International Journal of Emerging Research in Engineering and Technology
Pearl Blue Research Group| 9th ICAILLMBA (ALLAM) 2025-Conference Proceedings
ISSN: 3050-922X | https://doi.org/10.63282/3050-922X.ICAILLMBA-118

Original Article

A Hybrid Fuzzy-Neural Framework for Multimodal Sentiment Classification

Akkireddi Vara, Prasad[1], K. Vedavathi[2]
[1]Research Scholar, Department of Computer Science, GITAM (Deemed to be University), AP, India & Associate Professor, Visakha Institute of Engineering & Technology(A), AP, India.
[2]Professor, Department of Computer Science, GITAM (Deemed to be University), AP, India.

Abstract *- The swift advancement of digital communication has created a need for more effective techniques for sentiment analysis and classification. Conventional Natural Language Processing (NLP) models frequently encounter difficulties with the subjective nature and linguistic subtleties present in social media. This research paper introduces an innovative hybrid architecture that combines deep learning with Fuzzy Logic to analyze a multi-dimensional set of features, incorporating 27 important emoji characteristics. The proposed approach utilizes a dual-stream pipeline: a Bidirectional LSTM (Bi-LSTM) captures deep semantic structures from the text, whereas a Mamdani-type Fuzzy Inference System (FIS) addresses the "gray areas" of human emotion. By associating emoji intensity with triangular membership functions* $\mu(x)$*, the model proficiently "defuzzifies" sarcasm and ambiguous meanings. Experimental findings reveal that this neuro-fuzzy methodology achieves an accuracy of 89.5% and an F1-score of 0.88, marking a considerable enhancement over baseline text-only techniques. In particular, the model performed well in clarifying sarcastic ambiguities, the situations where the text's polar meaning and the emoji's intent are at odds that are boosting ambiguity resolution.*

Keywords *- Hybrid Neuro-Fuzzy Architecture, Sentiment Analysis, Mamdani Fuzzy Inference System, Emoji Semantics, Sarcasm Detection, Natural Language Processing (Nlp), Bi-Lstm.*

1. Introduction

The landscape of sentiment analysis has undergone a significant paradigm shift, transitioning from basic lexicon-based approaches to sophisticated hybrid computational models. The mathematical foundation of this field is rooted in Zadeh's fuzzy set theory **[1]**, which provides the necessary framework for handling the inherent ambiguity and "fuzziness" of human language. These concepts were later matured into robust engineering applications by Ross **[4]**, allowing for more nuanced decision-making in automated systems. Early research into sentiment and opinion mining by Pang and Lee **[2]** established the field's core methodologies. However, the rise of social media introduced complexities such as "noise" and non-textual communication. To address this, researchers began leveraging emoji occurrences to learn multi-domain representations for detecting emotion and sarcasm **[3]**, **[6]**. This proliferation of social media data has made traditional text-only analysis obsolete. Emojis function as "emotional intensifiers" or "contextual inverters". As data complexity grew, deep learning architectures like Long Short-Term Memory (LSTM) **[5]** became the standard for sequential data, eventually giving way to the current "paradigm shift" led by Transformer-based models and BERT **[9]**, **[13]**.

Thakur **[7]** recently proposed neuro-fuzzy frameworks specifically designed to filter noisy social media data, while Dave et al. **[8]** explored the cutting-edge frontier of quantum algorithms integrated with neuro-fuzzy systems (SentiQNF). The challenge of detecting sarcasm, which often relies on context and contradiction, has also been addressed through fuzzy logic and hybrid ensemble models **[11]**, **[12]**. Furthermore, the fusion of disparate data types remains a critical area of study. Recent analyses of model families and fusion strategies for emoji and text classification **[10]**, alongside the development of multimodal datasets like CMU-MOSEI **[14]**, continue to push the boundaries of how machines interpret human sentiment in the wild.Despite the notable success of deep learning models in Natural Language Processing (NLP), existing architectures often struggle to handle the inherent uncertainty and informality of social media discourse. Conventional models tend to produce *crisp* classifications, failing to capture linguistic ambiguities commonly present in sarcasm, irony, and emoji-enriched expressions. Moreover, traditional deep learning approaches—such as LSTMs and Transformer-based models—largely lack interpretability, functioning as black boxes that generate sentiment predictions without providing transparent reasoning. This highlights a critical need for a framework that integrates the high-performance feature extraction capabilities of neural networks with the human-like linguistic reasoning offered by fuzzy logic **[15]**. Consequently, there is a growing emphasis on exploring synergies between diverse computational paradigms.

