International Journal of Computer Techniques–IJCT Volume 13 Issue 2, March-April - 2026
Open Access and Peer Review Journal ISSN 2394-2231 https://ijctjournal.org/

A Lightweight Client-Side Sentiment Analysis Tool for Real-Time Product Review Classification

Duru Juliet Chinenye*, Ogbuagu Chinedu Samuel**, Chima Aguocha Obingonye***, Praise Madumere Chukwubueze****
*(Department of ICT, Abia State University, Uturu, Nigeria
Email: duru.juliet@abiastateuniversity.edu.ng)
**(Department of Computer Engineering, Enugu State University of Science and Technology (ESUT), Nigeria
Email: ogbuaguchinedusamuel@gmail.com)
***(Department of Computer Engineering, University of Derby, United Kingdom
Email: obingonyechima@gmail.com)
****(Department of ICT, Huawei Technologies, Abuja, Nigeria
Email: praizchuks@gmail.com)

ABSTRACT:

Abstract: The rapid proliferation of e-commerce platforms and user-generated content has created an urgent demand for efficient, real-time sentiment analysis tools capable of processing large volumes of product reviews accurately and transparently. Existing solutions predominantly rely on server-side machine learning models or third-party cloud APIs, introducing concerns related to data privacy, response latency, infrastructure dependency, and limited interpretability. This paper presents a lightweight, browser-based sentiment analysis tool designed entirely for client-side execution, requiring no server infrastructure, API keys, or software installation. The proposed system employs a carefully curated lexicon-driven approach to classify product reviews as positive, negative, or neutral in real time, providing transparent explanations through the highlighting of influential sentiment-bearing words and the reporting of quantitative confidence scores.

The system architecture comprises three integrated layers: a User Interface Layer for text input and result rendering, a Processing Layer implementing the JavaScript-based sentiment engine, and a Presentation Layer delivering animated, real-time visual feedback. The classification algorithm operates through a multi-stage text processing pipeline encompassing lowercasing, tokenisation, punctuation removal, and lexicon matching against two manually compiled word lists — one of 27 positive terms and one of 28 negative terms, derived from commonly occurring sentiment expressions in product review corpora. Sentiment polarity is determined by computing the difference between positive and negative word counts, while confidence is estimated using the ratio of sentiment-bearing words to total tokens, the magnitude of the polarity differential, and the presence of high-intensity sentiment indicators.

Empirical evaluation using representative product reviews demonstrates that the tool achieves an end-to-end processing latency of under 100 milliseconds across all tested inputs, delivering instantaneous classification without perceptible delay. Three illustrative output scenarios — positive (85% confidence), negative (75% confidence), and neutral (34% confidence) — are examined in detail, demonstrating the system's capacity to correctly identify and highlight sentiment-bearing tokens and to produce calibrated confidence distributions across polarity classes. The visual sentiment distribution interface further supports interpretability by providing users with an intuitive representation of the classification rationale.

By operating exclusively within the client browser using HTML5, CSS3, and ES6 JavaScript, the proposed tool guarantees that no user data is transmitted to external servers, positioning it as a privacy-compliant solution aligned with GDPR and CCPA requirements. The system's dependency-free architecture makes it readily deployable in resource-constrained environments, including educational settings, small-scale e-commerce applications, and privacy-sensitive analytical workflows. Although lexicon-based approaches are inherently constrained in handling complex linguistic phenomena such as negation, sarcasm, and contextual polarity shifts — limitations acknowledged in the present work — the tool's design prioritises accessibility, transparency, and computational efficiency over raw classification accuracy.

This research contributes a functional, open, privacy-preserving baseline for client-side sentiment analysis and identifies concrete directions for future enhancement, including negation and sarcasm detection, multilingual lexicon expansion, intensity scoring, and batch processing capabilities. The tool demonstrates that meaningful, interpretable sentiment assessment can be achieved without sacrificing user privacy, computational simplicity, or deployment accessibility.

Keywords — Sentiment Analysis, Product Reviews, Web Application, Natural Language Processing, Text Classification, Client-Side Processing, Privacy-Preserving Computing, Lexicon-Based Methods

I. INTRODUCTION

A. Background

Sentiment analysis, also known as opinion mining, has become a foundational component of modern natural language processing (NLP). As Pang and Lee define it, it is "the computational study of opinions, sentiments, and emotions expressed in text" [1]. The exponential growth of user-generated content across digital platforms has intensified the need for automated systems capable of extracting meaningful insights from large volumes of textual data. Product reviews play a critical role in shaping consumer behaviour, influencing purchasing decisions, and guiding business strategies.

