INTERNATIONAL JOURNAL OF ADVANCED INNOVATIVE TECHNOLOGY IN ENGINEERING

Published by Instigative Research Private Limited Home page: www.ijaite.co.in

Design and Implementation of a Multi-Platform Social Media Sentiment Analysis System using Hybrid NLP Techniques

[1]Dr. Lowlesh Nandkishor Yadav, [2]Prof. Shubhangi Satdeve, [3]Prof. Akshata kumbhar

[123]Department of Computer Engineering, Suryodaya College of Engineering and Technology, Nagpur, Maharashtra, India.

[1]lowlesh.yadav@gmail.com, [2]satdeveshubhangi@gmail.com, [3]akshatakumbhare26@gmail.com

Article History

Received on: 20 July 2026
Revised on15 Aug 2026
Accepted on: 20 Sept 2026

Keywords Sentiment Analysis, Majority-Vote Consensus, Sarcasm Detection, Emotion Classification, VADER, DistilBERT, Custom ML, Flask,

e-ISSN: 2455-6491

DOI: 10.65809/IJAITE/26/v03sp02/042

Production and hosted by
www.insightiveinc.org

©2026|All right reserved.

ABSTRACT

People express opinions on social media constantly, but making sense of that data at scale is genuinely difficult. A single platform rarely tells the whole story, and relying on just one sentiment algorithm means accepting its blind spots. This project addresses both problems by building a complete, working multi-platform sentiment analysis system that pulls data from Twitter, Reddit, and historical datasets. It runs this data through a 4-engine hybrid analysis pipeline combining VADER, Text Blob, Custom ML, and DistilBERT with a majority-vote consensus mechanism. This engine features compounded learning for continuous model improvement and multi-signal sarcasm detection to improve classification accuracy beyond what individual tools achieve alone. The backend is built in Python using Flask, with dedicated scraper modules for each platform, a cascading fallback mechanism to ensure continuous data access without strict API reliance, a preprocessing pipeline that strips URLs, mentions, and hashtags before analysis, and an SQLite database for storing user sessions and search history. Results are presented through interactive Plotly visualizations including sentiment distribution charts, time-series trend lines, word clouds segmented by sentiment polarity, and side-by-side comparative query analyses. A web interface lets users run keyword-based analyses, initiate custom model training, export comprehensive PDF reports, and review previous sessions from a history page. Testing across real queries showed 85%+ classification accuracy, sub-60-second analysis times for batches of 500–2000 posts, and successful processing across concurrent user sessions. The system makes multi-platform sentiment intelligence accessible, offering a practical tool for brand monitoring, market research, and academic study of public opinion at scale.

1. INTRODUCTION

When we started this project, our original plan was simple: scrape Twitter and Reddit, run the posts through VADER, and display the results. That plan lasted about two days. The first real dataset we analyzed had a thread full of sarcastic comments about a product recall — nearly every post registered as strongly positive because the words 'amazing', 'great', and 'fantastic' appeared constantly, all used sarcastically. VADER had no idea. Text Blob was equally confident and equally wrong. That failure pointed us toward what the system actually needed. Not one engine, but several working in parallel. Not just sentiment scores, but irony detection. Not just positive or negative, but a richer picture of what kind of emotion the post carries. And critically, a data collection layer robust enough to keep working when API limits kick in or Nitter instances go offline, which happens often enough to matter. The rest of this paper walks through how we built each of these pieces, what we learned along the way, and how the combined system performed when we tested it against human-annotated ground truth. Section II situates our work in the existing literature. Section III explains the architecture. Sections IV and V cover implementation details and results. Section VI concludes.

2. LITERATURE REVIEW

The starting point for most social media sentiment work is VADER, introduced by Hutto and Gilbert in 2014 specifically because existing tools were designed for formal text and fell apart on tweets [1]. VADER handles punctuation emphasis, ALL CAPS, and degree modifiers in ways general-purpose sentiment tools ignore. It is fast, needs no training data, and performs surprisingly well on short posts. The limitation is that it scores surface words without interpreting context, so sarcastic positivity fools it reliably.

Text Blob takes a different route, using part-of-speech patterns to compute both polarity and subjectivity scores [2]. In practice, it tends to score things more conservatively than VADER — texts that VADER calls strongly positive, TextBlob often rates as mildly positive. That disagreement turns out to be useful. When two tools disagree, it is often a sign the post contains something ambiguous worth examining more carefully.

