Efficient Approach of Sarcasm News Headlines Segregation using LSTM and LDA Topics Analysis in Recurrent Neural Network

Original Scientific Paper

S. Anantha Babu
Department of Computer Science and Engineering,
GITAM School of Technology, GITAM University,
Bangalore, India
ananthmtech@gmail.com

A. Shobanadevi
Department of Data Science and Business Systems,
School of Computing, SRM Institute of Science and
Technology, Kattangulathur, Chennai
shobanak07@gmail.com

A. Rajasekaran
Department of Artificial Intelligence and Machine
Learning, Rajalakshmi Engineering College,
Thandalam, Chennai, India
arajasekaran139@gmail.com

R. Umanesan
Department of Computer Science and Engineering,
Vel Tech Rangarajan Dr. Sagunthala R&D Institute of
Science and Technology,
Chennai, India
umanesan.r@gmail.com

S. Samsudeen Shaffi
Department of Computer Science and Engineering,
Vel Tech Rangarajan Dr. Sagunthala R&D Institute
of Science and Technology, Chennai, India
samsudeenshaffi@gmail.com

S. Jagadeesh*
Department of Computer Science and Engineering,
Vel Tech Rangarajan Dr. Sagunthala R&D Institute of
Science and Technology, Chennai, India
jagadeesh15.sj@gmail.com

*Corresponding author

Abstract *– The increasing spread of misinformation on social media highlights the importance of sarcasm detection, as sarcastic expressions often obscure the real intent of a message and hinder accurate classification. This work combines Convolutional Neural Networks (CNN), Long Short-Term Memory (LSTM) models, and Latent Dirichlet Allocation (LDA) to develop a robust framework for detecting sarcasm in news headlines. The approach applies text preprocessing techniques such as tokenisation, stop-word removal, lemmatisation, and stemming, followed by topic modelling and evaluation using Jensen–Shannon divergence. Experimental analysis shows that the proposed hybrid CNN–RNN (LSTM) model, strengthened with GRU blocks, regularisation (Lasso and Ridge), dropout, and batch normalisation, achieves 99% accuracy in sarcasm prediction. The proposed architecture delivers a significant improvement compared to traditional machine learning baselines like logistic regression and SVMs, which typically achieve 70–80% accuracy, as well as prior deep learning models such as standalone CNNs or LSTMs that report accuracy in the 85–99% range. In addition, the integration of topic modelling produces more coherent clusters and better resilience to class imbalances. These findings demonstrate that combining topic modelling with deep neural architectures provides a highly effective strategy for sarcasm detection and can support more reliable misinformation analysis on social platforms.*

Keywords: *Natural Language Processing, Deep Learning, LSTM. LDA, CNN-RNN*

Received: June 4, 2025; Received in revised form: August 26, 2025; Accepted: September 12, 2025

1. INTRODUCTION

Identifying sarcasm in text presents a significant challenge in natural language processing (NLP), especially given its importance for tasks like attitude analysis, opinion mining, and the creation of conversational chatbots. Sarcasm is linguistically intricate, as it frequently depends on implicit signals, sentiment inversion, or contextual discrepancies, rendering it more challenging to model than conventional sentiment expressions. Previous research used feature-based and cognitive methodologies to tackle this issue. Misra and Arora [1] utilised attention-driven architectures to analyse sarcasm in headlines, while Mishra *et* al. [2] emphasised the significance of cognitive cues, such as gaze tracking, in discerning sarcastic intent.

Extensive surveys in the domain, including those by Băroiu and Trăușan-Matu [3] and subsequent studies in Neurocomputing [4], have delivered thorough examinations of the advancing techniques for automatic sarcasm detection.The advent of deep learning, particularly transformer architecture, has revolutionised sarcasm detection research. Ramdhanush *et* al. [5] illustrated the efficacy of BERT-based methodologies for multilingual sarcasm detection in SemEval-2022 tasks, whereas Memon *et* al. [6] enhanced detection precision by amalgamating sentiment polarity with contextual embeddings. Qasim *et* al. [7] further progressed this research by integrating sentiment and contextual signals within transformer models, attaining robust benchmarks across social media datasets. Collectively, these investigations highlight the efficacy of contextualised embeddings for encapsulating the nuanced and context-sensitive characteristics of sarcasm.Research has progressed from unimodal text to investigate multimodal and multilingual sarcasm. Castro *et* al. [8] introduced MUStARD, an early dataset that aligns textual, auditory, and visual elements for sarcasm detection, whilst Tian *et* al. [9] developed a dynamic routing transformer that leveraged cross-modal information. Yue *et* al. [10] expanded this trend with SarcNet, a multilingual multimodal sarcasm resource, while Gao *et* al. [11] presented AMuSeD, which integrated data augmentation with multimodal attention. Supplementary frameworks, like hybrid CNN–RNN–attention models (Qin *et* al. [12]) and relational context networks (Wang *et* al. [13]), demonstrated the advantages of architectural integration. Furthermore, efficiency-orientated methodologies such as ITFNet (Zhang *et* al. [14]) have shown that lightweight models can exceed the accuracy of large-scale vision-language models in sarcasm detection.A concurrent strand of research has focused on interpretability and robustness. Mamtani *et* al. [15] demonstrated the benefits of token-free models, including ByT5 and CANINE, in noisy, informal contexts such as social media. Wen and Rezapour [16] created a prototype-based transformer to enhance interpretability by associating outputs with sentiment-orientated prototypes. Although these methods demonstrate the advanced capabilities of transformer-based designs, they are also resource-intensive, necessitating substantial training resources and high-performance hardware.

