphor detection had already laid critical foundations. A 2002 study by Turney applied pointwise mutual information (PMI) to classify reviews without labeled data [3], while Pang et al. [4] showed that support-vector machines outperform Naive Bayes and maximum-entropy classifiers for movie reviews. Since then, the field has exploded, encompassing lexicon construction [5], supervised and semi-supervised learning [6,7], pre-trained neural models [8,9], multimodal fusion [10,11]. Real-time opinion monitoring systems [12–15]. The goal of this survey is to synthesize this trajectory and provide researchers and practitioners with a comprehensive reference.

The evolution of sentiment analysis has spanned several decades, as illustrated in **Figure 1**. Early work in the 1990s focused on subjectivity detection. Hatzivassiloglou and McKeown [16] investigated polarity prediction based on linguistic constraints. Building on this, Turney [3] and Pang & Lee [4] advanced the field with unsupervised and supervised sentiment classification techniques. Nasukawa and Yi [1] subsequently formalized the notion of sentiment analysis. The introduction of neural embeddings like Word2Vec (2013) and GloVe (2014) [17, 18] revolutionized feature representation, while Vaswani et al.'s Transformer [19] enabled scalable attention mechanisms that paved the way for BERT (2018) [8], a context-aware language model. Most recently, ChatGPT (2022) exemplified the integration of generative AI with sentiment-aware capabilities, marking a significant shift toward interactive, contextually aware opinion mining.

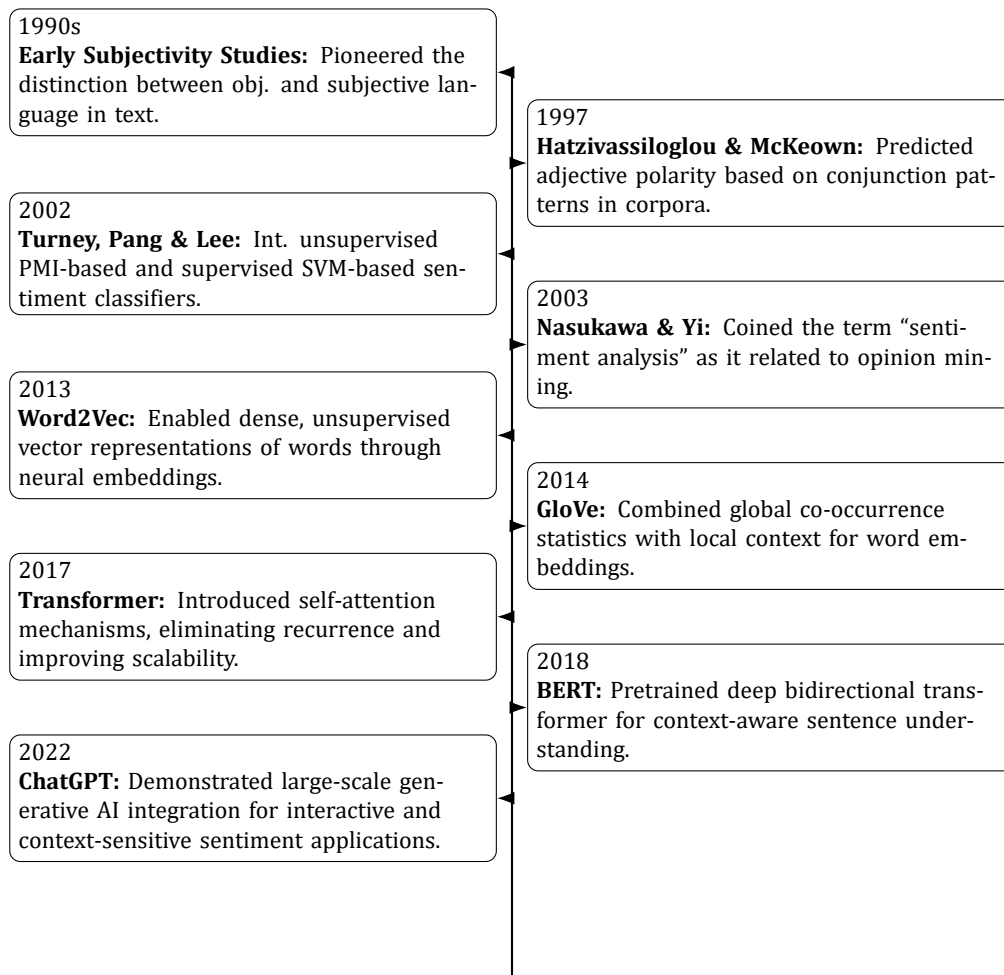


Figure 1. Timeline of relevant milestones in sentiment analysis.

The rest of this paper is organized as follows. Section 2 outlines the early foundations of sentiment analysis, focusing on lexicon-based methods, subjectivity detection, and semantic orientation. Section 3 covers classical machine learning techniques. Section 4 explores deep learning approaches, including neural networks and transformer architectures. Section 5 examines transfer learning, domain adaptation, and multimodal sentiment analysis. Section 6 investigates intelligent public opinion monitoring, detailing system architectures, key application domains, lessons learned, and illustrative case studies. Section 7 addresses ethical challenges, including bias, privacy, and representativeness. Section 8 discusses the limitations of our review. Section 9 concludes with practical takeaways and outlines future research directions in explainability, cross-lingual modeling, multimodal integration, and ethical AI for opinion mining.

