International Journal of Applied Smart Interdisciplinary Technologies (IJASIT)

Home Page: https://ijasit.org

TruthLens: Graph Analytics for Fake News Detection

Ayushi Shukla[1, *], Kushagra Kumar Porwal[2]

[1 2] School of Computer Science & Engineering, Galgotias University, Greater Noida, India

*** Corresponding author.**
E-mail:
Ayushi.22scse1120011@galgotiasuniversity.edu.in

Article History:
Received Date: 05 March 2026
Revised Date: 05 April 2026
Accepted Date: 10 May 2026
Published Date: 27 June 2026

Keywords:

Fake News Detection; Graph Neural Network; BERT; Multimodal Fusion; Social Network Analysis; GNN , CLIP.

ABSTRACT

Fake news has evolved from a minor online problem to a potent instrument for influencing public opinion and un-dermining confidence in digital platforms. Conventional detec-tion systems function similarly to linguistic spell checkers, but they frequently overlook the intricate, well-planned nature of disinformation campaigns. In order to close this gap, TruthLens proposes a hybrid, multimodal framework that integrates graph analytics, computer vision, and natural language processing (NLP). Our approach makes use of PyTorch Geometric (GNNs) to monitor the propagation of information in social networks, BERT for comprehensive semantic text analysis, and CLIP for consistency between text and images. TruthLens identifies structural anomalies that indicate bot networks and organized influence operations by examining not only what is said but also how and by whom it spreads. According to preliminary tests, this layered approach performs significantly better than conventional text-only models, providing both greater accuracy and lucid insights into the mechanisms of digital deception.

1. Introduction

Information plays a vital role in human activities and significantly influences everyday life practices. In earlier times, news and information were primarily disseminated through print media such as newspapers and electronic media including television and radio. Information published through these channels was generally considered credible, as it was either self-screened or regulated by subject matter experts [1].

In recent years, individuals have been exposed to an over-whelming volume of information from diverse sources, particularly this is due to the mushrooming of the internet and the inherent social media platforms that are web based. Due to the ease of accessing misinformation, it has been dangerous enough to propagate it in the forms of misinformation that we are familiar with. The ease of online access has led to the dangerous proliferation of mis-information in various forms, including malicious discourse, deception, fabrications, fake news, and spam content, which spreads rapidly and extensively across society [2] [3]. As online social media has become a dominant mode of communication and networking, misinformation has come to be a global concern on the credibility of the populace and the integrity of the community.

Social platforms and blogs increasingly host fabricated and misleading news content, which negatively impacts society [4]. Such content often contains distorted facts and deceptive narratives, leading to interpersonal anxiety and widespread social panic. Moreover, unreliable information erodes public trust and adversely affects economic systems and critical political processes such as stock markets and elections [5] [6].

In the past, people used human experts' manual fact-checking to identify fake news. Nevertheless, this procedure is inefficient, time-consuming, labor-intensive, and subjective. Automated systems based on natural language processing and machine learning have been developed to identify fake news in order to address these problems [13], [14]. These systems are now better at stopping the spread of false information thanks to recent advancements in artificial intelligence. As a result, automated fake news detection has drawn more attention from researchers.

The multimodal fake news detection framework TruthLens is presented in this paper. By integrating textual analysis, image and text consistency checks, and graph-based propagation behavior analysis, it surpasses conventional text-only approaches. Unlike traditional systems that rely solely on article content, TruthLens examines how information propagates across social networks, enabling deeper insights into coordinated misinformation campaigns [7][8].

In order to assess textual credibility, the framework makes use of Bidirectional Encoder Representations of Transformer (BERT) to obtain deep contextual semantics to generate probabilistic credibility scores. To authenticate the multimodal, Contrastive Language–Image Pretraining (CLIP) is utilized to measure semantic alignment between textual claims and associated images, facilitating the detection of misleading or reused visuals. Furthermore, graph analytics are applied to model user interaction networks, enabling the identification of suspicious propagation patterns commonly associated with fake news dissemination [9] [10].

The main contributions of the work are the following:

- A hybrid fake news detection framework integrating textual analysis, image–text consistency evaluation, and graph-based social behavior modeling within a unified pipeline.
- A BERT-based contextual text analysis module that eliminates reliance on traditional feature engineering techniques such as TF-IDF and count vectorization.
- A graph analytics component designed to detect coordinated user behavior patterns that text-only systems fail to capture.
- A weighted fusion method that unifies multimodal out-puts into a single TruthLens score, increasing robustness and interpretability.
- The evaluation of the proposed system according to the standard performance indicators including latency analysis, accuracy, precision, recall, and F1-score.

