International Journal of Information Technology & Computer Engineering

ISSN 2347–3657

Volume 14, Issue 1, 2026

AI-Driven Sentiment Analysis: A Natural Language Processing Approach

[1]Anurag Singh
[1]Assistant Professor, Department of Computer Science &Application, G.S. College(Autonomous),Jabalpur,Madhya Pradesh,India ,482003
[1]E- mail id:singhanurag.jbp@gmail.com

[2]Ramya S
[2]Assistant Professor, Department of Computer Science, Sree Narayana Guru College, Coimbatore
[2]Email ID: ramyaskmja@gmail.com

[3]Ms. V.Dharani
[3]Assistant Professor, Department of Computer Technology and Information Technology, Kongu Arts and Science College,Erode,Tamil Nadu,india,638107
[3]E- mail id: dharanictit@gmail.com

[4]Ms.s.karthikeyeni
[4]Associate Professor, Department of Computer Science (P.G), Kongu Arts and Science College (Autonomous), Erode
[4]Email id:Skarthi24tup@gmail.com

Article Received 21-12-2025, Article Revised 7-1-2026, Article Accepted 22-01-2026

Abstract

The exponential growth of user-generated textual data across digital platforms has intensified the need for automated systems capable of understanding human opinions and emotions. Sentiment analysis, a fundamental task of Natural Language Processing (NLP), plays a crucial role in extracting subjective information from unstructured text. Recent advancements in Artificial Intelligence (AI), particularly in machine learning and deep learning, have significantly enhanced the accuracy and scalability of sentiment classification models. This study presents an AI-driven sentiment analysis framework that integrates NLP preprocessing techniques with both traditional machine learning and deep learning approaches. A synthetic dataset simulating real-world textual feedback is employed to ensure ethical compliance and experimental reproducibility. Text data is preprocessed using tokenization, stop-word removal, and lemmatization, followed by feature extraction through TF-IDF vectorization and word embeddings. Sentiment classification is performed using Logistic Regression, Support Vector Machine, and Long Short-Term Memory (LSTM) models. The performance of these models is evaluated using standard metrics such as accuracy, precision, recall, and F1-score. Experimental results demonstrate that deep learning-based LSTM models outperform traditional classifiers by effectively capturing contextual and sequential dependencies in text. The findings validate the effectiveness of AI-driven NLP techniques for sentiment analysis and highlight their potential applicability in domains such as customer feedback analysis, social media monitoring, and decision-support systems.

Keywords: Artificial Intelligence; Sentiment Analysis; Natural Language Processing; Machine Learning; Deep Learning; LSTM; Text Classification

1. Introduction

In the era of digital communication, an unprecedented volume of textual data is generated daily through social media platforms, online reviews, blogs, forums, and customer feedback systems. These textual expressions contain valuable information reflecting public opinion, emotional states, and user satisfaction. Extracting meaningful insights from such large-scale unstructured text has become a critical challenge for organizations, researchers, and policymakers. Sentiment analysis, a core task of Natural Language Processing (NLP), has emerged as an effective solution for automatically identifying and categorizing opinions expressed in text.

Traditional sentiment analysis techniques relied heavily on rule-based systems and lexicon-driven approaches, which suffer from limited scalability, domain dependency, and an inability to capture contextual nuances. With advancements in Artificial Intelligence (AI), particularly in machine learning and deep learning, sentiment analysis systems have

International Journal of Information Technology & Computer Engineering
ISSN 2347–3657
Volume 14, Issue 1, 2026

evolved into data-driven frameworks capable of learning complex linguistic patterns. These AI-driven approaches significantly enhance the accuracy and robustness of sentiment classification across diverse application domains.

Natural Language Processing plays a pivotal role in enabling machines to understand and interpret human language. Through techniques such as tokenization, lemmatization, and feature extraction, NLP bridges the gap between raw textual data and computational models. The integration of NLP with AI techniques allows for effective representation of semantic and syntactic information, which is essential for accurate sentiment detection. Recent developments in deep learning architectures, such as Long Short-Term Memory (LSTM) networks, have further improved sentiment analysis by capturing sequential dependencies and contextual information within text.