To compile this review, we conducted a structured search of peer-reviewed journals, conference proceedings, and academic repositories spanning the period from 1990 to early 2025. Search terms included “sentiment analysis”, “opinion mining”, “multimodal sentiment”, “transfer learning”, and “domain adaptation”. We prioritized works that introduced influential algorithms, benchmarked performance across multiple datasets, or demonstrated novel applications across domains and languages. While not exhaustive, the selection captures representative methods and milestones across technologies and applications.

Unlike earlier surveys that examined lexical, machine-learning, or deep-learning approaches in isolation, our review integrates these perspectives, emphasizing the transition toward multimodal sentiment analysis and real-time opinion monitoring. We also discuss ethical considerations and practical lessons to provide a holistic perspective for researchers and practitioners.

To complement the timeline in **Figure 1**, **Table 1** summarizes representative NLP models and their core innovations, highlighting how advances in representation learning and architectures underpin modern sentiment analysis.

Table 1. Representative NLP models and their core innovations.

Model	Year	Core Innovation
Word2Vec	2013	Neural embeddings capturing distributional semantics via continuous bag-of-words and skip-gram architectures.
GloVe	2014	Global word co-occurrence statistics integrated with local context to produce robust embeddings.
Transformer	2017	Self-attention architecture enable parallel processing and long-range dependencies without recurrence.
BERT	2018	Deep bidirectional transformer pre-trained via masked language modeling and next-sentence prediction.
GPT (2–4)	2018–2023	Autoregressive transformer decoders trained with generative objectives, enabling zero-shot and few-shot tasks.
ChatGPT	2022	Large-scale generative AI fine-tuned with reinforcement learning from human feedback for interactive dialogue and sentiment-aware responses.

2. Early Foundations

Before delving into the specific techniques that shaped the nascent field of sentiment analysis, it is instructive to set the stage. The early foundations of the area were rooted in linguistics and semantic orientation, where researchers sought to understand how words and phrases convey subjectivity and polarity. This section surveys the seminal contributions that laid the groundwork for the transition from handcrafted lexica and linguistic heuristics toward more sophisticated computational models.

2.1. Lexicon and Linguistic Approaches

The earliest sentiment analysis research drew inspiration from linguistics and semantics. Wiebe et al. [20] annotated subjectivity in corpora, distinguishing subjective expressions from objective facts. Hatzivassiloglou and McKeown [16] demonstrated that the conjunctions linking adjectives (e.g., “good and bad” versus “good but bad”) can reveal their polarity, leading to algorithms that predict the semantic orientation of adjectives. Turney [3] extended this idea to multi-word phrases by applying PMI. Early lexicon-based methods compiled lists of positive and negative words, often manually or through bootstrapping from seed sets [21,22]. Building on these resources, Baccianella et al. [5] released SentiWordNet, a widely used lexicon that assigns sentiment scores to WordNet synsets. Cambria et al. [23] later introduced SenticNet, which leverages commonsense knowledge for concept-level sentiment analysis.

Lexicon approaches compute sentiment scores by summing or averaging the polarities of words. Given a document containing words w_1, \dots, w_n with base polarity s_i and context-dependent weight α_i (to account for intensifiers

and negation), a simple scoring function is

$$\text{Score} = \sum_{i=1}^n \alpha_i s_i \quad (1)$$

and words like “very” and “extremely” act as intensifiers ($\alpha_i > 1$), while negation flips the sign of subsequent sentiment words. Although transparent and interpretable, lexicon methods struggle with domain-specific vocabulary, sarcasm, and context dependency.

2.2. Subjectivity and Semantic Orientation

Identifying whether a sentence expresses an opinion or a fact—known as subjectivity classification—is a crucial task in sentiment analysis. Wiebe et al. [20] developed one of the first gold-standard corpora for subjectivity. Riloff and Wiebe [24] used bootstrapping to learn subjective words and patterns from unannotated data. Wilson et al. [7] introduced the notion of contextual polarity, noting that the sentiment of a word can shift depending on context (e.g., “not bad”).

Turney’s unsupervised PMI algorithm remains influential. It computes the semantic orientation of a phrase p relative to seed words p^+ (e.g., “excellent”) and p^- (e.g., “poor”):

$$\text{SO}(p) = \text{PMI}(p, p^+) - \text{PMI}(p, p^-) \quad (2)$$

where $\text{PMI}(x, y) = \log \frac{p(x, y)}{p(x)p(y)}$. A document is labeled positive if the average semantic orientation of its phrases exceeds zero [3]. This method requires only unannotated corpora and a set of seeded words.

3. Classical Machine Learning Techniques

The rise of supervised learning marked a turning point from handcrafted heuristics to data-driven sentiment models. Classical machine learning methods rely on transforming text into numerical feature vectors and using statistical algorithms to infer polarity from labeled examples [25]. This section reviews traditional approaches to sentiment classification, including supervised and semi-supervised learning, ensembles, and lexicon expansion strategies, highlighting their strengths and limitations.

3.1. Feature Engineering and Supervised Learning

With the availability of labeled datasets, researchers increasingly employed machine learning algorithms for sentiment classification. Pang et al. [4] compared Naive Bayes, maximum-entropy, and SVMs, demonstrating that SVMs outperformed the alternatives on movie review classification. They represented documents using a bag-of-words approach and bigrams, combined with term frequency-inverse document frequency (TF-IDF) weighting. Pang and Lee [6] later formulated star rating prediction as an ordinal regression problem.