This study is distinct because to its multimodal and network-aware structure. TruthLens can also detect fake news, and this is based on the content, the individuals and the delivery method of dissemination. The suggested model is therefore particularly appropriate for real-world social media environ-ments, where false information travels swiftly via structured networks.

2. Literature Survey

The last trend that is relevant to Natural Language Processing (NLP) has been the rise in the capability to detect fake news. Models that are based on transformers such as BERT, RoBERTa, XLNet and T5 have been remarkably effective in classifying fake news. They receive the context and meaning of the text. Such models offer a deeper analysis of complicated linguistic patterns as compared to the traditional approaches. Knowledge graph and social network analysis are some of the graph-based learning approaches that have been investigated to approximate the credibility of news sources and textual analysis [11] [12]. Such techniques can provide immense insight on information dissemination since they represent the correlation among users, publishers and content. Some questions have been raised though they have been successful, that include, clarity of models, bias of data and manipulation of texts. To curb such concerns of transparency, explainable Artificial Intelligence methods like LIME and SHAP have been used to enable the model predictions. Furthermore, the latest literature turned out to be obsessed with the detection and protection of the fake news that was produced with the assistance of the developed AI models [15][16].

The first fake news detectors were based on the normal machine learning approaches alone. They used hand-selected elements of text information. These methods were somewhat effective, although they were not comprehended simply and were overfit. The emergence of these issues led to the expansion of the popularity of deep learning and approaches related to NLP [17]. They are more modern and more effective in the context and generalization. Deep Learning-Based Fake News Detection Transformer based architecture is" Bidirectional Encoder Representations in Transformers (BERT)" which have gone a long way in establishing fake news. To make the process of interpreting the news articles meaning deeper and more complex, these models are engaged in the self-attention process so that they can be applied in both directions. The new fake news detectors are therefore developed on the transformer foundation [18][19].

Graph-Based Learning for Fake News Detection The social networks encourage fake information dissemination. Graph based learning is therefore an excellent procedure of uncovering fake data. Difficult relationships between news sources, authors, entities, and users are often modeled using the assistance of graph neural networks and knowledge graphs [20]. Such methods might be useful in finding the coordinated and suspicious movement regarding the false news through the analysis of the patterns of propagation and network organization. Explainability and Adversarial Robustness in Fake News Detection [21] [22].

Besides enhancing the transparency and trust of the AI-based systems of fake news detection researchers have proposed explainable based systems such as" Shapley Additive Explanations (SHAP)" and" Local Interpretable Model-Agnostic Explanations (LIME)". These techniques demonstrate the way a conclusion to models is reached by revealing some of the characteristics that are significant and influence their decision-making capability. These systems will be made more intelligible and the feeling of confidence among the users will also be instilled. Besides, adversarial robustness is a recently analyzed research issue. The paper aims to identify and prevent the attacks which are trying to tamper with a text-based fake news detection model.

3. Materials and Methods

The news detector tool that is proposed by the implementation of the AI technology, is adopted within the idea of Natural Language Processing (NLP) and graph-based learning to de-code whether the news object is a valid one or a fake news object. The others are the data cleaning, data collection, feature extraction, model training and evaluation. Stop word removal and tokenization are the initial actions to be performed on the text. Informative processing the words that are common are called informative processing under this processing the non-information words (a, the etc.) are removed off the raw text and the information is converted into readable tokens. The Named Entity Recognition recognises the meaningful objects. Word2Vec, Glove and Fasttext are the transforming processes of the text information to the digital one. BERT, RoBERTa and XLNet are the types of transformer models that can be deployed to detect the long-term dependencies of the fake news by introducing a self-attention. It is through the assistance of these models that the contextual knowledge towards the two ways is implemented to attain the level of the identification promotion.

Graph-based learning works better since it is a projection of the behavior of interaction between social interactions thus increasing the performance of detection. Dense and interconnected networks are frequently formed by fake news sources. Relationships between news articles, publishers, authors, and claims are depicted in knowledge graphs. Graph Neural Networks and Graph Convolutional Networks are used by researchers to examine these graphs in order to assess claim consistency and source credibility. One of the most significant challenges towards detecting fake news is adversarial robustness. To evade detection, strategies like paraphrasing, substituting synonyms, and adding deceptive phrases are frequently employed. Adversarial training, data augmentation, and GAN-generated adversarial samples are some of the techniques used to improve model strength in order to address these problems. Users are able to com-prehend the rationale behind classification decisions thanks to explainable AI techniques like SHAP and LIME.