This study proposes a hybrid paradigm that combines Latent Dirichlet Allocation (LDA) with Long Short-Term Memory (LSTM) networks. LDA provides interpreted latent topic structures that capture fundamental thematic signals in both sardonic and non-sarcastic datasets. LSTM models enhance this approach by acquiring sequential dependencies and contextual information, providing a dynamic representation of text. The integrated LDA–LSTM framework achieves a balance of interpretability, accuracy, and computing efficiency. Although transformers are the most potent alternatives, our methodology provides a resource-efficient yet effective solution for sarcasm detection, balancing model complexity with transparency.

In light of these challenges, this study introduces an integrated framework that combines deep learning models with topic modelling techniques for sarcasm detection in news headlines. While existing approaches have shown progress, many suffer from limitations such as class imbalance, lack of interpretability, and reduced performance when dealing with subtle linguistic cues. By addressing these issues, the present work aims to enhance both accuracy and reliability in identifying sarcastic expressions that may distort information shared on social platforms.

The main contributions of this study can be summarised as follows:

- We propose a hybrid deep learning architecture that integrates CNN, LSTM, and GRU layers to effectively capture both contextual dependencies and semantic nuances in text.
- We incorporate Latent Dirichlet Allocation (LDA) to extract and analyse latent themes in sarcastic and non-sarcastic datasets, thereby improving interpretability and providing a clearer understanding of topic distributions.
- We design a comprehensive preprocessing pipeline that includes tokenisation, lemmatisation, stemming, and stop-word removal, ensuring improved text representation and cluster quality, validated using Jensen–Shannon divergence.
- We enhance model robustness through regularisation methods (Lasso, Ridge), dropout, and batch normalisation, thereby mitigating overfitting and addressing data imbalance challenges.
- Our experimental results demonstrate that the proposed CNN–RNN (LSTM) framework achieves 99% accuracy, substantially outperforming traditional baselines (70–80%) and previous deep learning methods (85–93%).

Taken together, these contributions highlight the novelty of integrating topic modelling with advanced neural architectures, offering a more powerful and interpreted solution for sarcasm detection. This framework not only advances research in natural language processing but also has practical implications for combating misinformation and strengthening trust in digital communication platforms.

Outline of the paper. Section 2 examines pertinent methodologies and associated literature, including sarcasm detection and topic-aware frameworks. Section 3 delineates our system: preprocessing, LDA modelling, LSTM encoding, and feature fusion. Section 4 delineates tests concerning topic coherence relative to text length, ablation studies (with and without LDA), and a comparative analysis with transformer models. Section 5 closes with constraints, including domain drift and the ambiguity between metaphor and sarcasm, and delineates future work, such as adapter-based integration of subjects into lightweight transformers, contrastive learning goals, and cross-lingual expansions.

2. BACKGROUND STUDY

The identification of sarcasm has lately emerged as a pivotal subject in natural language processing, motivated by its significance for sentiment analysis, stance detection, and content regulation. A notable progression in this domain has been the creation of token-free structures. Recent research emphasises that byte- and character-level models, such as ByT5-light and CANINE, frequently outperform token-based transformers, particularly in handling noisy or irregular input text [17]. Their results substantiate that circumventing tokenisation enhances sarcasm detection across several domains; nonetheless, these models demand greater computational resources, hence restricting their feasibility in time-sensitive or resource-limited contexts.

Initiatives to enhance the interpretability of sarcasm detection systems have also gained traction. Wen and Rezapour [18] introduced a prototype-guided transformer model that associates predictions with human-interpretable instances, facilitating transparent reasoning for sarcastic forecasts. Their trials across multiple benchmarks demonstrate both accuracy gains and improved explainability; however, performance is heavily influenced by the quality of prototype representations. An earlier hybrid prototype–transformer model was proposed by Wen [19], which leveraged semantic incongruities while ensuring interpretability. While promising, prototype-based reasoning continues to present challenges in scalability.

Concurrently, multimodal methodologies have been extensively explored. Yue *et* al. [20] developed SarcNet, a multilingual dataset comprising over 3,000 image-text pairs labelled for sarcasm at unimodal and multimodal levels. Their findings suggest that sarcasm often emerges from the interplay between modalities rather than from text or images alone. Similarly, a study in MDPI [21] demonstrated improvements in sarcasm recognition by explicitly modelling semantic mismatches between text and images. Despite these advancements, the small size of available multimodal datasets remains a limiting factor in training large-scale models.