Given a document vector \mathbf{x} extracted from text features (e.g., n-grams, part-of-speech tags, syntactic patterns), classifiers estimate the probability of a document being assigned a positive label using logistic regression:

$$P(y = 1 | \mathbf{x}) = \sigma(\mathbf{w}^T \mathbf{x} + b) = \frac{1}{1 + e^{-(\mathbf{w}^T \mathbf{x} + b)}} \quad (3)$$

Here \mathbf{w} and b are parameters learned from training data. SVMs similarly learn a decision boundary that maximizes the margin between positive and negative classes. Traditional models rely heavily on feature engineering: unigrams, bigrams, part-of-speech tags [7], dependency relations, and syntactic patterns [24]. Feature selection and weighting (e.g., mutual information or chi-square) have been extensively studied.

3.2. Ensemble Methods and Semi-Supervised Learning

To enhance robustness, researchers have combined multiple classifiers. For instance, Boiy and Moens [26] integrated a polarity lexicon with an ensemble of SVMs, which improved classification performance. Semi-supervised approaches exploit unlabeled data; both the expectation-maximization algorithm and self-training have been applied to expand training sets. Furthermore, domain adaptation techniques aim to align distributions between labeled source domains (e.g., movie reviews) and unlabeled target domains (e.g., tweets) by identifying pivot features shared across domains [27].

3.3. Sentiment Lexicon Expansion

Extensive lexicons provide valuable resources for traditional classifiers. Mohammad and Turney [28] developed the NRC Emotion Lexicon through crowdsourcing, which associates words with eight basic emotions. Baccianella et al. [5] constructed SentiWordNet, assigning sentiment scores to all WordNet synsets. Other notable resources include SentiCircle [29], SenticNet [23], and VADER (Valence Aware Dictionary and sEntiment Reasoner), which integrate lexical knowledge with heuristic rules for punctuation and emoji, making them particularly effective for analyzing informal and social media text.

4. Deep Learning Approaches

Recent advances in neural computation have revolutionized sentiment analysis. Deep learning models automatically learn hierarchical and distributed representations of language from large corpora, capturing subtle semantic and syntactic patterns that elude manual feature engineering [30]. We begin by reviewing the development of word embeddings and early neural architectures, then discuss the advent of attention and transformer models, and conclude with aspect-level and targeted sentiment analysis, which offers fine-grained insights [31].

4.1. Word Embeddings and Neural Networks

The adoption of distributed representations has fundamentally changed sentiment analysis. Mikolov et al. [17] introduced Word2Vec, which learns dense vector embeddings that capture semantic relationships. Pennington et al. [18] proposed GloVe, which combines global co-occurrence statistics with local contextual information. These embeddings serve as effective inputs for neural classifiers. Socher et al. [32] developed recursive neural networks (RNNs) over parse trees to model compositionality, while Kim [33] applied convolutional neural networks (CNNs) to sentence classification, achieving competitive performance with minimal parameter tuning. Recurrent neural networks (RNNs) and long short-term memory (LSTM) units further address sequential dependencies, enabling the modeling of variable-length inputs.

4.2. Attention and Transformer Models

Attention mechanisms allow models to concentrate on the most informative parts of the input. Bahdanau et al. [34] introduced the attention mechanism for machine translation, which was subsequently adopted in sentiment analysis to emphasize sentiment-bearing phrases. The transformer architecture, proposed by Vaswani et al. [19], relies on multi-head self-attention and positional encoding, giving rise to large pre-trained models such as BERT [8], RoBERTa, and XLNet. These models are trained on massive corpora using masked language modeling and next-sentence prediction objectives, and are later fine-tuned for sentiment analysis tasks. Radford et al. [9] further advanced this paradigm through generative pre-training with GPT-2, which demonstrated strong transferability to classification problems. Comprehensive surveys by Young et al. [35] and Qiu et al. [36] summarize these developments. Although deep learning models frequently achieve state-of-the-art performance, they require substantial computational resources and are prone to overfitting when applied to small datasets.

4.3. Aspect and Targeted Sentiment Analysis

Beyond overall polarity detection, fine-grained sentiment analysis identifies opinions directed toward specific aspects (e.g., battery life or service quality) or targets (entities). Early aspect extraction methods subsequently introduced attention mechanisms to capture target-specific information better; for example, Tang et al. [37] applied neural attention for aspect-based sentiment classification. More recently, BERT-based models have been adapted to jointly extract aspects and predict their corresponding sentiment, providing a more integrated and accurate analysis.

5. Transfer Learning, Domain Adaptation, and Multimodal Sentiment Analysis

Models trained on a particular dataset or modality often fail to generalize to new domains, languages, or communication channels because the underlying feature distributions differ. Transfer learning and domain adaptation techniques reuse knowledge learned from source data to improve performance on related tasks or domains [38]. In

parallel, multimodal sentiment analysis integrates heterogeneous signals—text, audio, and visuals—to capture sentiment beyond written language. Combining these fields highlights a common challenge: handling distributional shifts across domains and modalities while leveraging shared structures.

5.1. Transfer Learning and Domain Adaptation

Sentiment classifiers often experience performance degradation when applied to new domains or languages due to distribution shifts. Domain adaptation techniques address this issue by mapping feature spaces across domains.