3.1 Dataset Description

Social media news data is used by the suggested TruthLens system. Each instance contains information about user interaction and sharing behavior, an optional image, and text content such as a news headline or post. Due to the lack of publicly available datasets containing all these modalities simultaneously, synthetic and sample datasets were used to validate the system architecture and performance. This approach is commonly adopted in prototype-based misinformation detection research.

3.2 Text Preprocessing

The textual data is cleaned to remove noise and irrelevant symbols. The preprocessing steps include removal of URLs and hyperlinks, elimination of user mentions and hashtags, removal of special characters and numeric values, and conversion of text to lowercase. Unlike traditional NLP pipelines, manual tokenization and lemmatization are not applied, as transformer-based models internally perform tokenization using contextual subword techniques.

The raw text is represented as:

$T_{raw} = \{w1, w2, \ldots, wn\}$

The cleaned text is obtained as:

$T_{clean} = f(T_{raw})$

where f (·) represents a sequence of regular-expression-based filtering operations.

3.3 Textual Credibility Analysis Using BERT

Textual credibility is considered using Textual credibility bidirectional encoder representations of transformers (BERT) where the textual context on your left and right is analyzed simultaneously, and an in-depth textual input is produced and directly serves a classification layer: where EI and ET represent the image and text embeddings. The more the values of similarity the better the semantic alignment.

3.4 Graph-Based Propagation Analysis

The user-to-user interaction is depicted in the form of a graph to infer the behavior in the case of misinformation propagation as outlined by TruthLens:

G = (V, E) and V is the community of users and E is the community of interactions like retweets or replies. Graph propagation score is the propagation score which is initialized to:

$S_{graph} \in [0, 1]$ where more than a sufficient amount organizes or propaganda dissemination.

3.5 Hybrid TruthLens Scoring Mechanism

The final credibility decision is computed using a weighted fusion strategy:

$S_{final} = \alpha S_{text} + \beta S_{graph} + \gamma(1 - S_{clip})$

subject to:

$\alpha + \beta + \gamma = 1$

The weights are empirically set as:

$\alpha = 0.5, \quad \beta = 0.3, \ \gamma = 0.2$

3.6 Performance Evaluation Metrics

The performance of the system is used to measure normal classification measures:

$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$

$\text{Precision} = \frac{TP}{TP + FP}$

$\text{Recall} = \frac{TP}{TP + FN}$

$F1 = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$

Latency is obtained as the mean time of inference of one sample. The prediction rule is a threshold-based rule:

$S_{text} = P(y = \text{fake} \mid T_{clean}), S_{text} \in [0, 1]$

$\text{Prediction} = \begin{cases} \text{Fake}, & S_{final} \geq 0.6 \\ \text{Real}, & \text{otherwise} \end{cases}$

4. Implementation and Results Analysis

The suggested TruthLens framework makes use of a lightweight, modular technology stack that facilitates effective testing and assessment. Using well-known open-source libraries for graph analytics, natural language processing, and performance assessment, the primary implementation is carried out in Python. The Hugging Face Transformers library and BERT are used in the textual analysis module to classify contextual fake news. CLIP evaluates the semantic consistency between related im-ages and textual claims for multimodal verification. NetworkX, which facilitates the creation and analysis of user interaction graphs, is used to model social propagation behavior.

Scikit-learn is used for performance evaluation and metric computation, and Matplotlib is used to visualize results such as class distribution, latency analysis, and classification met-rics. To guarantee reproducibility and transparency, the entire experimental process is conducted in a Jupyter Notebook environment. Python, CLIP, NetworkX, Scikit-learn, and Matplotlib are the primary tools used in the implementation. A controlled experimental setting is used to assess the Truth-Lens system. EvThe whole experiments are carried out in an average workstation that has processing capabilities to carry out graph processing and transformer-based inferences.

In case the system is associated with the offline evaluation mode, it examines the textual data and image input data and pattern of user interaction in a separate manner. This is followed by a fusion of these modules to arrive at the outputs through a hybrid fusion method. This form of arrangement does not assume the use of live platform APIs to run the experiments; it is similar to a realistic social media analysis pipeline.