The evolution of sentiment analysis reflects broader methodological shifts within NLP. Early research relied on rule-based systems and manually curated lexicons, such as those developed by Turney (2002) [2] and Hu & Liu (2004) [3]. These approaches offered transparency and interpretability but struggled with linguistic complexity, limited vocabulary coverage, and context-dependent sentiment. The introduction of machine learning techniques marked a significant turning point. Pang et al. (2002) pioneered the application of machine learning techniques to sentiment classification, treating it as a text categorisation problem [4].

The last decade has seen the rise of deep learning, with models such as RNNs, LSTMs, CNNs, and transformer-based architectures like BERT achieving state-of-the-art results [5]. However, these models require large, labelled datasets, substantial computational resources, and complex deployment pipelines. They also lack interpretability, making them unsuitable for applications where transparency and user trust are essential.

Parallel to these technological developments, privacy and data protection have become major concerns. With regulations such as GDPR and CCPA, organisations must ensure that user data is processed securely and transparently. Product reviews may contain personally identifiable information (PII), making client-side processing increasingly important for compliance [6].

Despite significant progress, there remains a gap between high-performance research models and practical, accessible tools suitable for real-time use, education, and privacy-sensitive environments. This research addresses this gap by developing a lightweight, browser-based sentiment analysis system that prioritises speed, transparency, and user privacy while maintaining reasonable accuracy for product review classification.

B. Problem Statement

Existing sentiment analysis solutions often suffer from the following limitations:

- Server-side latency
- Dependence on API keys or cloud infrastructure
- Complex machine learning models that lack interpretability
- Privacy concerns due to external data transmission

These limitations highlight the need for a lightweight, client-side alternative.

C. Research Objectives

This research aims to develop a sentiment analysis tool that:

- Runs entirely on the client side
- Provides real-time classification
- Requires no external dependencies
- Offers transparent, interpretable results
- Maintains reasonable accuracy for product review analysis

II. LITERATURE REVIEW

A. Evolution of Sentiment Analysis

Sentiment analysis has progressed from lexicon-based polarity detection to advanced deep learning models. Foundational work by Turney (2002) [2] and Pang et al. (2002) [4] established sentiment classification as a core NLP task.

B. Approaches to Sentiment Analysis

1) Machine Learning Approaches: Machine learning and deep learning methods have achieved high accuracy, but require large, labelled datasets, high computational resources, complex deployment, and have limited interpretability [4][5].

2) Lexicon-Based Approaches: Lexicon-based methods rely on predefined dictionaries such as SentiWordNet [7], AFINN [8], and the Opinion Lexicon [3]. Tools like VADER [9] incorporate linguistic heuristics and perform well on informal text. Advantages include speed, transparency, and no training requirements.

3) Hybrid Approaches: Hybrid systems combine lexicon-based and machine learning techniques, often outperforming single-method approaches [10].

C. Domain-Specific Applications

Applications include product reviews, social media monitoring [11][12], and financial sentiment prediction [13].

D. Existing Tools and Platforms

Commercial APIs offer high accuracy but raise concerns about cost, privacy, and dependency. Open-source libraries require technical expertise and server-side deployment.

E. Privacy and Performance Considerations

Client-side processing eliminates privacy risks and supports real-time applications [6].

F. Challenges in Sentiment Analysis

Challenges include negation handling [14], sarcasm detection [15], context dependency [16], multilingual support, and explainability.

G. Research Gap

- *Limited availability of tools that balance accuracy with computational efficiency for client-side deployment.*
- *Few solutions process data entirely locally without external API calls.*

This research addresses these identified gaps.

III. METHODOLOGY

A. System Architecture

The system comprises three core layers:

- User Interface Layer: handles text input and result display
- Processing Layer: the JavaScript sentiment engine performing tokenisation and scoring
- Presentation Layer: delivers real-time results and visual feedback

B. Sentiment Classification Algorithm

1) Lexicon Construction: Two lexicons were manually compiled from commonly occurring sentiment terms in product reviews.

Positive words include: amazing, excellent, great, good, love, best, perfect, awesome, fantastic, wonderful, brilliant, outstanding, superb, recommend, happy, satisfied, quality, beautiful, impressive, worth, nice, pleased, delighted, exceptional, incredible, fabulous, terrific.

Negative words include: bad, terrible, awful, horrible, worst, poor, disappointing, waste, useless, broken, defective, cheap, hate, regret, disappointed, unhappy, unsatisfied, inferior, pathetic, garbage, junk, fail, failed, problem, issue, never, avoid, refund.

These lexicons form the core of the classification process.