Blitzer et al. [27] proposed the Structural Correspondence Learning (SCL) framework using pivot features, while Daumé III [39] suggested feature augmentation for more straightforward adaptation.

For cross-lingual sentiment analysis, researchers have employed strategies such as translating training corpora or mapping embeddings across languages using bilingual dictionaries [40]. Chen et al. [41] developed multilingual variants of SentiWordNet.

Unsupervised domain adaptation with adversarial training introduces a domain classifier to encourage domain-invariant features [42]. In this minimax setup, the sentiment classifier learns to predict labels, while the domain discriminator attempts to distinguish between source and target domains. Joint training encourages the shared representation to preserve sentiment information while remaining invariant to domain differences.

5.2. Multimodal Sentiment Analysis

Human communication extends beyond text; prosody in speech and facial expressions also conveys sentiment.

Early studies in emotion recognition combined speech and facial cues [43]. Poria et al. [10] developed a multimodal CNN to fuse visual, auditory, and textual features for emotion recognition in video. Tian et al. [11] proposed reinforcement learning to weight modalities adaptively, while Felbo et al. [44] trained deep models on billions of emoji occurrences to learn universal representations.

Multimodal sentiment analysis remains challenging due to the heterogeneous nature of data sources and the need for effective synchronization and fusion. Datasets such as CMU-MOSI and CMU-MOSEI have played a central role in advancing this field.

6. Smart Public Opinion Monitoring

Modern sentiment analysis does not occur in isolation; it is a core component of end-to-end systems that gather, analyze, and present public opinion in near real time. These “smart” monitoring platforms ingest massive streams of social media posts, news articles, and other user-generated content, process them with sentiment and topic models, and aggregate the results into actionable insights. In this section, we outline typical system architectures, review the wide range of application domains, distill lessons learned from practice, and conclude with a set of illustrative case studies [45].

6.1. System Architecture

An innovative opinion-monitoring system typically consists of (i) data acquisition through APIs or scraping, (ii) preprocessing to filter noise, remove spam, detect language, and normalize text, (iii) sentiment and topic analysis via classifiers described in previous sections, (iv) aggregation to produce time series or geographic distributions, and (v) visualization and decision support. Modern systems process millions of posts per day and must handle concept drift and adversarial content. Temporal smoothing (e.g., moving averages or Kalman filters) reduces volatility, as demonstrated by O'Connor et al. [14]. Topic modeling (e.g., Latent Dirichlet Allocation) helps identify salient issues and subpopulations within a dataset. Many systems also integrate opinion scores with external indicators (such as poll results or sales numbers), thereby enabling predictive analytics.

6.2. Applications and Domains

Opinion monitoring has influenced politics, finance, marketing, public health, and disaster response. Tumasjan et al. [15] found that the volume of tweets mentioning political parties during the German federal election aligned

closely with election outcomes. Asur and Huberman [12] showed that tweet rate and sentiment predicted box-office revenue. Bollen et al. [13] correlated public mood with stock market indices. Ceron et al. [46] applied sentiment analysis to Italian political tweets to track approval of government decisions. Pak and Paroubek [47] built a Twitter sentiment corpus using emoticons as distant supervision. Kouloumpis et al. [48] investigated the role of features such as hashtags and emoticons. Go et al. [49] used distant supervision to train sentiment classifiers on tweets. Chen et al. [50] used sentiment to forecast stock prices. These diverse applications illustrate both the promise and the challenges of interpreting noisy, user-generated text.

6.3. Lessons Learned

Lessons from practice include: (i) Volume matters—tweet count often correlates with outcomes as much as sentiment does [15]; (ii) Preprocessing and smoothing reduce noise and improve correlations [14]; (iii) Domain adaptation is essential—models trained on one domain rarely generalize well to another; (iv) Combining sentiment with other features (such as user influence, temporal patterns, and network structure) often yields better predictions; (v) Ethical and privacy considerations are indispensable.

6.4. Case Studies

Table 2 summarizes representative sentiment analysis techniques, highlighting their characteristics, advantages, and limitations. **Table 3** presents a broad sample of case studies spanning different domains. To provide a broad sense of progress over time, **Table 4** reports indicative accuracy ranges for different model families across common benchmarks. These indicative ranges demonstrate a clear trend: as approaches evolve from lexicon-based heuristics to classical machine learning, deep neural networks, and multimodal fusion, accuracy tends to increase—albeit at the cost of larger datasets, more complex models, and higher computational overhead.

Table 2. Comparison of sentiment analysis techniques. Typical characteristics are listed alongside advantages and limitations.