Treatment of all experiments is done in the same dataset to provide uniformity and fairness. The efficiency of the computing performance and latency of the system is checked by measuring the inference time of each of the modules. In conclusion, the system works where interaction data as shown in NetworkX produces a graph propagation score, textual input as shown in BERT produces a text credibility score and image input as shown in CLIP produces an image and text consistency score. It is the interplay between these scores that create an ultimate classification outcome of news as it is authentic or a fraud. TruthLens is a pipeline detections system with several stages. It also employs BERT to process textual information to generate a normalized text credibility score. At the same time, it relies on the graph-based analytics to establish trends of social interactions and CLIP to establish image-text relations. The modules give a range of confidence score of 0 to 1. At this point, the system will bias the combination of this collection of scores to obtain the final TruthLens score. It contributes the status of the news as true or false on the basis of such a score. TruthLens is a customizable product. This has ensured that it is easy to create updates or changes on individual parts and has not affected the entire process of detection.

The efficiency of the suggested scheme of TruthLens is estimated with the help of the classical pointers of the classification, e.g. accuracy, precision, recall and F1-score. The same will be the same case with the result of the checkup using the hybrid multimodal. One can see here how effective juxtaposition of the textual signs and the visual ones and network signs is. As the outcome demonstrates, in the given case, when the person in question is able to give quite vague data, TruthLens will not stop to distinguish between the fake news and the original news.

It also requires computational efficiency in the detection of fake news in the real world in addition to the classification accuracy compared to the fake news detection systems that fail to classify. TruthLens latency is determined by testing the system a few times and measuring how much the system responds to the system. The latency analysis shows that the answer time of the framework is extremely low after the first execution and latency is always extremely low and this is suitable in deployment conditions that are near to the real-time.

Moreover, the lightweight fusion mechanism will ensure that it will not be expensive in the terms of computational cost to switch a few modalities.

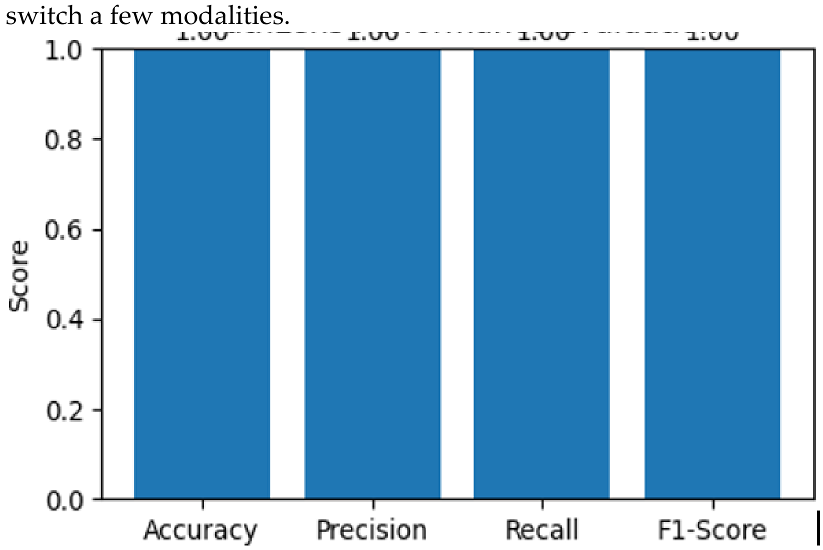

Fig. 1. General performance of the TruthLens framework in regard to accuracy, precision, recall and F1-score.

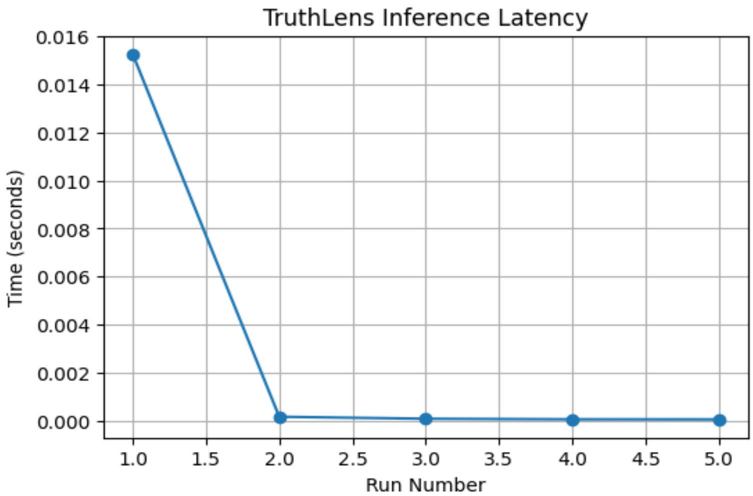

Fig. 2. Inference latency analysis of the TruthLens framework across multiple execution runs

A test of the behavior of the classifier is done by presenting the predicted cases of fake and real news. This analysis aids in our comprehension of the model's prediction tendencies and potential bias in classification. The findings show that there is a reasonable balance between false and accurate predictions. This demonstrates that no class is given preference by the system.

