iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** II **Month of publication:** February 2026

DOI: https://doi.org/10.22214/ijraset.2026.77345

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue II Feb 2026- Available at www.ijraset.com

Detection and Classification of Fake News Using Natural Language Processing and Deep Learning

Dr. P.K. Sharma[1], Mr. Manvendra Singh Divakar[2], Ritu Lodhi[3]
[1]Principal, [2]Assistant Professor, [3]Research Scholar, NRI Institute of Research & Technology, Bhopal (M.P)

Abstract: The exponential growth of digital media platforms and online information dissemination has significantly amplified the spread of fake news, posing serious threats to public trust, democratic processes, social harmony, and informed decision-making. Fake news is intentionally designed to mislead readers by mimicking the linguistic structure, narrative style, and presentation of legitimate journalism, making its automated detection a complex and evolving challenge. In recent years, Natural Language Processing (NLP) and deep learning techniques have emerged as powerful tools for addressing this problem by enabling large-scale, automated analysis of textual content. This review paper presents a comprehensive synthesis of contemporary research on fake news detection, with particular emphasis on NLP-driven feature extraction and deep learning-based classification models. Drawing upon dissertation-based empirical insights and recent scholarly contributions, the review examines linguistic characteristics of fake news, traditional and deep learning-based detection approaches, word embedding techniques, evaluation strategies, and ethical considerations. The analysis highlights the effectiveness of sequential models such as Long Short-Term Memory networks in capturing contextual and narrative inconsistencies inherent in deceptive content. Furthermore, the review identifies persistent challenges related to generalization, dataset bias, interpretability, and real-world deployment. The paper concludes that NLP and deep learning-based fake news detection systems represent a robust and scalable solution for combating misinformation, provided they are developed with methodological rigor, ethical responsibility, and human oversight.
Keywords: Fake News Detection, Natural Language Processing, Deep Learning, LSTM, Text Classification, Misinformation.

I. INTRODUCTION

The rapid expansion of digital media ecosystems has fundamentally transformed the production, distribution, and consumption of news and information. Online news portals, social media platforms, blogs, video-sharing sites, and instant messaging applications have enabled information to be disseminated to a global audience within seconds. This digital transformation has significantly democratized access to information, empowering individuals to consume, share, and create content without traditional gatekeeping mechanisms. However, alongside these benefits, the digital information landscape has also become increasingly vulnerable to the large-scale spread of fake news—deliberately fabricated or misleading content presented in the stylistic and structural form of legitimate journalism. The ease of content creation and rapid virality mechanisms embedded in digital platforms have made fake news a pervasive and persistent challenge. The proliferation of fake news has emerged as a serious societal concern with wide-ranging consequences. In political contexts, fake news has been shown to influence voter behavior, distort public discourse, and undermine trust in democratic institutions. In the domain of public health, misinformation related to diseases, treatments, and vaccines can lead to panic, mistrust, and harmful decision-making. Similarly, in financial and economic settings, false or manipulated information can impact stock markets, damage organizational reputations, and mislead investors. Beyond these tangible impacts, the continuous exposure to fake news erodes public confidence in credible journalism and contributes to social polarization by reinforcing ideological echo chambers. The unprecedented speed and scale at which misinformation spreads in digital environments have rendered traditional manual verification and fact-checking approaches inadequate, as they are often time-consuming and unable to cope with the sheer volume of online content generated daily.

Fake news is particularly challenging to detect because it is not merely incorrect information but is intentionally crafted to deceive. It often exploits journalistic conventions such as professional headlines, structured narratives, and authoritative language, making it difficult for both readers and automated systems to distinguish from genuine news. Additionally, fake news frequently employs emotionally charged language, sensational claims, and persuasive storytelling techniques designed to capture attention and encourage sharing. In fast-paced digital environments, users often consume and share content with minimal scrutiny, further amplifying the spread of deceptive information. These characteristics make fake news detection a complex and non-trivial problem that extends beyond simple fact verification or keyword matching.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue II Feb 2026- Available at www.ijraset.com

In response to these challenges, automated fake news detection has become a critical research area within artificial intelligence, particularly at the intersection of Natural Language Processing and deep learning.