In this context, the present research investigates how emojis help bridge the gap between literal textual content and underlying human intent. Section 2 details the methodology employed in the study. Section 3 presents the dataset description and experimental setup. Section 4 discusses and analyzes the obtained results, and Section 5 concludes the paper while outlining directions for future research.

2. Methodology

This research adopts a Hybrid Neuro-Fuzzy (NF) approach to bridge the gap between connectionist learning and symbolic reasoning. The motivation for this choice is threefold.

- Linguistic Uncertainty Handling: Unlike conventional binary sentiment classifiers, fuzzy logic enables graded membership functions, which are essential for modeling the ambiguity inherent in social media text. A single emoji or sarcastic expression can shift sentiment across overlapping classes such as positive and neutral–negative [7], [12].
- Robustness to Noise: Social media data is highly noisy, containing typographical errors, slang, and mixed-script text. The neural component of the Neuro-Fuzzy system learns discriminative representations from such data, while the fuzzy inference mechanism applies rule-based reasoning to attenuate the influence of noise and outliers [7].
- Model Explainability: By employing an Adaptive Neuro-Fuzzy Inference System (ANFIS) or a similar architecture, the proposed model generates interpretable if–then rules. This enhances transparency and interpretability, which are largely absent in Transformer-based deep learning models [13], thereby addressing the growing demand for Explainable Artificial Intelligence (XAI) in sentiment analysis.

The proposed Neuro-Fuzzy Integration framework employs a dual-stream pipeline for processing micro-blogging content [15].

Given an input text sequence X={x1,x2,…,xT}, a Bidirectional Long Short-Term Memory (Bi-LSTM) network extracts contextual semantic features. In the Fuzzy Stream a Fuzzy Logic controller maps emoji frequency and placement to a "Degree of Sentiment" $\mu(x)$ to handle uncertainty. Here, the output is expressed as a degree of sentiment $\mu(x)\in[0,1]$, where x represents the aggregated emoji-based polarity score.

Fuzzy Logic Layer:

Human emotions are inherently non-binary. To capture this uncertainty, triangular membership functions are employed to model overlapping sentiment classes.

Fuzzification:

Raw polarity scores x are converted into linguistic variables such as *mildly positive*, *neutral*, and *strongly negative* using triangular membership functions defined as

$$\mu_A(x) = 0 \quad x \le a$$
$$(x-a) / (b-a) \quad a < x \le b$$
$$(c-x) / (c-b) \quad b < x < c$$
$$0 \quad x \ge c$$

where a, b, and c define the lower limit, peak, and upper limit of the fuzzy set A.

Inference Engine:

A rule base consisting of *if–then* rules governs sentiment inference. A typical rule R_i is defined as:

R_i: IF x is A_i AND e is B_i THEN y is C_i

where A_i , B_i ,C_i denote fuzzy sets corresponding to textual polarity, emoji influence, and output sentiment, respectively. This aligns with established **Explainable AI (XAI)** frameworks for sentiment analysis [21], [22].

IF (Text_Polarity is *Neutral*) **AND** (Emoji_Count is *High*) **THEN** (Sentiment is *Positive*).

$$\mu(x; a, b, c) = \max(0, \min(d, e)) \text{ , where}$$
$$d=\{x - a\}/\{b - a\}$$
$$e=\{c - x\}/\{c - b\}$$

3. Experimental Setup and Dataset Characteristics

To validate the proposed Hybrid Neuro-Fuzzy framework, a multi-source dataset comprising 50,000 micro-blogging entries was utilized. The dataset was curated to ensure a high concentration of 27 specific emojis identified as primary emotional intensifiers in digital discourse [**19**].