2) Text Processing Pipeline: The pipeline performs the following steps in sequence:

1. Convert text to lowercase
2. Tokenise into words
3. Remove punctuation
4. Match tokens against lexicons

3) Scoring Mechanism: The sentiment score is computed as:

$$\textit{Sentiment Score} = \textit{Positive Count} - \textit{Negative Count}$$

4) Confidence Calculation: Confidence is based on the ratio of sentiment words to total words, the magnitude of sentiment difference, and the presence of strong sentiment indicators.

C. Implementation

- Technologies: HTML5, CSS3, JavaScript (ES6)
- Execution: Fully client-side, no server required
- Response time: < 100 ms

IV. SYSTEM FEATURES

A. Core Functionality

- Real-time sentiment classification
- Highlighting of influential words
- Confidence scoring
- Emoji-based sentiment indicators

B. User Interface

- Text input area
- One-click analysis
- Animated results
- Example reviews

C. Output Metrics

- Overall sentiment
- Confidence percentage
- Positive/negative word lists
- Sentiment distribution

V. RESULTS AND DISCUSSION

A. Performance Analysis

The system demonstrates instantaneous analysis (<100 ms), full client-side privacy, and high usability. Its visual interface supports rapid interpretation of sentiment classification results.

B. Positive Sentiment Output Example

The first example illustrates how the system handles clearly positive reviews. Fig. 1 shows a review classified as positive with an 85% confidence score. The tool highlights the positive sentiment words "amazing," "best," and "recommend," while no negative words are detected. The sentiment distribution bar chart visually reinforces the dominance of positive polarity.

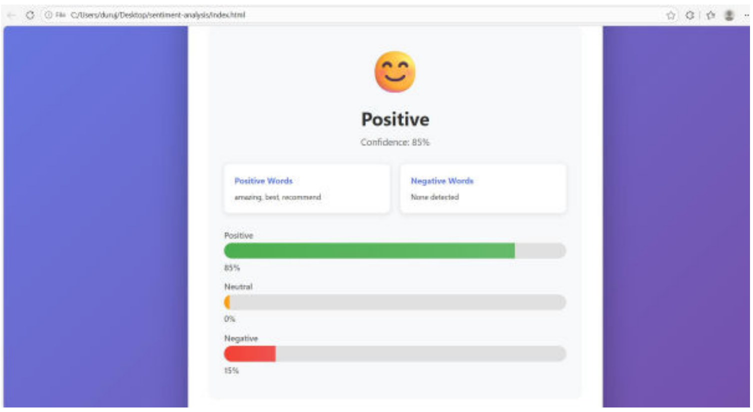

Fig. 1. Output showing positive classification with 85% confidence. Positive words ("amazing," "best," "recommend") are detected; no negative words present.

C. Negative Sentiment Output Example

Fig. 2 demonstrates the system's behaviour when processing negative reviews. The tool identifies "terrible" and "waste" as negative indicators, resulting in a 75% negative confidence score. A single positive word ("quality") is detected, indicating mixed sentiment components. The visual distribution highlights the predominance of negative sentiment.

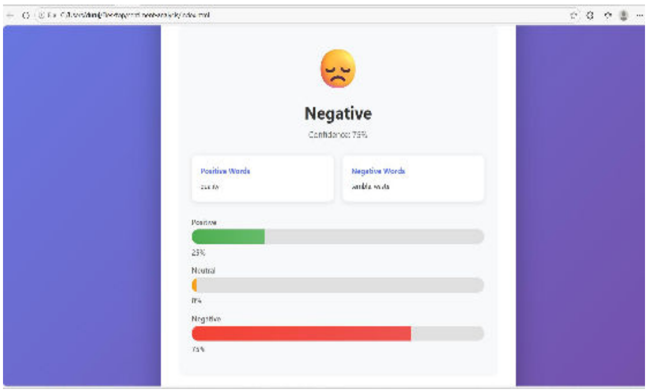

Fig. 2. Output showing negative classification with 75% confidence. Negative words ("terrible," "waste") detected; "quality" appears as a positive word. Distribution reflects strong negative polarity.

D. Neutral Sentiment Output Example

Fig. 3 illustrates how the system handles ambiguous or balanced reviews. A neutral classification is shown with a confidence score of 34%. No positive or negative words are detected, and the sentiment distribution bars show an almost even split across positive (33%), neutral (34%), and negative (33%) categories, demonstrating the system's ability to recognise the absence of strong sentiment cues.