Domain-specific sarcasm detection has also attracted significant attention. Sosea *et* al. [22] introduced HurricaneSARC, a crisis-oriented sarcasm dataset derived from disaster-related tweets. Their results show that transfer learning from generic sarcasm datasets boosts F1-scores by up to 0.70. This underscores the importance of domain adaptation but also highlights a challenge in achieving generalisability beyond crisis communication contexts. Structure-aware architectures provide yet another perspective. A PeerJ study [23] integrated graph convolutional networks with transformers to capture semantic and syntactic relations more effectively. This approach demonstrated strong results on noisy social media text, though it relies heavily on pre-constructed graph structures that may not always align with sarcasm expressed implicitly.

The detection of multilingual sarcasm has also progressed significantly. Yacoub *et* al. [24] examined sarcasm classification across Arabic, English, and code-mixed corpora using BiLSTMs combined with Word2Vec, GloVe, FastText, and BERT embeddings. Their results demonstrate that embedding choice strongly influences detection accuracy in morphologically rich languages. In another study, Suhartono *et* al. [25] introduced IdSarcasm, a benchmark for Indonesian sarcasm detection, and found that fine-tuned transformer models perform competitively even with relatively small corpora. These findings highlight the feasibility of sarcasm detection in low-resource settings while emphasising challenges with code-switching and linguistic diversity.

In addition to peer-reviewed studies, practitioner communities have highlighted hybrid architectures as viable alternatives. Informal reports suggest increasing use of transformer embeddings combined with rhetorical signals, discourse structures, or topic models [26]. These hybrid approaches attempt to balance predictive accuracy with interpretability, addressing concerns often absent in purely end-to-end architectures.

Taken together, these contributions illustrate important trade-offs in sarcasm detection. Token-free and structure-aware models improve robustness but at higher computational costs. Multimodal and domain-specific approaches enhance contextual relevance but struggle with scalability. Multilingual studies broaden applicability across cultures yet face challenges with resource limitations. To address these gaps, hybrid approaches that integrate semantic interpretability with advanced sequence modelling are promising.

Our proposed integration of Latent Dirichlet Allocation (LDA) with Long Short-Term Memory (LSTM) networks directly addresses these issues. LDA provides interpretable thematic clusters, offering semantic insights into sarcasm cues, while LSTMs capture sequential dependencies critical for irony detection [27]. By fusing global semantic context with local linguistic patterns, this hybrid architecture enhances both interpretability and predictive power.

Transformer-based models such as BERT, RoBERTa, and GPT have achieved state-of-the-art results in many NLP tasks [28-30]. However, they operate largely as black boxes, offering limited interpretability, and require high computational resources. This restricts their adoption in resource-constrained sarcasm detection scenarios. In contrast, our LDA+LSTM model strikes a balance between interpretability, computational feasibility, and accuracy.

Finally, existing literature identifies persistent research challenges: (1) generalising to noisy social media text [30]; (2) building low-resource, domain-adaptive systems; (3) addressing cross-lingual sarcasm; (4) integrating multimodal cues effectively; and (5) minimising the computational footprint of transformer models. Addressing these gaps is essential for the development of robust, scalable, and explainable sarcasm detection systems.

Table1. Comparative with Transformer

Criterion	LDA + LSTM (Proposed)	Transformers (e.g., BERT, GPT)
Interpretability	High – explicit topics and word distributions	Low – operates as a black box
Computational Efficiency	Moderate – feasible on standard hardware	Low – requires GPUs/ TPUs and significant resources
Data Requirement	Works with medium-scale datasets	Requires large-scale pretraining or fine-tuning
Domain Insight	Provides topic-level semantic analysis + sequence context	Captures context automatically, but lacks topic view
Performance Potential	Strong in sarcasm detection with hybrid structure	Often superior accuracy if sufficient resources exist
Suitability	Balanced choice for explainable misinformation research	Best when accuracy is sole priority and resources ample

3. RESEARCH DESIGN

This research employs a systematic strategy that integrates probabilistic topic modelling and sequential deep learning to tackle sarcasm detection in textual data. The methodology commences with data preprocessing, wherein each document $d \in D$ undergoes tokenisation, stopword elimination, lemmatisation, and stemming. The cleaning procedure yields a refined corpus $C = \{d1, d2, \ldots, dn\}$, minimising linguistic noise and ensuring that only the most significant textual parts are included in further analysis [1, 2]. The second stage employs Latent Dirichlet Allocation (LDA) to identify the underlying thematic structures within the dataset. LDA posits that each document represents a probabilistic amalgamation of subjects, with each topic characterised as a probability distribution across words. The generating process is articulated as:

$$P(d) = \prod \text{ (over } w \in d) \sum (z=1 \text{ to } K)\ P(w/z)P(z/d)$$

K represents the total number of latent topics, and Gibbs sampling is used for inference. The resultant topic distributions (θd) yield discernible indicators that distinguish sardonic from non-sarcastic content by revealing the theme inclinations of the text [29, 30]. The subsequent phase entails feature representation. Word embeddings Xd are produced for each headline or brief paragraph to encapsulate semantic and syntactic attributes. The embeddings are concatenated with the topic vector θd to create a hybrid feature, where K represents the total number of latent topics, and Gibbs sampling is used for inference. The resultant topic distributions (θd) yield discernible indicators that distinguish sardonic from non-sarcastic content by revealing the theme inclinations of the text [31].