Category	Example References	Typical Characteristics	Advantages	Limitations
Lexicon based	[5,21,23]	Lists of positive/negative words and rules for intensifiers and negation; scores computed by aggregating word polarities	Transparent and interpretable; requires no labeled data; simple to implement	Domain-dependent; struggles with context, sarcasm, and irony; lexicon maintenance overhead
Classical ML	[4,7,26]	Hand-crafted features (n-grams, part-of-speech tags, syntactic patterns); supervised algorithms such as SVMs and Naive Bayes	Works with small datasets; flexible feature engineering; efficient training	Requires labeled data; feature selection and tuning are labor-intensive; limited semantic understanding
Deep learning (CNN/RNN/LSTM)	[32,33]	Pre-trained embeddings (e.g., Word2Vec, GloVe) fed into neural architectures such as CNNs, RNNs, or LSTMs	Learns representations automatically; captures local and long-range dependencies; delivers strong accuracy	Data-hungry; less interpretable; computationally expensive; prone to overfitting on small domains
Transformer models	[8,9,35]	Multi-head self-attention and positional encoding; large pre-trained models (e.g., BERT, GPT) fine-tuned for downstream tasks	Excellent transfer learning; captures long-range interactions; often state of the art	Requires massive computational resources; potential for encoded biases; limited interpretability; fine-tuning can be costly [51]
Multimodal fusion	[10,11,44]	Joint modeling of text with audio, visual, or physiological signals via multimodal encoders or fusion networks	Exploits non-textual cues (prosody, facial expressions); improves robustness on video/audio data	Necessitates synchronized multimodal datasets; architectures are complex; scarcity of large multimodal corpora

Table 3. Illustrative case studies of sentiment analysis for public opinion monitoring. For each study, we outline the domain, describe the data and methodology, report key outcomes, and give the principal reference.

Study	Domain	Data and Methodology	Outcome	Key Reference
O'Connor et al. (2010)	Politics	Data: 1 M tweets/day; Method: sentiment scoring plus Kalman smoothing	Sentiment time series correlated up to 80% with consumer confidence and presidential approval	[14]
Tumasjan et al. (2010)	Elections	Data: 100 k German election tweets; Method: LIWC lexicon and mention counts	Tweet volume predicted vote share; sentiment reflected political orientation	[15]
Asur & Huberman (2010)	Entertainment	Data: pre- and post-release tweets; Method: tweet rates and sentiment via regression	Pre-release tweet volume predicted opening-weekend revenue; sentiment improved predictions	[12]
Bollen et al. (2011)	Finance	Data: millions of tweets; Method: mood scores derived with OpinionFinder and GPOMS	Including mood dimensions improved Dow Jones forecasting accuracy to 87.6%	[13]
Ceron et al. (2014)	Politics	Data: Italian political tweets; Method: supervised classifiers	Sentiment trends matched public approval of government policies	[46]
Pak & Paroubek (2010)	Social media	Data: 1.6 M tweets labeled via emoticons; Method: Naive Bayes classifier	Created publicly available Twitter corpus; baseline accuracy of 62%	[47]
Kouloumpis et al. (2011)	Social media	Data: tweets with hashtags, emoticons, punctuation, and POS tags; Method: feature combination	Emoticons and hashtags were strong indicators; lexical features alone performed poorly	[48]
Go et al. (2009)	Social media	Data: tweets labeled by emoticons; Method: linear classifiers with distant supervision	Achieved 80% accuracy using simple features; established distant supervision paradigm	[49]
Baccianella et al. (2010)	Lexicon evaluation	Data: cross-language benchmark; Method: evaluate SentiWordNet lexicon	Demonstrated lexicon usefulness across domains; highlighted limitations in neutrality detection	[5]
Chen et al. (2014)	Finance	Data: news and social media sentiment; Method: regression models with macroeconomic variables	Combining sentiment with macroeconomic variables improved stock index forecasting	[50]
Poria et al. (2015)	Multimodal	Data: YouTube reviews with text, audio, and video; Method: deep CNN with multimodal fusion	Achieved significant gains over unimodal baselines; introduced concept-level fusion	[10]

Table 4. Indicative accuracy ranges for common sentiment analysis model families. The ranges are approximate and vary across datasets and domains.

Model Family	Typical Accuracy Range	Notes
Lexicon-based	50–65%	Effective for coarse polarity but limited by domain vocabulary and sarcasm.
Classical ML (e.g., SVM, Naive Bayes)	60–80%	Dependent on feature engineering; robust on specific domains with curated features.
CNN/RNN/LSTM	70–85%	Benefit from distributed embeddings; capture local and sequential dependencies.
Transformer-based (e.g., BERT, GPT)	80–90%	Pre-trained language models fine-tuned for sentiment; strong generalisation but computationally intensive.
Multimodal fusion	70–95%	Combines text, audio, and vision; performance depends on modality alignment and data quality.

7. Ethical Considerations and Challenges

As sentiment analysis systems increasingly inform high-impact decisions—spanning elections, financial forecasting, and public health responses—ethical considerations have moved to the forefront. Models trained on human language corpora can reflect and amplify social biases, infringe on individual privacy, and misrepresent public opin-

ion due to sampling noise. The following subsections discuss the key ethical and practical challenges that must be addressed to ensure fair, responsible, and trustworthy sentiment analysis [52].

7.1. Bias and Fairness

Sentiment classifiers trained on existing corpora may encode societal biases. Such models can inadvertently perpetuate stereotypes (e.g., associating specific names or dialects with negative sentiment). Efforts to mitigate bias include curating balanced training data, employing debiasing techniques, and applying fairness metrics to ensure equitable outcomes. Algorithmic transparency is also crucial when sentiment analysis is used to inform high-impact decisions.

7.2. Privacy and Consent

Opinion monitoring often relies on publicly available social media posts, but users may not anticipate that their expressions will be analyzed at scale. Aggregation and anonymization reduce privacy risks, yet questions remain regarding informed consent. Regulations such as the GDPR impose strict requirements for data processing and storage. Researchers must also consider ethical guidelines and institutional review procedures when collecting and analyzing user-generated [53].