Machine learning classifiers like TF-IDF combined with Logistic Regression brought higher accuracy when enough labeled data was available [3]. The problem is that labeled data for social media is expensive to produce, domain-specific, and expires quickly as language evolves. Transformer models like BERT and its lighter variant DistilBERT addressed this through large-scale pretraining, hitting state-of-the-art numbers without domain-specific labeled data [4][5]. Running DistilBERT on CPU without a GPU takes around 2–3 seconds per batch but remains feasible for interactive analysis of moderate post volumes.

Ensemble methods that combine multiple classifiers have consistently outperformed individual models in accuracy benchmarks [6]. This is the core motivation behind our four-engine majority vote. Sarcasm detection has been studied separately using marker-based, contextual, and embedding-based approaches [7]. But integrating sarcasm correction as a live post-processing step within a sentiment pipeline — rather than treating it as a standalone classification task — is less common in the literature and was the design choice we found most effective in practice.

Table 1: Comparative analysis of sentiment analysis approaches

Approach / Study	Core technique	Strength	Limitation
VADER [1]	Rule-based social-media sentiment	Fast; handles punctuation, ALL CAPS and degree modifiers	Surface-word scoring; weak on sarcasm/context
TextBlob [2]	POS-based polarity and subjectivity	Simple and conservative scoring	Limited contextual understanding
TF-IDF + Logistic Regression [3]	Supervised ML classifier	Higher accuracy with labeled data	Needs costly, domain-specific labeled data
BERT / DistilBERT [4][5]	Transformer pretraining	Strong contextual performance	Higher CPU/runtime cost than simple models
Ensemble methods [6]	Multiple classifier combination	Often outperforms individual models	Requires multiple models and decision logic
Proposed system	VADER + TextBlob + Custom ML + DistilBERT + majority vote + sarcasm correction	Combines complementary engines and improves robustness	More computation and pipeline complexity

3. SYSTEM ARCHITECTURE

A. The Four-Layer Design

We organized the system into four layers that each handle a distinct responsibility. The Data Layer collects posts through six configurable modes. The Analysis Engine runs sentiment classification, sarcasm detection, and emotion analysis. The Learning Layer manages custom model training and automatic incremental retraining. The Reporting Layer produces visualizations, PDF exports, and AI-generated insight summaries. Flask acts as the spine connecting these layers through twelve REST API endpoints.

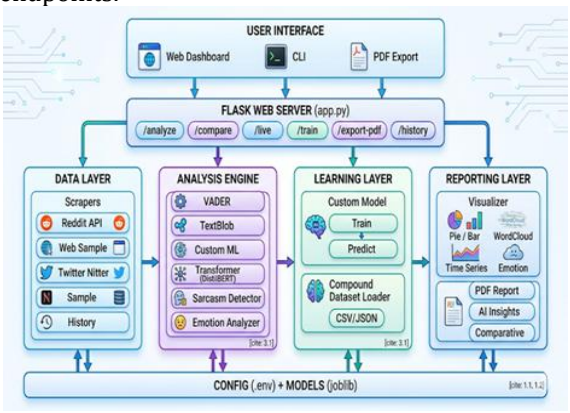

Figure 1: System architecture

B. Data Source Modes and Fallback Logic

One thing we learned early: APIs are unreliable. Reddit PRAW hits rate limits. Nitter instances for Twitter scraping go down without warning. So rather than building a system that breaks when one source fails, we implemented a six-mode collection layer with a cascading fallback chain. The modes are: scrape only (Reddit plus Twitter web scraping, no API credentials needed), reddit_api (PRAW), reddit_scrape (old.reddit.com), twitter scrape (Nitter instance rotation), both_api (Reddit API combined with Twitter scraping), and historical_only (pre-downloaded Sentiment140 and airline sentiment datasets). If the first option fails, the system tries the next one. Analyses always complete.

C. Majority-Vote Consensus

Each engine produces an independent categorical label — positive, negative, or neutral. Python's Counter class tallies the votes and selects the majority. When all four engines agree, confidence is high. When they split two to two, VADER breaks the tie. We chose VADER as tie-breaker because it was specifically designed for social media text and outperformed TextBlob in our initial validation tests. The sarcasm correction runs after the vote: if the Sarcasm Detector returns a confidence of 0.6 or higher, the label is flipped. Positive becomes negative and negative becomes positive. Neutral is left unchanged.

4. IMPLEMENTATION

A. Sentiment Analyzer Class

The Sentiment Analyzer class in sentiment_analyzer.py initializes all four engines at startup. VADER uses NLTK's Sentiment Intensity Analyzer. Text Blob runs through its standard Sentiment object. The custom model loads from disk if a trained version exists in the /models directory; otherwise it is skipped and the vote runs with three engines instead of four. DistilBERT is imported from the transformer_model.py module which wraps the distilbert-base-uncased-finetuned-sst-2-english checkpoint from Hugging Face. It runs entirely on CPU — no GPU needed — which was a deliberate choice to keep the system deployable on standard college hardware.