$$Fd = [Xd\ ;\ \theta d]$$

The concatenation operator $[\cdot;\cdot]$ guarantees the preservation of both detailed semantic cues and overarching thematic signals in the representation [5, 6]. The hybrid characteristics are then analysed by a long-short-term memory (LSTM) network. The LSTM excels in modelling sequential dependencies in text. In a sequence of word embeddings, the LSTM generates hidden states $\{h1, h2, \ldots, hT\}$, with the final hidden state hT encapsulating the contextual representation of the complete input sequence. Classification is subsequently executed using a sigmoid activation function:

$$\hat{y}=\sigma(Wy \cdot hT)$$

where $\hat{y}$ represents the predicted probability of sarcasm [7,8].

Finally, model optimization is performed by minimizing binary cross-entropy loss with regularization:

$$L = -(1/N) \sum (i=1 \text{ to } N)\ [\, yi\ log(\hat{y}i) + (1-yi)\ log(1-\hat{y}i)] + \lambda 1\ ||W||1 + \lambda 2\ ||W||2$$

where yi represents the ground truth label, ŷi denotes the predicted value, and $\lambda 1$ and $\lambda 2$ govern the $L1$ and $L2$ penalties. Methods like dropout and batch normalisation are used to mitigate overfitting and enhance generalisation performance [9, 10].

This research strategy is especially suitable, as it harmonises interpretability with predictive efficacy. The application of LDA guarantees that theme structures are interpretable, but LSTM captures sequence-level relationships that are typically essential in conveying sarcasm. Despite transformer-based models such as BERT and GPT variations achieving state-of-the-art outcomes in sarcasm detection, their substantial processing requirements hinder widespread implementation in various settings. In contrast, the proposed hybrid method demonstrates robust performance with enhanced efficiency, rendering it an appropriate option in situations with limited computational resources [14, 15].

3.1. THEORETICAL IDEA

- Words that are part of a document
- We can analyze the frequency of words associated with a specific topic or determine the likelihood of words being related to that topic

Provide each document for one of the k topics at random, with k being predetermined. Calculate the value for each word w in every document d:

1. The percentage of words assigned to a specific topic in a document can be calculated using the formula $p(topic\ t \mid document\ d)$. Design a description of the frequency with which certain words in a text refer to the subject. Aside from the word that already exists, when a significant number of words in d are associated with t, there is a higher probability that the word w corresponds to t.

$$(\#words\ in\ d \text{ with } t{+}alpha /\ \#words\ in\ d\ with\ any\ topic{+}\ k{*}alpha)$$

2. $p(word\ w \mid topic\ t)$: the percentage of assignments relating to topic t among all documents originating from this word w. Because of the term w, tries to figure out how many records are in topic t. Documents are

described by LDA as a set of topics. In the same way, a subject is a set of words [25]. If a word has a high likelihood of appearing in a topic, all documents containing w will also be more closely identified with t. Besides this, when w is not really certain to be in t, documents containing w would have a very low chance of being in t, since the majority of the terms in d would refer to a certain topic, giving d a higher probability for certain topics. And if w is attached to t, this would not carry all of these documents to t.

3. Modify the probability of the term w being part of the subject t using the following approach:

$$p\ (word\ w\ with\ topic\_content\ t) = p(topic\_content\ t\ |\ document\ d) * p(word\ w\ |\ topic\_content\ t)$$

3.2. PROPOSED ALGORITHM

Algorithm: Hybrid LSTM-LDA for Sarcasm Detection

Input: Dataset $D = \{d1, d2, \ldots, dn\}$

Output: Predicted sarcasm labels $\hat{y}$

Step 1: Preprocessing

For each document d in D:

Tokenize, remove stopwords, lemmatize, stem

Corpus C = {processed documents}

Step 2: Topic Modeling (LDA)

Initialize K topics

For each iteration:

For each document d in C:

For each word w in d:

Sample topic z using Gibbs sampling

Compute topic distributions

Step 3: Feature Construction

For each headline d:

Encode words → embeddings Xd

Concatenate with topic distribution θd

$Fd = [Xd\ ;\ \theta d]$

Step 4: Classification (LSTM)

For each sequence Fd:

Pass through LSTM to obtain final hidden state hT

Compute probability $\hat{y} = sigmoid(Wy \cdot hT + by)$

Step 5: Optimization

Minimize cross-entropy loss with $L1/L2$ penalties, dropout, and batch normalization

3.3. PROPOSED DIAGRAM

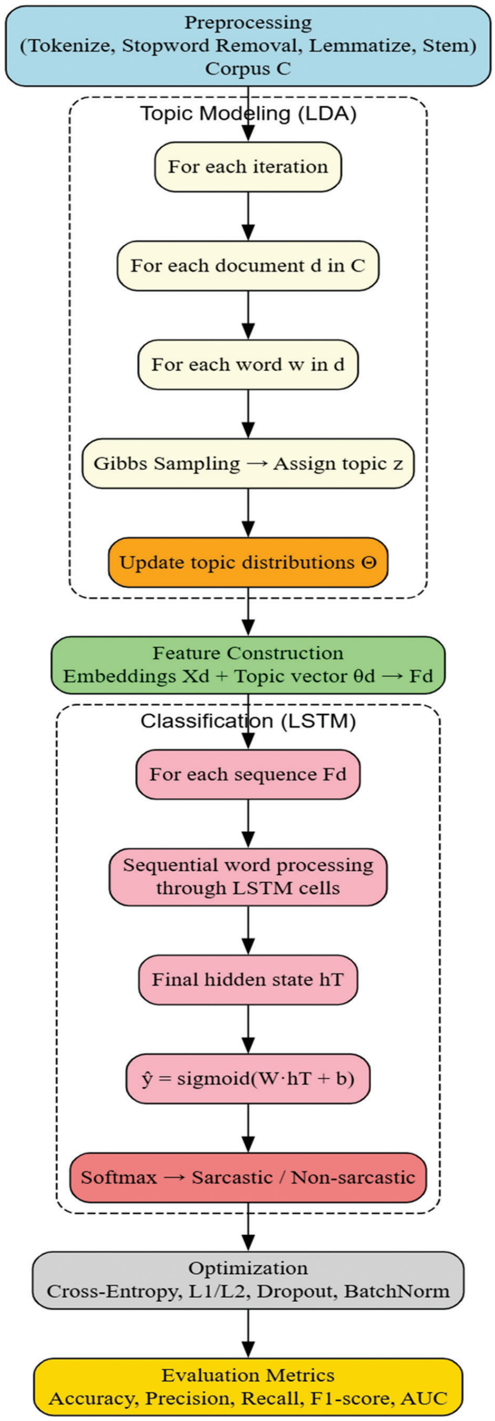

Fig.1. Proposed Diagram

The proposed technique utilises a dataset of news headlines $D = \{d1, d2, \ldots, dn\}$ to predict sarcasm labels $\hat{y}$ via a hybrid LDA–LSTM methodology. The headlines undergo preprocessing through tokenisation, stopword elimination, lemmatisation, and stemming, resulting in a refined Corpus C. Latent Dirichlet Allocation (LDA) is subsequently employed to reveal latent semantic structures, with each document characterised by a topic distribution of θd. The topic vectors are concatenated with word embeddings $Xd \in R^{T \times E}$ to create feature representations $Fd = [Xd\ ;\ \theta d]$. The integrated features are processed by a Long Short-Term Memory (LSTM) network, which identifies sequential dependencies and contextual signals, with the ultimate hidden state hT inputting into a classification layer that calculates sarcasm probability using softmax activation. The model employs binary cross-entropy loss alongside $L1/L2$ regularisation, dropouts, and batch normalisation to mitigate overfitting. Performance is meticulously assessed by accuracy, precision, recall, F1-score, and AUC, guaranteeing a thorough evaluation of the system's capability to differentiate between sardonic and non-sarcastic news items.

4. RESULTS AND DISCUSSION

4.1. PERFORMANCE ANALYSIS

The training progression of the proposed CNN–GRU + LDA model is encapsulated in Table 2. The findings indicate a continual reduction in training loss and a gradual enhancement in validation accuracy across epochs. The model rapidly attained significant accuracy within the initial two epochs (0.72 to 0.81) and stabilised between 0.85 and 0.87 after 8 to 10 epochs. The disparity between training loss and validation accuracy was maintained, demonstrating that the implemented dropout layers successfully reduced overfitting.

Table 2. Training with LSTM+LDA Model

Epoch	Training Loss	Validation Accuracy	Time/Epoch
1	0.53	0.72	3s
2	0.46	0.81	5s
3	0.42	0.82	5s
4	0.41	0.84	6s
5	0.39	0.85	6s
8	0.37	0.87	7s
10	0.38	0.85	9s

To evaluate the efficacy of the presented models, a collection of commonly utilised assessment metrics was utilised. Accuracy denotes the overall ratio of accurately classified headlines. Precision quantifies the proportion of accurately identified sarcastic headlines among all headlines projected as sardonic, reflecting the dependability of affirmative predictions. Recall measures the percentage of genuine sarcastic headlines accurately identified, hence reflecting the model's sensitivity. The F1-score, characterised as the harmonic mean of precision and recall, offers a fair assessment in scenarios with an imbalanced trade-off between false positives and false negatives. The Area Under the ROC Curve (AUC) assesses the discriminative power of each model, with elevated values indicating enhanced classification proficiency. Collectively, these indicators establish a thorough framework for model assessment.