7.3. Representativeness and Noise

Social media users do not accurately represent the general population, as their demographics tend to skew toward younger and more technologically engaged individuals. Bots and spammers further distort sentiment signals. Filtering, bot detection, and demographic weighting can mitigate these issues, but they are not foolproof. Triangulating social media sentiment with traditional surveys and other data sources can provide more reliable and balanced insights.

8. Limitations of the Review

While this survey strives to provide a comprehensive overview, several limitations warrant acknowledgment. First, our literature search cannot cover all papers published in the rapidly expanding field of sentiment analysis. We selected representative works based on citation impact and methodological diversity, meaning that some relevant studies may not be included. Second, performance figures reported here summarize typical ranges across datasets; individual results vary depending on preprocessing, model tuning, and domain. Third, our survey focuses primarily on English-language sources and high-resource languages; sentiment analysis in low-resource languages and cross-cultural contexts remains underexplored. Finally, we do not provide exhaustive coverage of all emerging techniques, such as large-scale generative models and few-shot prompting strategies, which continue to evolve rapidly.

9. Conclusions

Sentiment analysis has matured dramatically over the past three decades. What began as handcrafted lexicon lookups and simple statistical heuristics has evolved into a rich ecosystem of algorithms spanning classical and deep learning, domain adaptation, and multimodal fusion. In this paper, we have traced this evolution, linking early linguistic insights to the data-driven paradigms that now power commercial opinion monitoring systems. Along the way, we reviewed representative techniques, summarized their strengths and weaknesses in modern tables, and examined case studies that demonstrate both the promise and pitfalls of deploying sentiment models in real-world applications. Our review also highlighted the ethical considerations of public opinion mining, including bias, fairness, privacy, and representativeness. Taken together, these threads paint a holistic picture of a field that bridges linguistics, machine learning, human-computer interaction, and social science. Continued progress will hinge not only on technical breakthroughs but also on responsible deployment that respects individual rights and societal norms.

The following practical takeaways distill these insights for researchers and practitioners:

- **Align method with domain and data:** Lexicon-based and traditional machine-learning approaches are sufficient for coarse sentiment analysis in well-defined domains. In contrast, deep learning and transformer-based models excel when large, labeled datasets and computational resources are available.

- Mitigate domain shift: Employ transfer learning and domain adaptation techniques—including adversarial training and cross-lingual embedding alignment—when applying models across domains or languages.
- Leverage multimodality: Where possible, combine textual signals with audio, visual, or physiological cues to capture richer sentiment expressions and improve robustness.
- Prioritize fairness and privacy: Use balanced datasets, debiasing strategies, and anonymization to mitigate bias and protect user privacy in opinion monitoring applications.

Future Directions

Future research should address several open challenges: Explainability: developing interpretable models that reveal why a prediction was made will increase trust and facilitate debugging. Low-resource languages: cross-lingual transfer and multilingual pre-training can democratize sentiment analysis beyond English [40]. Multimodal understanding: integrating text with images, audio, and physiological signals will enable richer sentiment detection. Real-time adaptation: models must adapt to concept drift and emerging slang without extensive retraining. Causal inference: distinguishing correlation from causation in opinion dynamics could transform how sentiment informs policy and marketing. Ultimately, ethical frameworks for data use and algorithmic fairness must evolve in tandem with technical advancements.

Funding

This work received no external funding.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

No new data were created.

Conflict of Interest

The author declares no conflict of interest.

References

1. Nasukawa, T.; Yi, J. Sentiment analysis: Capturing favorability using natural language processing. In Proceedings of the 2nd International Conference on Knowledge Capture, Sanibel Island, FL, USA, 23–25 October 2003; pp. 70–77. [[CrossRef](#)]
2. Dave, K.; Lawrence, S.; Pennock, D.M. Mining the peanut gallery: Opinion extraction and semantic classification of product reviews. In Proceedings of the 12th International Conference on World Wide Web, Budapest, Hungary, 20–24 May 2003; pp. 519–528. [[CrossRef](#)]
3. Turney, P.D. Thumbs up or thumbs down? semantic orientation applied to unsupervised classification of reviews. In Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, Philadelphia, PA, USA, 6–12 July 2002; pp. 417–424.
4. Pang, B.; Lee, L.; Vaithyanathan, S. Thumbs up? sentiment classification using machine learning techniques. In Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing (EMNLP 2002), Philadelphia, PA, USA, 6–7 July 2002; pp. 79–86. [[CrossRef](#)]
5. Baccianella, S.; Esuli, A.; Sebastiani, F. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC 2010), Valletta, Malta, 17–23 May 2010; pp. 2200–2204.
6. Pang, B.; Lee, L. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating