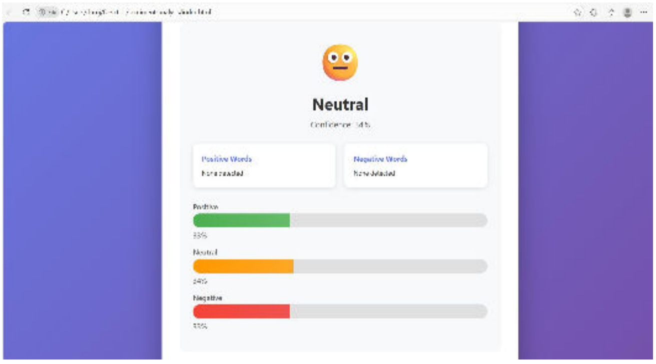

Fig. 3. Output showing neutral classification with 34% confidence. No sentiment-bearing words detected; distribution is nearly balanced across all categories.

E. Interpretation of Visual Outputs

Together, Figs. 1–3 demonstrate the system's ability to identify sentiment-bearing words, provide transparent and interpretable classifications, display confidence levels, visualise sentiment distribution, and handle positive, negative, and neutral reviews consistently. These outputs highlight the tool's suitability for real-time review assessment and privacy-sensitive environments.

VI. FUTURE ENHANCEMENTS

Proposed improvements to the system include:

- Negation detection
- Expanded lexicons
- Multilingual support
- Intensity scoring
- Sarcasm detection
- Export functionality
- Batch processing

VII. CONCLUSIONS

This research introduces a lightweight, privacy-preserving sentiment analysis tool that operates entirely on the client side. While lexicon-based methods cannot match the accuracy of advanced machine learning models, they offer significant advantages in speed, transparency, and ease of deployment. The tool is well-suited for rapid sentiment assessment, educational use, and privacy-sensitive environments. Future work will focus on improving linguistic coverage and contextual understanding while preserving the system's simplicity and efficiency.

APPENDIX A: SENTIMENT LEXICONS

Positive Words: amazing, excellent, great, good, love, best, perfect, awesome, fantastic, wonderful, brilliant, outstanding, superb, recommend, happy, satisfied, quality, beautiful, impressive, worth, nice, pleased, delighted, exceptional, incredible, fabulous, terrific.

Negative Words: bad, terrible, awful, horrible, worst, poor, disappointing, waste, useless, broken, defective, cheap, hate, regret, disappointed, unhappy, unsatisfied, inferior, pathetic, garbage, junk, fail, failed, problem, issue, never, avoid, refund.

REFERENCES

[1] 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.

[2] P. D. Turney, "Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews," in *Proc. ACL,* 2002, pp. 417–424.

[3] M. Hu and B. Liu, "Mining and summarizing customer reviews," in *Proc. ACM SIGKDD,* 2004, pp. 168–177.

[4] B. Pang, L. Lee, and S. Vaithyanathan, "Thumbs up? Sentiment classification using machine learning techniques," in *Proc. EMNLP,* 2002, pp. 79–86.

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

[6] S. Kokolakis, "Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon," *Computers & Security,* vol. 64, pp. 122–134, 2017.

[7] A. Esuli and F. Sebastiani, "SentiWordNet: A publicly available lexical resource for opinion mining," in *Proc. LREC,* 2006, pp. 417–422.

[8] F. Å. Nielsen, "A new ANEW: Evaluation of a word list for sentiment analysis in microblogs," in *Proc. ESWC Workshop,* 2011, pp. 93–98.

[9] C. Hutto and E. Gilbert, "VADER: A parsimonious rule-based model for sentiment analysis of social media text," in *Proc. ICWSM,* vol. 8, no. 1, pp. 216–225, 2014.

[10] O. Appel, F. Chiclana, J. Carter, and H. Fujita, "A hybrid approach to the sentiment analysis problem at the sentence level," *Knowledge-Based Systems,* vol. 108, pp. 110–124, 2016.

[11] A. Pak and P. Paroubek, "Twitter as a corpus for sentiment analysis and opinion mining," in *Proc. LREC,* 2010, pp. 1320–1326.

[12] D. Tang, B. Qin, and T. Liu, "Learning sentiment-specific word embedding for Twitter sentiment classification," in *Proc. ACL,* 2015, pp. 1555–1565.

[13] J. Bollen, H. Mao, and X. Zeng, "Twitter mood predicts the stock market," *Journal of Computational Science,* vol. 2, no. 1, pp. 1–8, 2011.

[14] M. Wiegand, A. Balahur, B. Roth, D. Klakow, and A. Montoyo, "A survey on the role of negation in sentiment analysis," in *Proc. Workshop on Negation and Speculation in NLP,* 2010, pp. 60–68.

[15] D. Maynard and M. A. Greenwood, "Who cares about sarcastic tweets? Investigating the impact of sarcasm on sentiment analysis," in *Proc. LREC,* 2014, pp. 4238–4243.

[16] T. Wilson, J. Wiebe, and P. Hoffmann, "Recognizing contextual polarity in phrase-level sentiment analysis," in *Proc. HLT/EMNLP,* 2005, pp. 347–354.