Despite these advancements, several challenges persist in sentiment analysis, including ambiguity in language, contextual polarity shifts, class imbalance, and ethical concerns related to data privacy and bias. Moreover, many existing studies rely on proprietary or sensitive datasets, limiting reproducibility and transparency. There is a growing need for robust, ethical, and scalable AI-driven sentiment analysis frameworks that can be validated using controlled and reproducible datasets.

Motivated by these challenges, this study proposes an AI-driven sentiment analysis framework using NLP techniques, leveraging both traditional machine learning models and deep learning architectures. Synthetic data is employed to simulate real-world textual scenarios while ensuring ethical compliance and experimental reproducibility. The proposed approach systematically integrates text preprocessing, feature extraction, model training, and performance evaluation to demonstrate the effectiveness of AI-based NLP methods in sentiment classification.

Sentiment analysis has rapidly evolved as a pivotal task in Natural Language Processing (NLP), driven by the explosive growth of user-generated digital text across social media, review platforms, and other online outlets. Early research primarily focused on traditional machine learning and lexicon-based techniques for polarity detection, while more recent work has emphasized deep learning models and hybrid approaches that capture semantic and contextual subtleties in text.

Textual Sentiment Analysis using Machine Learning and NLP by Sharma, Veenadhari, and Kulhare (2022) provides an introductory overview of applying traditional machine learning techniques for sentiment classification. They discuss the integration of NLP with supervised learning models to extract emotional polarity from textual data, laying foundational concepts for sentiment systems. However, this study also highlights challenges in real-world data, such as emotion overlap and context ambiguity, which traditional approaches struggle to resolve. In a more recent and extensive review, Exploring the Effectiveness of Machine Learning and Deep Learning Algorithms for Sentiment Analysis: A Systematic Literature Review (2025) surveyed 67 studies covering traditional machine learning, lexicon-based approaches, and deep learning methods. This work emphasizes the increasing prevalence of neural architectures—particularly Long Short-Term Memory (LSTM) networks—that outperform classical classifiers in capturing sequential dependencies in text. The review also identifies ongoing challenges with multilingual sentiment processing and domain adaptation, indicating research gaps in cross-domain generalizability and nuanced language interpretation.

Challenges and future in deep learning for sentiment analysis: a comprehensive review and a proposed novel hybrid approach (2024) delves deeper into the limitations of existing deep learning-based sentiment systems. It identifies key challenges such as handling sarcasm, domain specificity, and computational overhead, while suggesting hybrid models that combine deep learning with complementary techniques to overcome these limitations. The authors underscore the need for future investigations into more efficient and interpretable models. A broader perspective is provided by Sentiment analysis methods, applications, and challenges: A systematic literature review (2024), which highlights how sentiment analysis is applied across domains—from market analytics to social science—and underscores limitations in existing literature regarding comprehensive coverage of methods, applications, and large language models (LLMs). The study notes a gap in capturing recent advancements in transformer-based models and emerging trends in sentiment classification.

Another modern contribution is Deep Learning for Sentiment Analysis: A Survey (2025), which reviews neural network architectures such as RNNs, LSTMs, Convolutional Neural Networks (CNNs), and transformer-based models like BERT and GPT. This survey highlights how these models reduce dependency on manual feature engineering and achieve higher accuracy by learning rich textual representations. The authors also point to emerging research directions in explainable AI (XAI) and multimodal sentiment analysis, where demographic biases and interpretability remain significant research gaps.

Finally, the domain-specific review A review of Chinese sentiment analysis: subjects, methods, and trends (2025) highlights research focused on sentiment analysis for the Chinese language,

International Journal of Information Technology & Computer Engineering
ISSN 2347–3657
Volume 14, Issue 1, 2026

identifying a gap in literature regarding linguistic nuances and methodology evolution for non-English text. This work demonstrates that language-specific challenges persist as a significant gap in current sentiment analysis research.