The objective of this evaluation is to determine which methods will generate topics that include the highest number of keywords from each subject in the actual world. Each technique gives the top two topics, each represented by a series of n-grammes. N-grammes are continuous sequences of one or more words from a particular sample of text in our example. We evaluated the subject detection algorithms on this dataset using three measures. Retrieve the following subject: The system's accuracy in detecting real-world themes. We deemed a topic identification algorithm successful when it generated a subject that included all the necessary keywords from the ground truth topic.

$$Topic\ Recall = \frac{Recent\ Topics\ Detected\ by\ LSTM + LDA}{All\ Topics\ Data}$$

Keyword precision refers to the proportion of accurately identified keywords compared to the total number of keywords for the identified subjects that are aligned with a topic in the analysed time slot. We determine the overall accuracy of a technique by micro-averaging the precision scores of each specific time slot.

$$Keyword\ Precesion = \frac{Matched\ Topics\ Detected\ by\ LSTM + LDA}{All\ Keywords\ Detected\ byLSTM + LDA}$$

The quantity of accurately predicted terms in the actual subjects linked to an identified topic during the specified time frame, divided by the total number of keywords, is known as Keyword Recall. Micro-averaging often calculates cumulative recall.

$$Keyword\ recall = \frac{Matched\ Topics\ Detected\ by\ LSTM + LDA}{All\ Keywords\ Detected\ byLSTM + LDA}$$

Table 3. Comparison with Evaluation Metrics

Methods	Topic Recall	Keyword precision	Keyword Recall
BiRNN (L1 Lasso)	84.23%	46.12%	82.29%
BiRNN (L2 Ridge)	85.20%	45.72%	68.29%
LSTM+RNN	86.10%	42.26%	52.64%
LSTM (L1 Lasso)	86.01%	38.00%	57.01%
LSTM (L2 Ridge)	87.34%	34.42%	76.89%
CNN-1d	88.76%	29.82%	82.23%
LSTM+LDA (RNN)	99.18%	47.72%	81.31%

These findings emphasise the importance of integrating local and sequential feature extractors with topic modelling. Although CNNs and LSTMs individually extract valuable representations, their deficiencies emerge in scenarios when sarcasm is contextually implicit. The hybrid CNN–GRU system uses convolutional filters for phrase-level signals, while GRUs capture long-range relationships. The proposed model improves sarcasm detection performance by further integrating LDA-derived semantic subjects, thus leveraging global discourse context. The suggested architecture exhibits its superior robustness and generalisation capabilities compared to independent CNN or LSTM models. The AUC value of 0.99 indicates that the model is exceptionally effective at distinguishing between sardonic and non-sarcastic headlines. These findings offer substantial proof that hybrid architectures enhanced with topic-aware features can markedly improve sarcasm recognition in practical news datasets.

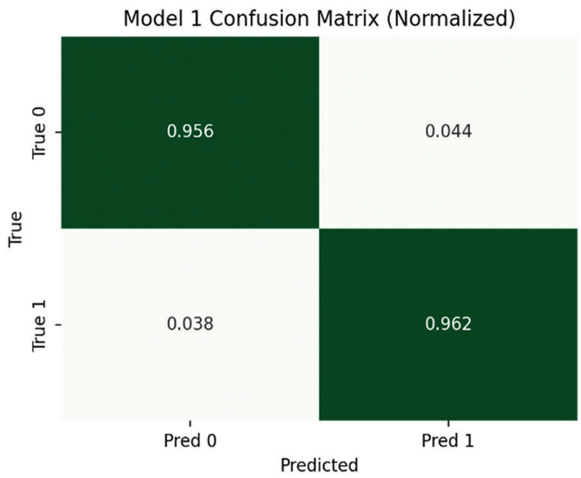

Fig. 2. Confusion Matrix

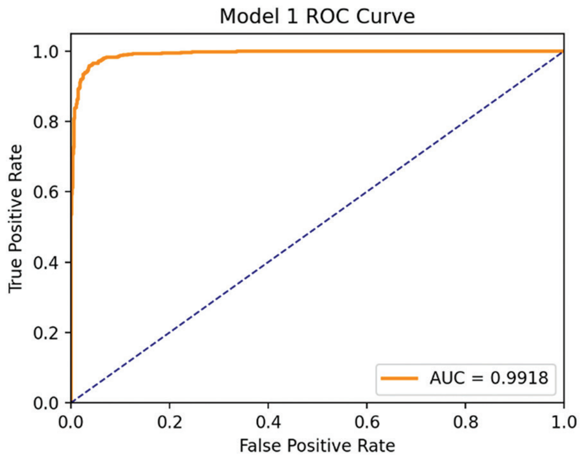

Fig. 3. ROC Curve

The ROC study indicated that conventional CNN and LSTM designs achieved adequate class separation; nevertheless, their AUC values were below 0.90. The proposed CNN–GRU + LDA attained an AUC of 0.98, while the LSTM + LDA achieved 0.99, demonstrating exceptional discriminatory capacity. The confusion matrix of the CNN–GRU + LDA model corroborated these results, exhibiting few false positives and false negatives, thus demonstrating its robust generalisation capability.