- scales. In Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05), Ann Arbor, MI, USA, 25–30 June 2005; pp. 115–124. [CrossRef]
7. Wilson, T.; Wiebe, J.; Hoffmann, P. Recognizing contextual polarity in phrase-level sentiment analysis. In Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing, Vancouver, BC, Canada, 6–8 October 2005; pp. 347–354. [CrossRef]
 8. Devlin, J.; Chang, M.-W.; Lee, K.; et al. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv* **2018**, arXiv:1810.04805.
 9. Radford, A.; Narasimhan, K.; Salimans, T.; et al. Improving language understanding by generative pre-training. OpenAI Tech. Rep. 2018. Available from: https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf
 10. Poria, S.; Cambria, E.; Bajpai, R.; et al. Deep convolutional neural network text representations and multi-modal fusion for emotion recognition. In Proceedings of the IEEE International Conference on Automatic Face & Gesture Recognition (FG), Ljubljana, Slovenia, 4–8 May 2015; pp. 873–878.
 11. Tian, Y.; Xia, F.; Song, Y. Learning multimodal contrast with cross-modal memory and reinforced contrast recognition. In *Findings of the Association for Computational Linguistics ACL 2024*; Association for Computational Linguistics: Bangkok, Thailand, 2024; pp. 6561–6573.
 12. Asur, S.; Huberman, B.A. Predicting the future with social media. In Proceedings of the 2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology, Toronto, ON, Canada, 31 August–3 September 2010; pp. 492–499. [CrossRef]
 13. Bollen, J.; Mao, H.; Zeng, X.-J. Twitter mood predicts the stock market. *J. Comput. Sci.* **2011**, *2*, 1–8. [CrossRef]
 14. O'Connor, B.; Balasubramanian, R.; Routledge, B.; et al. From tweets to polls: Linking text sentiment to public opinion time series. *Proc. Int. AAAI Conf. Web Soc. Media* **2010**, *4*, 122–129. [CrossRef]
 15. Tumasjan, A.; Sprenger, T.O.; Sandner, P.; et al. Predicting elections with twitter: What 140 characters reveal about political sentiment. *Proc. Int. AAAI Conf. Web Soc. Media* **2010**, *4*, 178–185. [CrossRef]
 16. Hatzivassiloglou, V.; McKeown, K.R. Predicting the semantic orientation of adjectives. In Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics and 8th Conference of the European Chapter of the Association for Computational Linguistics, Madrid, Spain, 7–12 July 1997; Association for Computational Linguistics: Madrid, Spain, 1997; pp. 174–181. [CrossRef]
 17. Mikolov, T.; Chen, K.; Corrado, G.; et al. Efficient estimation of word representations in vector space. In Proceedings of the International Conference on Learning Representations (ICLR), Scottsdale, AZ, USA, 2–4 May 2013.
 18. Pennington, J.; Socher, R.; Manning, C.D. GloVe: Global vectors for word representation. In Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), Doha, Qatar, 25–29 October 2014; pp. 1532–1543.
 19. Vaswani, A.; Shazeer, N.; Parmar, N.; et al. Attention is all you need. In Proceedings of the 31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA, 4–9 December 2017; pp. 5998–6008.
 20. Wiebe, J.M.; Bruce, R.F.; O'Hara, T.P. Development and use of a gold-standard data set for subjectivity classifications. In Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics, College Park, MD, USA, 20–26 June 1999; pp. 246–253.
 21. Hu, M.; Liu, B. Mining and summarizing customer reviews. In Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Seattle, WA, USA, 22–25 August 2004; pp. 168–177. [CrossRef]
 22. Kim, S.-M.; Hovy, E. Determining the sentiment of opinions. In *COLING 2004: Proceedings of the 20th International Conference on Computational Linguistics*; COLING: Geneva, Switzerland, 2004; pp. 1367–1373.
 23. Cambria, E., Hussain, A. *Sentic Computing: A Common-Sense-Based Framework for Concept-Level Sentiment Analysis*; Springer: Cham, Switzerland, 2015. [CrossRef]
 24. Riloff, E.; Wiebe, J. Learning subjective nouns using extraction pattern bootstrapping. In Proceedings of the Seventh Conference on Natural Language Learning (CoNLL), Edmonton, AB, Canada, 31 May–1 June: 2003; pp. 25–32.
 25. Nápoles, G.; Hoitsma, F.; Knoblen, A.; et al. Prolog-based agnostic explanation module for structured pattern classification. *Inf. Sci.* **2023**, *622*, 1196–1227.
 26. Boiy, E.; Moens, M.-F. Boiy, E.; Moens, M.F. A machine learning approach to sentiment analysis in multilingual web texts. *Inform. Retrieval* **2009**, *12*, 526–558.
 27. Blitzer, J.; Dredze, M.; Pereira, F. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for