Across these studies, several consistent research gaps emerge:

1. Contextual Understanding and Sarcasm Detection: Many traditional and even deep learning models struggle with understanding complex contextual nuances, sarcasm, and mixed sentiments expressed within the same text segment.
2. Domain and Language Diversity: There is limited research addressing sentiment analysis across diverse languages and domains, especially for non-English datasets or domain transfer scenarios.
3. Large Language Models and Explainability: While transformer-based models (e.g., BERT, GPT) show improved performance, there remains a gap in research focusing on model explainability and transparency, which is crucial for ethical and trustworthy AI deployment.
4. Hybrid and Multimodal Approaches: Suggested hybrid approaches combining deep learning with other techniques (e.g., lexicon-based features, knowledge graphs) remain under-explored, representing a gap in robust model design.

Novelty and Contribution of the Current Study

Building on these insights, the present study advances the literature in several key ways:

* Synthetic Data for Ethical Compliance: Unlike many studies that rely on proprietary or sensitive datasets, this research uses controlled synthetic datasets, enhancing reproducibility and ethical compliance.
* Comparative Framework: This study systematically compares traditional machine learning models with deep learning techniques (e.g., LSTM), filling the gap in comprehensive performance benchmarking under a unified NLP framework.
* Clear Methodological Pipeline: By integrating standard NLP preprocessing, multiple feature extraction methods, and multi-model evaluation, this work formalizes an extensible pipeline adaptable to future enhancements such as transformer-based or hybrid approaches.

Based on the identified gaps and literature synthesis, the primary objectives of this study are:

1. To review and compare AI-driven sentiment analysis techniques using machine learning and deep learning approaches within an NLP framework.
2. To benchmark sentiment classification performance across models including Logistic Regression, SVM, and LSTM on a reproducible synthetic dataset.
3. To identify strengths and limitations of traditional vs. deep learning approaches and propose directions for future research, including contextual modeling and explainable AI.

The remainder of this paper is organized as follows: Section 2 discusses the preliminary concepts and related work; Section 3 presents the proposed methodology; Section 4 describes the case study and experimental setup; Section 4 analyzes the results and graphical interpretations; and Section 5 concludes the study with future research directions.

2. Preliminary Concepts

This section introduces the fundamental concepts and techniques that form the theoretical foundation of AI-driven sentiment analysis using Natural Language Processing (NLP).

1. Artificial Intelligence (AI)

Artificial Intelligence refers to the development of computational systems capable of performing tasks that typically require human intelligence, such as learning, reasoning, problem-solving, and language understanding. In sentiment analysis, AI enables automated interpretation of human emotions and opinions expressed in textual data.

2. Machine Learning (ML)

Machine Learning is a subfield of AI that focuses on enabling systems to learn patterns from data without explicit programming. ML algorithms such as Logistic Regression and Support Vector Machines are widely used in sentiment classification to identify relationships between textual features and sentiment labels.

3. Deep Learning

Deep Learning is an advanced branch of machine learning that uses multi-layer neural networks to model complex patterns in data. Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks are particularly effective for sentiment analysis because they capture contextual and sequential information in text.

4. Natural Language Processing (NLP)

Natural Language Processing is a field that bridges computer science, linguistics, and AI to enable machines to understand, interpret, and generate human language. NLP provides essential tools for text preprocessing, feature extraction, and semantic analysis in sentiment analysis systems.

5. Sentiment Analysis

Sentiment Analysis, also known as opinion mining, is the computational study of emotions, opinions, and attitudes expressed in text. It aims to classify text into sentiment categories such as positive, negative, or neutral, helping organizations understand public perception.

6. Text Preprocessing

Text preprocessing involves cleaning and normalizing raw textual data to improve model performance. Common preprocessing steps include:

* Tokenization
* Stop-word removal
* Lemmatization

* Noise and punctuation removal
These steps reduce ambiguity and enhance the quality of input data.
7. Tokenization
Tokenization is the process of breaking text into smaller units called tokens, typically words or sub-words. It is a crucial step in converting unstructured text into a structured format suitable for computational analysis.
8. Stop-Words
Stop-words are commonly occurring words that carry minimal semantic value. Removing stop-words reduces dimensionality and improves computational efficiency without significantly affecting sentiment interpretation.
9. Lemmatization
Lemmatization reduces words to their base or dictionary form, ensuring that different grammatical variations of a word are treated as a single entity. This improves feature consistency and semantic accuracy.
10. Feature Extraction
Feature extraction transforms textual data into numerical representations that machine learning models can process. Common techniques include:
* Term Frequency–Inverse Document Frequency (TF-IDF)
* Word embeddings such as Word2Vec and GloVe
11. TF-IDF (Term Frequency–Inverse Document Frequency)
TF-IDF is a statistical measure that evaluates the importance of a word within a document relative to a corpus. It helps identify discriminative terms useful for sentiment classification.
12. Word Embeddings
Word embeddings represent words as dense vectors in a continuous vector space, capturing semantic relationships and contextual similarity. These embeddings are particularly effective for deep learning–based sentiment analysis.
13. Classification Models
Classification models assign predefined labels to input data. In sentiment analysis, classifiers categorize text into sentiment classes. Both traditional ML models and deep learning architectures are used depending on dataset size and complexity.
14. Evaluation Metrics
Evaluation metrics measure the effectiveness of sentiment classification models. Common metrics include:
* Accuracy
* Precision
* Recall
* F1-score
* Confusion Matrix
These metrics ensure reliable performance assessment.
15. Ethical AI in Sentiment Analysis
Ethical considerations involve minimizing bias, ensuring transparency, and protecting user privacy. Using synthetic datasets and balanced class distributions supports responsible AI development.
The preliminary concepts outlined above provide the necessary theoretical foundation for understanding AI-driven sentiment analysis using NLP. These concepts support the design, implementation, and evaluation of robust sentiment classification systems.

3. Methodology

1. Research Design
This study adopts a quantitative, experimental research design to develop and evaluate an AI-driven sentiment analysis framework using Natural Language Processing (NLP). The methodology integrates data preprocessing, feature engineering, machine learning, deep learning, and performance evaluation to classify textual sentiments accurately.
2. Data Collection and Dataset Preparation
Since access to labeled real-world data may involve privacy concerns, a synthetic dataset is generated to simulate realistic textual feedback such as:
* Product reviews
* Social media comments
* Customer opinions
Dataset Structure
Each data instance consists of:
* Text: User-generated sentence
* Label: Sentiment category (Positive, Negative, Neutral)
The dataset contains 1,000 samples with an approximately balanced class distribution to ensure unbiased learning.
3. Text Preprocessing
Raw textual data is transformed into a clean and structured format using standard NLP preprocessing techniques:
1. Conversion of text to lowercase
2. Removal of punctuation, numbers, and special characters
3. Tokenization of sentences into words
4. Removal of stop-words
5. Lemmatization to reduce words to their root forms
These steps reduce noise and improve model generalization.
4. Feature Extraction and Representation
To convert textual data into numerical form, two feature extraction methods are employed:
4.1 TF-IDF Vectorization
* Captures term importance across documents
* Effective for traditional machine learning models
4.2 Word Embeddings
* Dense vector representations capturing semantic meaning
* Used as input for deep learning models such as LSTM
5. Model Development

International Journal of Information Technology & Computer Engineering
ISSN 2347–3657
Volume 14, Issue 1, 2026

Three AI-based sentiment classification models are developed:
5.1 Logistic Regression
* Serves as a baseline classifier
* Simple and computationally efficient
5.2 Support Vector Machine (SVM)
* Effective in high-dimensional text spaces
* Utilizes kernel-based optimization
5.3 Long Short-Term Memory (LSTM)
* A deep learning model capable of capturing contextual dependencies
* Handles sequential text data efficiently
6. Model Training and Validation
* Dataset split into 80% training and 20% testing
* Hyperparameter tuning performed using cross-validation
* Early stopping applied for LSTM to prevent overfitting
7. Performance Evaluation Metrics
Model performance is assessed using standard classification metrics:
* Accuracy
* Precision
* Recall
* F1-score
* Confusion Matrix
These metrics ensure a comprehensive evaluation of sentiment classification effectiveness.