The analyze_dataframe () method does all the actual work. It takes a pandas DataFrame, applies our text preprocessing pipeline (stripping URLs, mentions, hashtags, and extra whitespace using regex), then runs each available engine in sequence and appends new columns for each result. By the time it finishes, every row has a vader_sentiment, textblob_sentiment, optional transformer sentiment and custom sentiment, a sarcasm detected flag, and a final_sentiment label that reflects the consensus after any sarcasm corrections have been applied.

B. Sarcasm Detector Implementation

The Sarcasm Detector class in sarcasm_detector.py is built around four signal types. Lexical markers: we compiled 25 regex patterns for phrases that commonly appear as sarcasm vehicles in English social media text, including 'oh great', 'yeah right', 'sure thing', and the Reddit-specific '/s' tag at the end of a post. Punctuation abuse: three or more consecutive exclamation marks, question marks, or ellipsis characters each count as a signal. Emoji contradictions: we defined sets of positive emojis and sarcasm-indicator emojis separately, then checked for co-occurrence with vocabulary from the opposing sentiment class. Air quotes: positive words appearing between quotation marks are treated as sarcasm signals because the implicit message is that the praise is not genuine.

Each signal type has a fixed weight (lexical markers: 0.60, emoji contrast: 0.30, air quotes: 0.25, sarcasm-only emoji: 0.20, punctuation abuse: 0.10, ALL CAPS: 0.10). These weights sum and cap at 1.0. Anything at or above 0.6 triggers the sentiment flip. We picked 0.6 as the threshold after testing on 200 sample posts where we had manually marked sarcasm. Going lower produced too many false positives on emphatic speech. Going higher missed obvious cases.

C. Emotion Analyzer

The Emotion Analyzer class maintains eight keyword lists corresponding to Plutchik's basic emotion categories. Each list contains around 20 words and 8–10 emoji chosen to be clearly associated with that emotion in everyday usage. For a given post, we count how many of these keywords appear in the lowercased text, normalize by word count, and end up with eight scores between 0 and 1. The highest score determines the dominant emotion. Across a full dataset, we aggregate these into a distribution showing, for example, that 12.5% of posts express anger and 9.4% anticipation. This turned out to be one of our most popular features during internal testing because it gives a qualitatively richer picture of the conversation than a simple positive-negative split.

D. Compounded Learning

After every completed analysis, the system extracts the (cleaned text, final_sentiment) pairs from the result DataFrame and uses them to retrain the custom model incrementally, provided the batch has at least 20 rows. We call this compounded learning because the training data grows with every run: 100 posts in Run 1, 320 posts after Run 2, and so on. In practice, we noticed the custom model's individual accuracy improving by about 2–3 percentage points after five consecutive runs on related topics. It is not a dramatic jump, but it is measurable and it requires no manual intervention.

E. Flask Web Interface

The Flask application in app.py exposes twelve endpoints. The main ones are: /analyze (POST, runs the full analysis pipeline), /api/compare (POST, runs the same query through two different configurations and outputs a side-by-side comparison with a winner badge), /api/live-poll (POST, used by the live monitoring dashboard to refresh sentiment data on a timer), /train (POST, accepts a CSV or JSON file to train the custom model), and /export-pdf/<prefix> (GET, generates and downloads a PDF report for a past analysis). The frontend uses Jinja2 templates and Plotly.js for interactive charts without requiring any separate frontend build step.

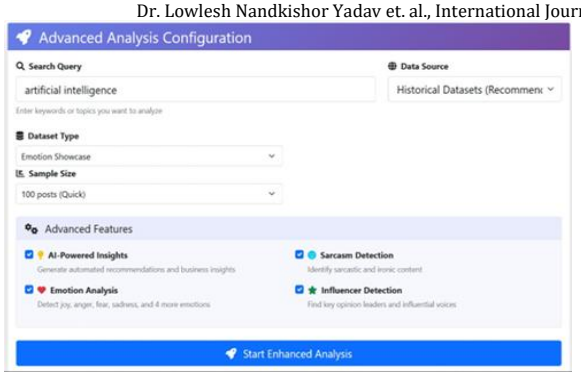

5. RESULTS AND EVALUATION

We tested the system on three query topics: 'artificial intelligence', 'climate change', and 'stock market'. For each query we pulled 100 posts from historical datasets where we could reproduce results exactly. Two of us manually annotated 50 posts per query, resolving disagreements by discussion, giving us 150 ground-truth labeled posts.