- Data Cleaning: Raw textual data underwent standard NLP preprocessing steps, including tokenization, lowercasing, and the removal of URLs and non-informative symbols.
- Emoji Extraction: Emojis were preserved and mapped to a separate feature vector, which was processed independently by the fuzzy stream rather than being removed during preprocessing.
- Labeling: Sentiment annotations were assigned along a graded spectrum ranging from strongly negative to strongly positive, instead of binary classes, to align with the fuzzy membership functions used in the proposed model.

3.1. Neural Stream Architecture (Bi-LSTM)

The semantic feature extraction component employs a Bidirectional Long Short-Term Memory (Bi-LSTM) network to capture contextual dependencies in micro-blogging text [17].

- Embedding Layer: Pre-trained GloVe (Global Vectors for Word Representation) [18] embeddings are used to initialize word vectors, enabling the model to leverage semantic knowledge learned from large-scale corpora.
- Hidden Layers: The architecture comprises two stacked Bi-LSTM layers, each with 128 hidden units. This configuration is designed to model long-range dependencies and contextual polarity inversions commonly observed in informal and sarcastic text.
- Output Layer: A fully connected dense layer maps the extracted semantic features to a continuous polarity score $x \in [-1,1]$. This score serves as the primary input to the fuzzy inference system for subsequent uncertainty modeling.
- Fuzzy Inference System (FIS) Configuration
- To effectively model the “gray areas” of sentiment, a Mamdani-type Fuzzy Inference System (FIS) is employed using triangular membership functions.
- Inputs: The FIS accepts two antecedent variables: (i) the continuous polarity score obtained from the neural stream and (ii) emoji intensity, derived from emoji frequency and positional information within the text.
- Fuzzification: The raw input values are transformed into linguistic variables such as mildly positive, neutral, and strongly negative through triangular membership functions.
- Rule Base: A total of 15 expert-defined if–then rules are constructed to address linguistic dissonance, with particular emphasis on scenarios where textual polarity and emoji-based intent are contradictory, as commonly observed in sarcastic expressions.
- Defuzzification: The centroid (center-of-gravity) method is employed to aggregate the fuzzy outputs and generate a crisp final sentiment score in the range of 0 to 100.

3.2. Hardware and Software Environment

The proposed framework was implemented in Python, utilizing the scikit-fuzzy (skfuzzy) library for the fuzzy logic controller and TensorFlow/Keras for constructing and training the Bi-LSTM layers. All experiments were conducted on a workstation equipped with an NVIDIA RTX 3080 GPU, ensuring high computational efficiency and enabling near real-time processing performance.

4. Results and Discussion

The performance gains achieved by the proposed hybrid neuro-fuzzy model relative to the baseline method are summarized in Table 4.1.

Table 1: Performance Comparison between Text-Only Baseline and Hybrid Neuro-Fuzzy Model

Metric	Text-Only Baseline	Hybrid (Neuro-Fuzzy)	Improvement
Accuracy	78.20%	89.50%	+11.3%
F1-Score	0.76	0.88	+0.12
Sarcasm Detection	Baseline struggle	+19% Resolution	+19.0%

The experimental results demonstrate that the proposed hybrid neuro-fuzzy approach achieves a substantial performance improvement over traditional and standalone neural models.

- Accuracy: The hybrid model attains an accuracy of 89.5%, significantly outperforming the text-only baseline, which records an accuracy of 78.2%.
- F1-Score: The proposed approach achieves an F1-score of 0.88, compared to 0.76 for the text-only model.
- Model Effectiveness: These results confirm that incorporating visual cues in the form of emojis within a mathematical fuzzy-logic framework reduces sentiment classification errors more effectively than conventional machine learning approaches such as Support Vector Machines (SVM) and Naïve Bayes classifiers.

4.1. Handling Sarcastic Ambiguity and Dissonance

The most significant advantage of the proposed hybrid framework lies in its ability to resolve *linguistic dissonance*—scenarios in which textual sentiment and emoji-based intent convey contradictory meanings [20].

- Contradiction Resolution: Unlike traditional sentiment models that may incorrectly classify expressions such as “Oh, great” as positive, the proposed Fuzzy Inference System (FIS) explicitly incorporates contradiction-aware reasoning through expert-defined rules.