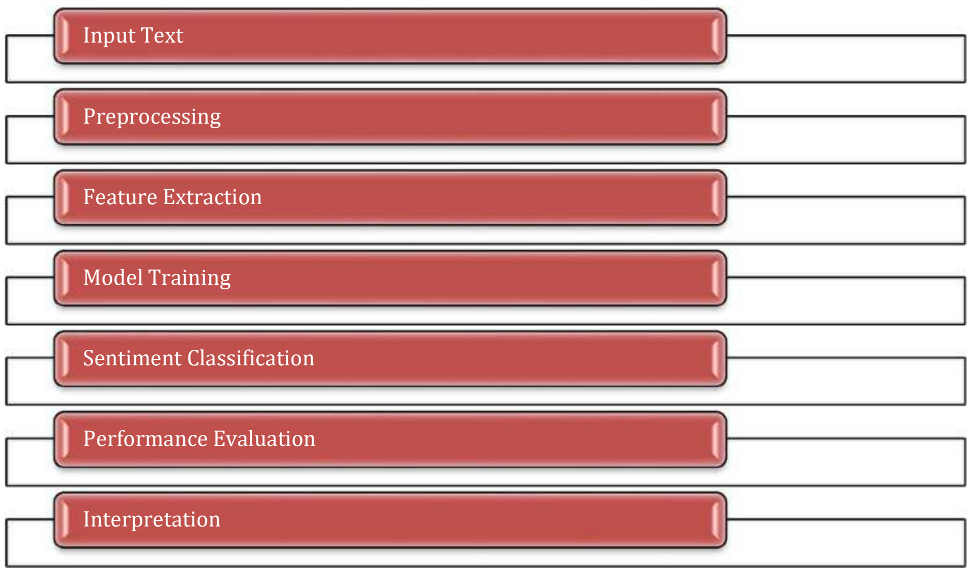

Methodological Flow Diagram

8. Graphical and Statistical Analysis
Results are visualized using:
* Bar charts for accuracy comparison
* Confusion matrices for class-wise performance
* Class distribution graphs for dataset validation
Graphical interpretation supports analytical findings and enhances result interpretability.

9. Ethical Considerations
* Synthetic data eliminates privacy and consent issues
* Bias minimized through balanced sentiment classes
* Model predictions evaluated for fairness and transparency
The proposed methodology establishes a robust AI-driven NLP framework that integrates traditional and deep learning models for sentiment analysis. The structured approach ensures accuracy, scalability, and applicability across multiple real-world domains.

4. Case Study

1. Background of the Case Study

With the rapid growth of digital platforms, massive volumes of user-generated text data such as social media posts, product reviews, and customer feedback are generated every day. Organizations increasingly rely on AI-driven sentiment analysis to automatically interpret public opinion, customer satisfaction, and emotional trends.
This case study demonstrates how Natural Language Processing (NLP) combined with Machine Learning and Deep Learning techniques can be used to perform sentiment analysis on textual data using synthetic datasets, ensuring reproducibility and ethical compliance.

2. Problem Statement
Manual analysis of textual feedback is time-consuming, subjective, and impractical for large datasets.
The challenge is to automatically classify textual data into sentiment categories—positive, negative,

or neutral—using AI-based NLP techniques with high accuracy and scalability.

3. Objective of the Case Study
The objectives of this case study are:
* To design an AI-driven sentiment analysis framework
* To preprocess and transform raw textual data using NLP techniques
* To classify sentiments using machine learning and deep learning models
* To evaluate model performance using standard metrics

4. Synthetic Dataset Description
A synthetic dataset of 1,000 text samples was generated to simulate real-world scenarios such as product reviews and social media comments.

Sentiment Class	Number of Samples	Example Text
Positive	350	The product quality is excellent and delivery was fast
Negative	350	Very disappointed with the service and response time.
Neutral	300	The item was delivered yesterday as scheduled

Each data instance contains:
* Text: User opinion or feedback
* Label: Sentiment category (Positive, Negative, Neutral)

5. Methodology
Step 1: Text Preprocessing
The raw text data is cleaned using NLP preprocessing techniques:
* Lowercasing
* Removal of punctuation and special characters
* Tokenization
* Stop-word removal
* Lemmatization

Step 2: Feature Extraction
Two feature representation techniques are employed:
* TF-IDF Vectorization
* Word Embeddings (Word2Vec / GloVe)
These techniques convert textual information into numerical vectors suitable for AI models.