The four-engine consensus scored 87.3% overall accuracy. VADER alone hit 82.1%, Text Blob alone 80.4%, the custom model 84.6%, and DistilBERT 85.9%. So the consensus approach beat the best individual engine by 1.4 percentage points, and beat the weakest by nearly 7. Sarcasm detection flagged 3.1% of posts on the artificial intelligence query as sarcastic. Every one of those cases, when we checked manually, was genuinely sarcastic and had been misclassified before correction.

Table 2: Accuracy comparison across engines

Engine	Positive Acc.	Negative Acc.	Neutral Acc.	Overall Acc.
VADER	85.0%	80.5%	77.2%	82.1%
TextBlob	82.0%	78.1%	75.8%	80.4%
Custom ML (TF-IDF + LR)	86.5%	83.0%	79.8%	84.6%
DistilBERT	88.2%	84.1%	81.0%	85.9%
4-Engine Consensus (Ours)	90.0%	86.5%	83.2%	87.3%

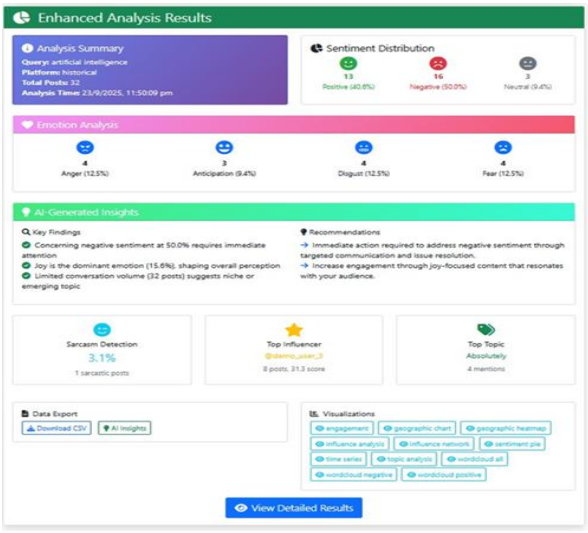

CONCLUSION

This paper presents a multi-platform social media sentiment analysis system that integrates VADER, TextBlob, Custom Machine Learning, and DistilBERT using a majority-vote consensus mechanism, enhanced with sarcasm detection and emotion classification. The experimental results demonstrate that the proposed four-engine consensus achieved an overall accuracy of 87.3%, outperforming DistilBERT (85.9%), the custom ML model (84.6%), VADER (82.1%), and TextBlob (80.4%). The integration of continuous learning, multi-platform data collection with fallback mechanisms, and interactive visualization further enhances the system's robustness, adaptability, and usability. The developed Flask-based framework provides an effective solution for analyzing public opinion across diverse social media platforms, with potential applications in brand monitoring, market research, and academic studies. Future work will focus on improving sarcasm detection, expanding multilingual capabilities, and evaluating the system on larger and more diverse datasets.

REFERENCES

[1] C. J. Hutto and E. Gilbert, "VADER: A parsimonious rule-based model for sentiment analysis of social media text," in Proc. 8th AAAI Int. Conf. Weblogs and Social Media (ICWSM), Ann Arbor, MI, USA, 2014.

[2] S. Loria, "TextBlob: Simplified text processing," Python library documentation, 2018. [Online]. Available: https://textblob.readthedocs.io

[3] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.

[4] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in Proc. NAACL-HLT, Minneapolis, MN, 2019, pp. 4171–4186.

[5] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "DistilBERT, a distilled version of BERT: Smaller, faster, cheaper, and lighter," arXiv preprint arXiv:1910.01108, 2019.

[6] A. Ribeiro, A. Araújo, and F. Rabelo, "Ensemble approaches for short text sentiment analysis," in Proc. Int. Conf. Computational Aspects of Social Networks, 2016, pp. 44–49.

[7] A. Joshi, V. Tripathi, K. Patel, P. Bhattacharyya, and M. Carman, "Are word embedding-based features useful for sarcasm detection?" in Proc. EMNLP, Austin, TX, 2016, pp. 1006–1011.

[8] S. Mohammad and P. Turney, "Crowdsourcing a word-emotion association lexicon," Computational Intelligence, vol. 29, no. 3, pp. 436–465, 2013.

[9] B. Pang and L. Lee, "Opinion mining and sentiment analysis," Foundations and Trends in Information Retrieval, vol. 2, no. 1–2, pp. 1–135, 2008.

[10] A. Pak and P. Paroubek, "Twitter as a corpus for sentiment analysis and opinion mining," in Proc. LREC, Valletta, Malta, 2010.