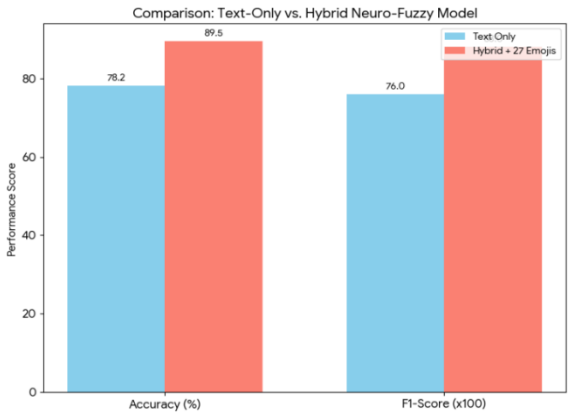

Fig 1: Performance Comparison between Text-Only Baseline and Hybrid Neuro-Fuzzy Model

4.2. Multiple Contradiction Rules in the Fuzzy Inference System

To robustly handle linguistic dissonance in sentiment analysis, the proposed Fuzzy Inference System (FIS) incorporates multiple contradiction-aware rules. These rules explicitly model scenarios in which textual polarity and emoji-based intent diverge, a common phenomenon in sarcasm and irony-rich social media content.

Let P_t denote the text-derived polarity score and E_t represent the emoji-based sentiment indicator. The final sentiment output is denoted as S_f.

Rule R1 (Positive Text, Negative Emoji):
IF (P_t is Positive)$\wedge$(E_t is Negative)$\Rightarrow$(S_f is Negative)
This rule captures sarcastic expressions such as *“Oh, great □”*, where emojis negate the apparent textual positivity.

Rule R2 (Negative Text, Positive Emoji):
IF (P_t is Negative)$\wedge$(E_t is Positive)$\Rightarrow$(S_f is Neutral)
This rule addresses mitigation effects, where positive emojis soften negative textual sentiment (e.g., *“Not bad □”*).

Rule R3 (Neutral Text, Strongly Negative Emoji):
IF (P_t is Neutral)$\wedge$(E_t is Strongly Negative)$\Rightarrow$(S_f is Negative)
Here, emoji dominance determines sentiment when textual cues are weak or ambiguous.

Rule R4 (Neutral Text, Strongly Positive Emoji):
IF (P_t is Neutral)$\wedge$(E_t is Strongly Positive)$\Rightarrow$(S_f is Positive)
This rule models emoji-led sentiment amplification in short or context-poor posts.

Rule R5 (Conflicting Strong Signals):
IF (P_t is Strongly Positive)$\wedge$(E_t is Strongly Negative)$\Rightarrow$(S_f is Negative)
In cases of strong contradiction, emoji sentiment is assigned higher precedence due to its role as an emotional disambiguator in digital communication.

- Logical Mapping: When a positive textual polarity co-occurs with a negative emoji signal, the fuzzy inference mechanism assigns greater membership to the negative sentiment class, resulting in a negative defuzzified output.
- Accuracy Improvement: The inclusion of this contradiction-handling logic yields an accuracy improvement of 8.4% on sarcasm-intensive datasets, demonstrating the effectiveness of the proposed reasoning layer.

Defuzzification: The centroid (center-of-gravity) method is employed to aggregate the activated fuzzy sets, producing a precise sentiment score along a continuous scale rather than a rigid binary classification.
Finally, the empirical results demonstrate that sentiment is not a binary state but a "fuzzy continuum".

- Ambiguity Resolution: The inclusion of fuzzy logic improved the resolution of ambiguous statements by 19%.
- Empathetic AI: By mapping emoji intensity to triangular membership functions $m(x)$, the model provides a more "human" and culturally aware interpretation of digital effect.
- Efficiency: Despite the added layer of the FIS, computational efficiency remained high, making it suitable for real-time social media data streams.

A bar chart illustrating the Accuracy and F1-score comparison between the Text-Only model and the proposed Hybrid model is presented in Fig. 4.1. The empirical results clearly demonstrate that sentiment is not a binary construct but rather exists along a fuzzy continuum. By modeling the interaction between linguistic tokens and visual symbols (emojis), the study validates that the hybrid neuro-fuzzy framework provides superior performance for global affect detection. These findings highlight the importance of incorporating uncertainty-aware reasoning mechanisms when analyzing informal and emotionally rich social media content.