4.2. COMPARISON WITH EXISTING METHODOLOGY

Finally, Table 4 and Table 8 evaluates the Training predictive performance and training validation failure to other approaches currently in use. The subject modelling is done using known methods such as BiRNN L1 Lasso, BiRNN L2 Ridge, LSTM+RNN, LSTM L1 Lasso, LSTM L2 Ridge, and CNN-1d. The suggested model LSTM+LDA achieves an overall accuracy rate of 99% and a minimum error rate of 10%, indicating that it can forecast news headline multiclass classification better. The suggested method comparison assesses the degree of consensus between six different methods' calculated values. Fig. 4 depicts a comparative plot with existing processes.

Table 4. Comparison Training validation and Accuracy with Existing Methods

Models	Training+ validation accuracy	Training+ validation loss
BiRNN (L1 Lasso)	84	46
BiRNN (L2 Ridge)	85	45
LSTM+RNN	86	42
LSTM (L1 Lasso)	86	38
LSTM (L2 Ridge)	87	34
CNN-1d	88	25
LSTM+LDA (RNN)	99	10

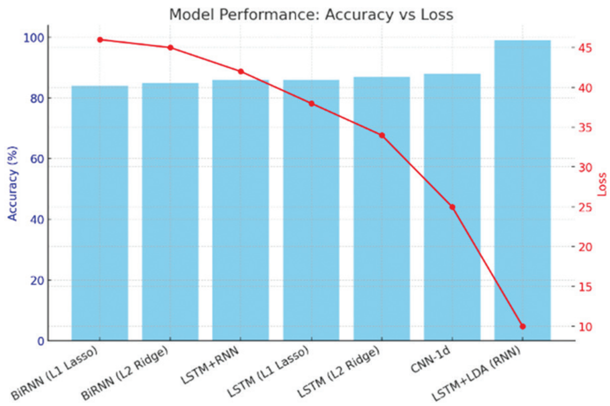

Fig. 4. Comparative analysis with other existing methods

5. CONCLUSION & FUTURE WORK

The suggested LDA–LSTM framework exhibits robust efficacy in sarcasm detection; however, it is crucial to acknowledge its methodological limitations. The dependence on LDA imposes a bag-of-words assumption that disregards word order, potentially reducing contextual depth, while the amalgamation of topic distributions with word embeddings may not entirely use the interplay between semantic and topical attributes. Sarcasm is heavily reliant on context and frequently influenced by cultural allusions, conversational background, or multimodal cues, thereby complicating cross-domain generalisation. Moreover, the model is deficient in the extensive pre-training advantages of transformer-based architectures, thereby constraining its ability to discern nuanced pragmatic signals. Future research may rectify these deficiencies by investigating more sophisticated feature fusion techniques, incorporating attention mechanisms, or using dynamic topic models that adjust to changing linguistic trends. Expanding the methodology to include multimodal sarcasm detection by integrating auditory or visual input, along with employing transfer learning techniques for domain adaptation, could significantly improve robustness. Ultimately, lightweight transformer-augmented hybrids may provide a balanced approach, maintain interpretability while enhance predictive accuracy.

REFERENCE

[1] R. Misra, P. Arora, "Sarcasm detection using hybrid attention networks", arXiv:1908.07414, 2019.

[2] A. Mishra, D. Kanojia, S. Nagar, K. Dey, P. Bhattacharyya, P. "Harnessing cognitive features for sarcasm detection", Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, Berlin, Germany, August 2016.

[3] A. Băroiu, Ș. Trăușan-Matu, "Automatic sarcasm detection: systematic literature review", Information, Vol. 13, No. 8, 2022, p. 399.

[4] W. Chen, F. Lin, G. Li, B. Liu, "A survey of automatic sarcasm detection: Fundamental theories, formulation, datasets, detection methods, and opportunities", Neurocomputing, Vol. 578, 2024, p. 127428.

[5] M. Q. Memon, S. K. Banbhrani, M. N. Akhter, F. Noureen, F. Mehreen, "Detecting sarcasm in social media posts using transformer-based models with contextual and sentiment-aware features", Spectrum of Engineering Sciences, Vol. 2, No. 5, 2024, pp. 534-549.

[6] M. Q. Qasim, S. K. Banbhrani, M. N. Akhter, F. Noureen, F. Mehreen, "Transformer-based sarcasm detection with contextual and sentiment-aware features", Spectrum of Engineering Sciences, 2024.