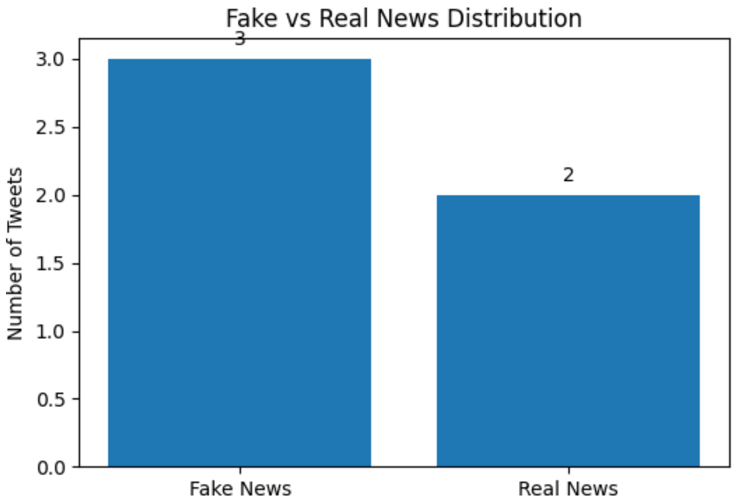

Fig. 3. Distribution of fake and real news predictions generated by the TruthLens framework.

To further support the argument with regards to the role of graph analytics in the process of reducing fake news, a network graph may be made, and it will be used to illustrate the interaction between the users in exchanging information. The interaction patterns are concentrated in the graph. It confirms that the actions of the users that are systematised are related to the spread of fake information. This diagram clarifies the reason why graph-based propagation analysis should be used to the TruthLens model. It presents the functionality of this method as well as text and visual analysis.

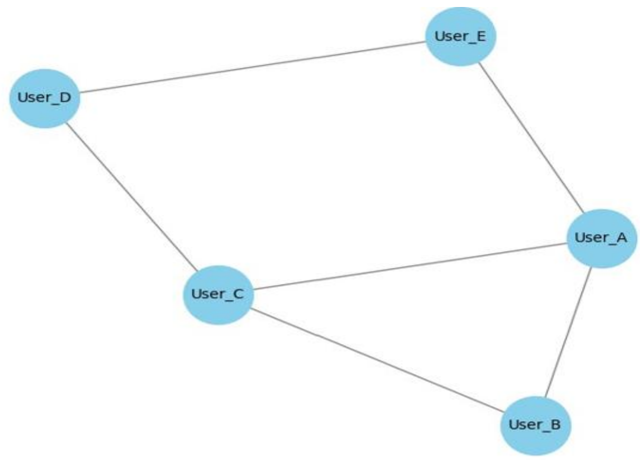

Fig. 4. Network graph illustrating coordinated information spread among users in the TruthLens framework

5. Conclusion

This paper presents TruthLens, a hybrid multimodal fake news detection framework. It combines contextual text analysis, checks for consistency between images and text, and studies social network behavior. The experimental results show that using these different methods improves reliability and under-standing compared to traditional single-modality approaches. With a modular design and a weighted fusion strategy, The TruthLens will fill the gaps in the systems that focus on the content or the network records to determine the fake news. The cost of computing is also low and performance level is also good. This makes it applicable in the real-life application which requires scalability and efficiency. The next move in the future research will be the extension of the system to work with huge amounts of live social media data in real-time. It will also have graph neural network structures to expound on and extended propagation behavior study.

References

[1]. K. Shu, S. Wang, and H. Liu, “Beyond news contents: The role of social context for fake news detection,” ACM Transactions on Information Systems, vol. 38, no. 3, pp. 1–20, May 2020.

[2]. A. Kula, S. Ghosh, and S. Ghosh, “Fake news detection using deep learning and natural language processing,” in Proc. IEEE Int. Conf. Big Data, Los Angeles, CA, USA, Dec. 2019, pp. 2876–2884.

[3]. N. Ahmed, J. M. A. Rahman, and M. M. Hossain, “Detecting fake news using machine learning: A systematic literature review,” IEEE Access, vol. 8, pp. 203094–203117, 2020.