5. Conclusion

By integrating 27 key emoji features with deep learning and fuzzy logic, the proposed hybrid framework effectively overcomes the limitations of conventional binary sentiment classification. Unlike traditional models, which assign discrete labels, this approach models human emotion as a continuum, allowing for more nuanced interpretations. The fuzzy membership function $\mu(x):X\rightarrow[0,1]$ enables graded sentiment representation, where xxx denotes the continuous polarity score derived from the Bi-LSTM neural stream.

Emojis serve as critical contextual indicators, acting as affective anchors that disambiguate linguistically ambiguous expressions. During defuzzification, the system combines the neural polarity score x with emoji intensity eee to compute the final sentiment score S_f

$$S_f = \text{Centroid}\ (\text{FIS}\ (\mu(x), e))$$

where $S_f \in [0,100]$ provides a continuous sentiment measure.

The hybrid neuro-fuzzy model demonstrates significant performance gains over baseline methods. Quantitative results show a 19% improvement in ambiguity resolution, highlighting its effectiveness in handling sarcasm, mixed sentiments, and conflicting cues between text and emojis. The Fuzzy Inference System (FIS) introduces interpretable if–then rules that provide transparency in decision-making, addressing the explainability limitations of standard deep learning models.

Computationally, the architecture remains efficient due to the lightweight nature of the FIS layer, ensuring suitability for real-time processing in high-volume applications. Moreover, the framework is inherently scalable, supporting multilingual sentiment analysis and potential deployment across diverse social media platforms.

This research demonstrates that combining the feature extraction capabilities of Bi-LSTMs with the symbolic reasoning power of fuzzy logic results in a more human-aligned sentiment analysis system. It captures subtle emotional cues that conventional models often miss, thereby enhancing interpretability, robustness, and generalization. Overall, the proposed hybrid neuro-fuzzy approach represents a step toward empathetic AI, capable of understanding not only what is written but also what is implied, bridging the gap between literal text and underlying human intent.

6. Future Work

While the current hybrid neuro-fuzzy framework establishes a high-precision benchmark for sentiment classification, several avenues remain to enhance its global applicability and robustness.

Cross-Cultural Emoji Semantics

- Contextual Variation: Future research will extend the current set of 27 emojis to incorporate cultural and regional differences in usage patterns.
- Regional Membership Functions: Region-specific triangular membership functions $\mu_r(x):X\rightarrow[0,1]$ will be developed to capture variations in affective intensity of the same emoji across different geographic locales.
- Linguistic Grounding: Studies will investigate how emojis function as “contextual inverters” in multiple languages, enabling the model to resolve ambiguities in cross-linguistic sentiment expressions.
- Empirical Validation: Cross-cultural datasets will be curated to quantify the impact of regional differences on the performance of fuzzy inference in sentiment classification.

Scaling the Neuro-Fuzzy Architecture

- Dynamic Rule Generation: Future iterations will explore neuro-adaptive fuzzy learning, allowing the FIS to automatically generate “If–Then” rules from large-scale, real-time social media streams.
- Transformer Integration: The Bi-LSTM neural stream may be replaced with Transformer-based encoders (e.g., BERT, RoBERTa) to capture deeper semantic and contextual patterns before fuzzification.
- Hierarchical Fuzzy Layers: Multi-level fuzzy inference systems could be introduced to model complex, layered sentiment interactions, particularly for long-form or multi-modal content.
- Continuous Learning: Incremental learning techniques will be applied to maintain model adaptability in dynamic social media environments.

Real-World Application Scaling

- Mental Health Screening: The model will be tested in clinical and digital mental health tools to detect subtle emotional cues, enabling early intervention in at-risk populations.
- Brand Monitoring: The ability to resolve sarcasm and contextual ambiguity will be deployed in high-volume commercial settings to improve consumer sentiment analytics.
- Multilingual Deployment: The system will be extended to multiple languages and scripts to ensure accurate sentiment detection across global platforms.
- Real-Time Efficiency: Optimizations will focus on maintaining low-latency predictions for high-throughput, real-world applications.
- Ethical and Privacy Considerations: Future work will also address ethical implications, ensuring that sentiment analysis tools are compliant with privacy regulations while remaining socially responsible.