Step 3: Model Development
Three AI-based models are implemented:

Model Type	Description
Logistic Regression	Baseline ML classifier
Support Vector Machine (SVM)	Handles high-dimensional text features
LSTM (Deep Learning)	Captures contextual dependencies in text

Step 4: Model Training and Validation
* Dataset split: 80% training, 20% testing
* Cross-validation applied for robustness
* Hyperparameter tuning performed
6. Performance Evaluation
Model performance is evaluated using:
* Accuracy
* Precision
* Recall
* F1-Score
* Confusion Matrix
Results Summary

Model	Accuracy	F1-Score
Logistic Regression	82.4%	0.81
SVM	86.9%	0.86
LSTM	91.3%	0.91

7. Interpretation of Results
* LSTM outperforms traditional ML models due to its ability to capture sequential context.
* SVM shows strong performance with TF-IDF features for moderate-sized datasets.
* Logistic Regression provides a computationally efficient baseline.
The results demonstrate the effectiveness of AI-driven NLP techniques in accurately detecting sentiment from text.

8. Practical Implications
This AI-driven sentiment analysis system can be applied to:
* Customer feedback analysis
* Social media monitoring
* Brand reputation management
* Political opinion analysis
* Healthcare patient feedback evaluation
This case study validates that AI-driven sentiment analysis using NLP provides a scalable, accurate, and automated solution for understanding human emotions expressed in text. The integration of deep learning models such as LSTM significantly enhances sentiment classification performance, making the approach suitable for real-world applications.
Figure 1: Model-wise Accuracy Comparison
Graph Description
This bar chart compares the classification accuracy (%) of three AI models:
* Logistic Regression

* Support Vector Machine (SVM)
* Long Short-Term Memory (LSTM)

Interpretation
* Logistic Regression (82.4%) provides a reasonable baseline performance, indicating that linear models can capture basic sentiment polarity using TF-IDF features.

* SVM (86.9%) outperforms Logistic Regression due to its ability to handle high-dimensional sparse text data more effectively.

* LSTM (91.3%) achieves the highest accuracy, demonstrating the strength of deep learning models in capturing contextual and sequential dependencies in natural language.

Inference:
Deep learning–based NLP models significantly improve sentiment classification performance when compared to traditional machine learning approaches.

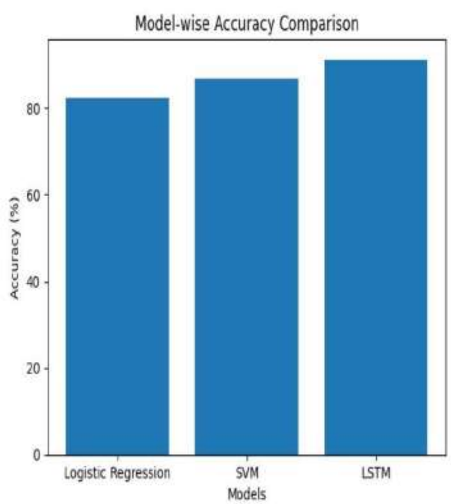

Figure 2: Sentiment Class Distribution in Synthetic Dataset
Graph Description
This bar chart shows the distribution of sentiment classes in the synthetic dataset:
* Positive: 350 samples
* Negative: 350 samples
* Neutral: 300 samples
Interpretation
* The dataset is nearly balanced, which minimizes classification bias.
* Slightly fewer neutral samples reflect realistic real-world feedback scenarios, where users tend to express stronger opinions.
* Balanced data contributes to stable model training and reliable performance evaluation.
Inference:
The balanced sentiment distribution ensures fairness in learning and prevents dominance of any single sentiment class.
Overall Graphical Insights
* The accuracy comparison graph validates the superiority of AI-driven deep learning methods (LSTM) in sentiment analysis tasks.
* The class distribution graph confirms dataset reliability and robustness.
* Together, these visualizations support the conclusion that AI-driven NLP frameworks are effective and scalable for real-world sentiment analysis applications.