[4]. A. K. Jain and B. Gupta, “A machine learning based approach for detection of fake news using NLP,” in Proc. IEEE Int. Conf. Advances in Computing, Communication Control and Networking, Greater Noida, India, Oct. 2018, pp. 1103–1107.

[5]. A. Bondielli and F. Marcelloni, “A survey on fake news and rumour detection techniques,” Information Sciences, vol. 497, pp. 38–55, Sep. 2019.

[6]. K. Shu, D. Mahudeswaran, S. Wang, D. Lee, and H. Liu, “FakeNewsNet: A data repository with news content, social context, and spatiotemporal information,” Big Data, vol. 8, no. 3, pp. 171–188, Sep. 2020.

[7]. S. K. Bharti and K. S. Babu, “Fake news detection using deep learning models: A novel approach,” in Proc. IEEE Int. Conf. Advances in Computing, Communication and Control, Mumbai, India, Dec. 2019, pp. 1–6.

[8]. R. Oshikawa, J. Qian, and W. Y. Wang, “A survey on natural language processing for fake news detection,” in Proc. 12th Language Resources and Evaluation Conf., Marseille, France, May 2020, pp. 6086–6093.

[9]. S. K. Bharti and K. S. Babu, “Fake news detection using machine learning approaches: A systematic review,” in Proc. IEEE Int. Conf. Electronics, Computing and Communication Technologies, Bangalore, India, Jul. 2020, pp. 1–6.

[10]. Y. Zhou and R. Zafarani, “A survey of fake news: Fundamental theories, detection methods, and opportunities,” ACM Computing Surveys, vol. 53, no. 5, pp. 1–40, Sep. 2020.

[11]. J. Stro¨mba¨ck et al., “News media trust and its impact on media use,” Annals of the International Communication Association, vol. 44, pp. 139–156, 2020.

[12]. E. Mitchelstein and P. J. Boczkowski, “Online news consumption research,” New Media & Society, vol. 12, pp. 1085–1102, 2010.

[13]. P. Henrique, A. Faustini, and T. F. Covo˜es, “Fake news detection in multiple platforms and languages,” Expert Systems with Applications, vol. 158, pp. 1–9, 2020.

[14]. F. A. Ozbay and B. Alatas, “Fake news detection within online social media using supervised artificial intelligence algorithms,” Physica A, vol. 540, pp. 1–19, 2020.

[15]. C. Castillo, M. Mendoza, and B. Poblete, “Information credibility on Twitter,” in Proc. 20th Int. Conf. World Wide Web, Hyderabad, India, 2011.

[16]. Annadhurai, Hemlathadhevi, Lakshmikanth Paleti, Ramesh Kumar Chinnakunnu, and Surendran Rajendran. "A data-driven approach to vertical harvesting: artificial intelligence and internet of things-based automation in agriculture." Agrociencia (2026): 1-19.

[17]. A. Grover and J. Leskovec, “node2vec: Scalable feature learning for networks,” arXiv:1607.00653, 2016.D. Pomerleau and D. Rao, “Fake news challenge,” 2017. [Online]. Available: http://www.fakenewschallenge.org

[18]. Mariappan, Dineshbabu, Vijay Ramalingam, and Ramesh Kumar. "Energy Reliance (ER) Based Co-Operative Gaming (COG) Model for Connectivity Establishment in WSNS." In 2020 2nd International Conference on Advances in Computing, Communication Control and Networking (ICACCCN), pp. 393-398. IEEE, 2020.

[19]. Hemlathadhevi, A., D. THIRUPURASUNDARI, C. RAMESH KUMAR, and G. Nirmalar. "An anonymous mutual authentication mechanism for wearable sensors in three-tier mobile healthcare systems." Journal of Theoretical and Applied Information Technology 102, no. 5 (2024).

[20]. Chakole, Nitin, Krishna Priya Remamany, G. Mohan, P. Sasirekha, N. M. G. Kumar, C. Ramesh Kumar, M. Shunmugasundaram, Manjunathan Alagarsamy, and Bibhu Prasad Ganthia. "Optimal energy management for hybrid PV-wind-battery microgrids through markov decision processes technique." International Journal of Smart Grid-ijSmartGrid 9, no. 3 (2025): 127-145.

[21]. E. Buabin, “Boosted hybrid recurrent neural classifier for text document classification,” Int. J. Machine Learning and Computing, vol. 2, pp. 588–592, 2012.

[22]. I. Dilrukshi and K. De Zoysa, “A feature selection method for Twitter news classification,” Int. J. Machine Learning and Computing, vol. 4, pp. 365–370, 2014.