- sentiment classification. In Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics, Prague, Czech Republic, 25–27 June 2007; pp. 440–447.
28. Mohammad, S.M.; Turney, P.D. Crowdsourcing a word–emotion association lexicon. *Comput. Intell.* **2013**, *29*, 436–465. [[CrossRef](#)]
29. Poria, S.; Cambria, E.; Winterstein, G.; et al. Sentic patterns: Dependency-based rules for concept-level sentiment analysis. *Knowl.-Based Syst.* **2014**, *69*, 45–63. [[CrossRef](#)]
30. Leon, M.; Mkrtchyan, L.; Depaire, B.; et al. Learning and clustering of fuzzy cognitive maps for travel behaviour analysis. *Knowl. Inf. Syst.* **2014**, *39*, 435–462.
31. DeSimone, H.; Leon, M. Explainable ai: The quest for transparency in business and beyond. In Proceedings of the 2024 7th IEEE International Conference on Information and Computer Technologies (ICICT), Honolulu, HI, USA, 15–17 March 2024; pp. 532–538.
32. Socher, R.; Perelygin, A.; Wu, J.; et al. Recursive deep models for semantic compositionality over a sentiment treebank. In Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP), Seattle, WA, USA, 18–21 October 2013; pp. 1631–1642.
33. Kim, Y. Convolutional neural networks for sentence classification. In Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), Doha, Qatar, 25–29 October 2014; pp. 1746–1751.
34. Bahdanau, D.; Cho, K.; Bengio, Y. Neural machine translation by jointly learning to align and translate. In Proceedings of the International Conference on Learning Representations (ICLR), San Diego, CA, USA, 7–9 May 2015.
35. Young, T.; Hazarika, D.; Poria, S.; et al. Recent trends in deep learning based natural language processing. *IEEE Comput. Intell. Mag.* **2018**, *13*, 55–75.
36. Qiu, X.; Sun, T.; Xu, Y.; et al. Pre-trained models for natural language processing: A survey. *Sci. China Technol. Sci.* **2020**, *63*, 1872–1897. [[CrossRef](#)]
37. Tang, D.; Qin, B.; Liu, T. Aspect level sentiment classification with deep memory network. In Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, Austin, TX, USA, 1–4 November 2016; pp. 214–224.
38. Leon, M.; Depaire, B.; Vanhoof, K. Fuzzy cognitive maps with rough concepts. In *Artificial Intelligence Applications and Innovations: 9th IFIP WG 12.5 International Conference, AIAI 2013, Proceedings*; Springer Berlin Heidelberg: Berlin/Heidelberg, Germany, 2013; pp. 527–536.
39. Daumé, H., III. Frustratingly easy domain adaptation. In Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics, Prague, Czech Republic, 25–27 June 2007; pp. 256–263.
40. Abdalla, M.; Hirst, G. 2017. Cross-Lingual Sentiment Analysis Without (Good) Translation. In Proceedings of the Eighth International Joint Conference on Natural Language Processing, Taipei, Taiwan, 27 November–1 December 2017.
41. Abdellatif, A.; Sahmoud, S.; Nizam, A. A Unified Framework for Multi-Language Sentiment Analysis. In Proceedings of the 2023 3rd International Conference on Computing and Information Technology (ICCIT), Tabuk, Saudi Arabia, 13–14 September 2023.
42. Ganin, Y.; Lempitsky, V. Domain-adversarial training of neural networks. In Proceedings of the 33rd International Conference on Machine Learning, New York, NY, USA, 20–22 June 2016; pp. 209–219.
43. Ho, T.T.; Huang, Y. Stock price movement prediction using sentiment analysis and CandleStick chart representation. *Sensors* **2021**, *21*, 7957.
44. Felbo, B.; Mislove, A.; Søgaard, A.; et al. Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP), Copenhagen, Denmark, 7–11 September 2017; pp. 1616–1626.
45. Leon, M.; Martinez Sanchez, N.; Garcia Valdivia, Z.; et al. Concept maps combined with case-based reasoning in order to elaborate intelligent teaching/learning systems. In Proceedings of the Seventh International Conference on Intelligent Systems Design and Applications (ISDA 2007), Rio de Janeiro, Brazil, 20–24 October 2007; pp. 205–210.
46. Ceron, A.; Curini, L.; Iacus, S.M.; et al. Every tweet counts? how sentiment analysis of social media can improve our knowledge of citizens’ political preferences. *New Media Soc.* **2014**, *16*, 340–358. [[CrossRef](#)]
47. Pak, A.; Paroubek, P. Twitter as a corpus for sentiment analysis and opinion mining. In Proceedings of the Seventh Conference on International Language Resources and Evaluation (LREC 2010), Valletta, Malta, 17–23 May 2010; pp. 1320–1326.
48. Kouloumpis, E.; Wilson, T.; Moore, J.D. Twitter sentiment analysis: The good the bad and the omg! *Proc. Int.*

- AAAI Conf. Web Soc. Media **2011**, 5, 538–541.
49. Go, A.; Bhayani, R.; Huang, L. Twitter sentiment classification using distant supervision. 2009. Available from: <https://cs.stanford.edu/people/alecmgo/papers/TwitterDistantSupervision09.pdf>
50. Maqbool, J.; Aggarwal, P.; Kaur, R.; et al. Stock prediction by integrating sentiment scores of financial news and MLP-regressor: A machine learning approach. *Procedia Comput. Sci.* **2023**, 218, 1067–1078.
51. Leon, M. The escalating ai's energy demands and the imperative need for sustainable solutions. *WSEAS Trans. Syst.* **2024**, 23, 444–457.
52. Leon, M.; Nápoles, G.; García, M.M.; et al. Two steps individuals travel behavior modeling through fuzzy cognitive maps pre-definition and learning. In *Advances in Soft Computing: 10th Mexican International Conference on Artificial Intelligence, MICAI 2011, Proceedings, Part II*; Springer Berlin Heidelberg: Berlin/Heidelberg, Germany, 2011; pp. 82–94.
53. Leon, M. Aggregating procedure for fuzzy cognitive maps. *Int. FLAIRS Conf. Proc.* **2023**, 36. [[CrossRef](#)]



Copyright © 2025 by the author(s). Published by UK Scientific Publishing Limited. This is an open access article under the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Publisher's Note: The views, opinions, and information presented in all publications are the sole responsibility of the respective authors and contributors, and do not necessarily reflect the views of UK Scientific Publishing Limited and/or its editors. UK Scientific Publishing Limited and/or its editors hereby disclaim any liability for any harm or damage to individuals or property arising from the implementation of ideas, methods, instructions, or products mentioned in the content.