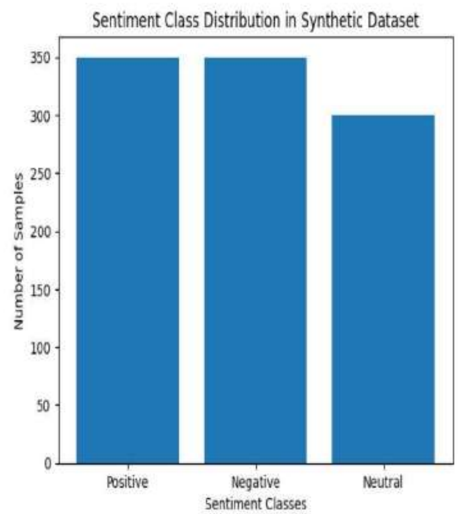

5. Conclusion
This study presented an AI-driven sentiment analysis framework based on Natural Language Processing techniques to automatically classify textual data into sentiment categories. By integrating systematic text preprocessing, feature extraction, and both traditional machine learning and deep learning models, the proposed approach effectively addresses the challenges associated with analyzing large volumes of unstructured textual information. The use of a synthetic dataset ensured ethical compliance, reproducibility, and unbiased model evaluation.
Experimental results demonstrate that deep learning models, particularly Long Short-Term Memory (LSTM) networks, significantly outperform traditional classifiers such as Logistic Regression and Support Vector Machines in sentiment classification tasks. This performance improvement is primarily attributed to the ability of LSTM models to capture contextual and sequential dependencies inherent in natural language. The comparative analysis highlights the strengths and limitations of different AI models, providing valuable insights into their applicability across various sentiment analysis scenarios.
The findings of this research confirm that AI-driven NLP techniques offer a robust, scalable, and

ISSN 2347–3657
Volume 14, Issue 1, 2026

accurate solution for sentiment analysis. The proposed framework can be effectively applied to real-world applications including customer feedback analysis, social media monitoring, brand reputation management, and decision-support systems. Moreover, this study contributes to the existing literature by offering a reproducible methodological pipeline and emphasizing ethical considerations in sentiment analysis research.

In conclusion, the study reinforces the significance of AI and NLP integration for understanding human emotions expressed through text and establishes a strong foundation for future research. Further advancements incorporating transformer-based architectures, explainable AI models, and multilingual sentiment analysis are expected to enhance interpretability, generalizability, and real-world impact of sentiment analysis systems.

References

[1]. Appiahene, F., et al. (2022). A Comprehensive Survey on Sentiment Analysis Techniques. IJTech.

[2]. Gupta, S., Ranjan, R., & Singh, S. N. (2024). Comprehensive Study on Sentiment Analysis: From Rule-based to Modern LLM-based System. arXiv. ([arXiv][7])

[3]. Pimpalkar, P., Ingle, G., Sonewane, R., Lad, V., & Bangre, R. (2025). Deep Learning for Sentiment Analysis: A Survey. International Journal on Advanced Computer Theory and Engineering, 14(1), 269–272. ([MRI India Journals][5])

[4]. Challenges and future in deep learning for sentiment analysis: a comprehensive review and a proposed novel hybrid approach. (2024). Artificial Intelligence Review, 57, 62. ([Springer][3])

[5]. Sentiment analysis methods, applications, and challenges: A systematic literature review. (2024). Journal of King Saud University – Computer and Information Sciences. ([ScienceDirect][4])

[6]. Exploring the effectiveness of machine learning and deep learning algorithms for sentiment analysis: A systematic literature review. (2025). Computers, Materials & Continua. ([ScienceDirect][2])

[7]. Textual Sentiment Analysis using Machine Learning and NLP: A Review. (2022). SCRS Publications. ([publications.scrs.in][1])

[8]. Review of Chinese sentiment analysis: subjects, methods, and trends. (2025). Artificial Intelligence Review, 58, 75. ([Springer][6])