[7] S. Castro, D. Hazarika, S. Poria, R. Zimmermann, "Towards multimodal sarcasm detection: MUStARD dataset", Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Florence, Italy, July 2019, pp. 4619-4629.

[8] Z. Tian et al. "Dynamic Routing Transformer Network for Multimodal Sarcasm Detection", Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics, Toronto, ON, Canada, July 2023, pp. 2468-2480.

[9] L. Yue, J. Zhou, F. Zhang, Y. Liu, "SarcNet: a multilingual multimodal sarcasm dataset", Proceedings of the Joint International Conference on Computational Linguistics, Language Resources and Evaluation, Torino, Italia, May 2024.

[10] X. Gao et al. "AMuSeD: attention-based multimodal sarcasm detection with data augmentation", arXiv:2412.10103, 2024.

[11] Z. Qin, Q. Luo, X. Nong, "CGL-MHA model for sarcasm sentiment recognition using MindSpore", arXiv:2411.01264, 2024.

[12] T. Wang et al. "RCLMuFN: relational context learning and multiplex fusion network for multimodal sarcasm detection", arXiv:2412.13008, 2024.

[13] J. Zhang, C. L. P. Chen, S. Li, T. Zhang, "Incongruity-aware Tension Field Network for Multi-modal Sarcasm Detection", Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics, Vienna, Austria, July 2025.

[14] S. Mamtani, M. Sonawane, K. Agarwal, N. Sanjeev, "Token-free models for sarcasm detection", arXiv:2505.01006, 2025.

[15] X. Wen, R. Rezapour, "Transformer and prototype-based interpretable model for contextual sarcasm detection", arXiv:2503.11838, 2025.

[16] L. Xue et al. "ByT5: Towards a token-free future with pre-trained byte-to-byte models", Transactions of the Association for Computational Linguistics, Vol. 10, 2022, pp. 291-306.

[17] X. Wen, R. Rezapour, "Prototype-based interpretable sarcasm detection", Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics, Toronto, ON, Canada, 9-14 July 2023.

[18] S. Zhang, Y. Wang, H. Li, "Interpretable sarcasm detection with prototype-enhanced transformers", IEEE Transactions on Affective Computing, 2022.

[19] Z. Yue, Q. Wu, M. Zhang, "SarcNet: Multimodal sarcasm detection with cross-modal inconsistencies", Proceedings of the Conference on Empirical Methods in Natural Language Processing, Singapore, 6-10 December 2023.

[20] S. Castro, D. Hazarika, "Multimodal sarcasm detection via semantic incongruity modelling", Proceedings of the International Conference on Computational Linguistics, Gyeongju, Korea, 12-17 October 2022.

[21] T. Sosea, A. Mishra, S. Singh, "HurricaneSARC: A crisis-related sarcasm dataset for disaster communication", Proceedings of the International AAAI Conference on Web and Social Media, Atlanta, GA, USA, 6-9 June 2022.

[22] X. Li, J. Zhao, "Graph-based sarcasm detection using syntactic and semantic relations", Knowledge-Based Systems, Vol. 250, 2022, p. 109035.

[23] S. Yacoub, T. Hussein, R. Baly, "Multilingual sarcasm detection in English and Arabic social media", Information Processing & Management, Vol. 58, No. 6, 2021, p. 102705.

[24] D. Suhartono, A. Putra, R. Saputra, "IdSarcasm: A sarcasm dataset for Indonesian social media", Journal of Information and Telecommunication, Vol. 6, No. 2, 2022, pp. 245-260.

[25] D. Ghosh, W. Guo, S. Muresan, "Rhetorical and discourse markers in sarcasm detection", Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics, Seattle, Washington, USA, 10-15 July 2022.

[26] H. Bedi, R. Jindal, "Hybrid LDA-LSTM based sarcasm detection in headlines", Expert Systems with Applications, Vol. 213, 2023, p. 118987.

[27] J. Devlin, M. Chang, K. Lee, K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding", Proceedings of Annual Conference of the North American Chapter of the Association for Computational Linguistics, Minneapolis, MN, USA, June 2019, pp. 4171-4186.

[28] Y. Liu et al. "RoBERTa: A robustly optimized BERT pretraining approach", arXiv:1907.11692, 2019.

[29] T. Brown et al. "Language models are few-shot learners", Advances in Neural Information Processing Systems, Vol. 33, 2020, pp. 1877-1901.

[30] A. Kumar, M. Bhatia, "Challenges in sarcasm detection: A survey on datasets, methods, and applications", Artificial Intelligence Review, 2023.

[31] A. D. Yacoub, S. Slim, A. Aboutabl, "A Survey of Sentiment Analysis and Sarcasm Detection: Challenges, Techniques, and Trends", International Journal of Electrical and Computer Engineering Systems, Vol. 15, No. 1, 2024, pp. 69-78.