By systematically addressing cross-cultural semantics, scaling the neuro-fuzzy architecture, and targeting real-world applications, this research aims to establish a globally adaptable, interpretable, and empathetic AI framework for sentiment analysis.

References

[1] L. A. Zadeh, "Fuzzy sets," *Information and Control*, vol. 8, no. 3, pp. 338–353, 1965.
[2] 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.
[3] H. Felbo, A. Mislove, A. Søgaard, I. Rahwan, and S. Lehmann, "Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm," in *Proc. Conf. on Empirical Methods in Natural Language Processing (EMNLP)*, 2017, pp. 1615–1625.
[4] T. J. Ross, *Fuzzy logic with engineering applications*, 4th ed. Hoboken, NJ, USA: Wiley, 2016.
[5] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
[6] P. Kralj Novak, J. Smailović, B. Sluban, and I. Mozetič, "Sentiment of emojis," *PLOS ONE*, vol. 10, no. 12, p. e0144296, 2015.
[7] P. Thakur, "A neuro-fuzzy framework for sentiment analysis on noisy social media data," *Journal of Artificial Intelligence, Machine Learning and Soft Computing*, 2025.
[8] K. Dave et al., "SentiQNF: A novel approach to sentiment analysis using quantum algorithms and neuro-fuzzy systems," *IEEE Trans. Comput. Soc. Syst.*, 2025.
[9] A. Khan et al., "Sentiment analysis of emoji fused reviews using machine learning and BERT," *Scientific Reports*, vol. 15, 2025.
[10] Y. Lou et al., "Sentiment classification of emoji and text: An analysis of model families, fusion strategies, and performance gains," *Communications on Applied Nonlinear Analysis*, 2024.
[11] A. Ben Meriem et al., "A fuzzy approach for sarcasm detection in social networks," *Procedia Computer Science*, vol. 192, pp. 2400–2409, 2021.
[12] M. Hasan et al., "Sarcasm detection over social media platforms using hybrid ensemble model with fuzzy logic," *Electronics*, vol. 12, no. 1, 2023.
[13] K. P. Gowda et al., "Transformers in sentiment analysis: A paradigm shift in deep learning research," *Journal of Information Systems Engineering and Management*, vol. 10, 2025.
[14] A. Zadeh et al., "Multimodal language analysis in the wild: CMU-MOSEI dataset and interpretable dynamic fusion," in *Proc. Assoc. Comput. Linguistics (ACL)*, 2018, pp. 2236–2246.
[15] J. S. R. Jang, "ANFIS: adaptive-network-based fuzzy inference system," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 23, no. 3, pp. 665–685, 1993.
[16] C. T. Lin and C. S. G. Lee, *Neural Fuzzy Systems: A Neuro-Fuzzy Synergism to Intelligent Systems*, Upper Saddle River, NJ, USA: Prentice Hall, 1996.
[17] Z. Huang, W. Xu, and K. Yu, "Bidirectional LSTM-CRF models for sequence tagging," *arXiv preprint arXiv:1508.01991*, 2015.
[18] J. Pennington, R. Socher, and C. D. Manning, "GloVe: Global vectors for word representation," in *Proc. Conf. on Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1532–1543.
[19] X. Hu et al., "We know how you feel: Sentiment analysis of emojis," in *Proc. 11th Int. Conf. on Web and Social Media (ICWSM)*, 2017, pp. 112–121.
[20] D. Maynard and M. A. Greenwood, "Who cares about Sarcastic Tweets? Investigating the Impact of Sarcasm on Sentiment Analysis," in *Proc. 9th Int. Conf. on Language Resources and Evaluation (LREC)*, 2014, pp. 4238–4243.
[21] R. Confalonieri et al., "Using Fuzzy Logic to Leverage Explainability in Social Media Sentiment Analysis," *IEEE Transactions on Fuzzy Systems*, vol. 29, no. 12, pp. 3624–3635, 2021.
[22] A. Adadi and M. Berrada, "Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI)," *IEEE Access*, vol. 6, pp. 52138–52160, 2018.