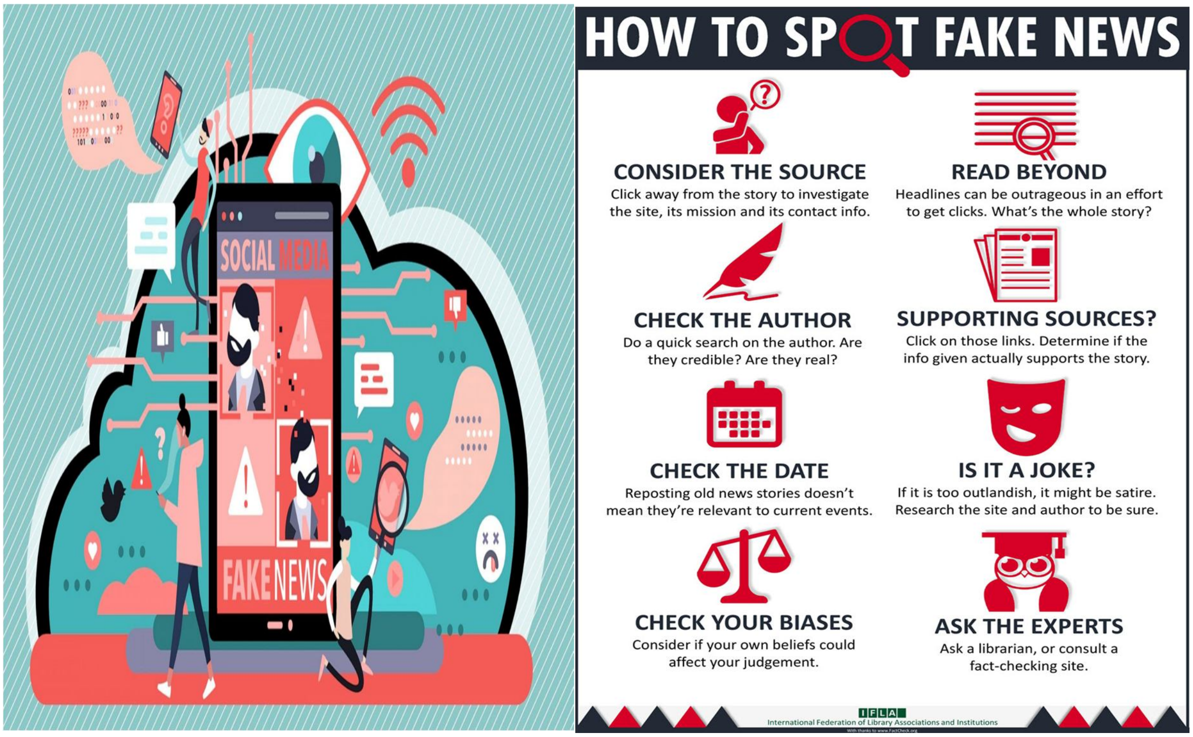

Figure 1.1: Illustration showing the spread and societal impact of fake news across digital platforms.

NLP techniques enable machines to analyze textual data at multiple linguistic levels, including lexical, syntactic, semantic, and discourse dimensions. When combined with deep learning models, NLP-based systems can automatically learn complex patterns and representations from large volumes of news text. Deep learning architectures, especially those designed for sequential modeling, have demonstrated strong potential in capturing contextual dependencies and narrative inconsistencies that are characteristic of deceptive content. Unlike traditional rule-based or shallow machine learning approaches, these models reduce reliance on handcrafted features and offer greater adaptability to evolving misinformation strategies. This review paper examines the convergence of fake news detection research with advances in Natural Language Processing and deep learning. It synthesizes existing literature to analyze how computational techniques have been applied to identify linguistic patterns, semantic distortions, and contextual cues in news articles. The review further highlights key methodological trends, evaluation practices, and persistent challenges associated with automated fake news detection. By consolidating current knowledge and identifying research gaps, this paper aims to provide a comprehensive understanding of the state of the art and to guide future research toward developing robust, interpretable, and ethically responsible fake news detection system.

II. LINGUISTIC CHARACTERISTICS AND SOCIETAL IMPACT OF FAKE NEWS

Fake news exhibits distinct linguistic and stylistic characteristics that differentiate it from genuine journalism, although these differences are often subtle. Studies consistently report that fake news articles tend to employ sensational headlines, emotionally charged language, exaggerated claims, and simplified narrative structures to attract attention and encourage sharing. Such content often appeals to fear, anger, or curiosity, exploiting cognitive biases and emotional triggers. In contrast, legitimate news articles typically emphasize factual consistency, balanced reporting, and neutral tone. However, sophisticated fake news increasingly blurs this distinction, making automated detection a non-trivial task. The societal impact of fake news is profound. In political contexts, misinformation has been shown to influence elections, polarize public opinion, and erode trust in democratic institutions. During public health crises, fake news can spread dangerous misinformation regarding treatments, vaccines, or disease origins, leading to panic and harmful behaviors. Financial misinformation can manipulate markets and damage reputations. These consequences underscore the urgency of developing reliable, scalable, and objective fake news detection systems capable of operating in real-world digital environments.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue II Feb 2026- Available at www.ijraset.com

Table 1: Summary of Recent Studies on Fake News Detection Using NLP and Deep Learning

Author(s) & Year	Dataset / Domain	Method Used	Key Findings
Alshuwaier & Alsulaiman (2025)	Online news articles	ML & DL review	Deep learning outperforms traditional ML
Roumeliotis et al. (2025)	News & social media	CNN, LSTM, LLMs	LSTM shows strong contextual learning
Embarak (2025)	Facebook data	LSTM-based DL	Effective sequential text modeling
Cavus et al. (2024)	Social networks	Hybrid NLP–DL	Real-time detection feasibility
Yakkundi (2025)	Social media text	ML classifiers	NLP improves detection accuracy
Kumari (2024)	Multimodal news	Deep learning	Contextual features improve robustness

III. ROLE OF NATURAL LANGUAGE PROCESSING IN FAKE NEWS DETECTION

Natural Language Processing forms the foundation of automated fake news detection by enabling machines to process, analyze, and interpret textual data. NLP techniques transform raw text into structured representations that capture lexical, syntactic, semantic, and discourse-level features. Early NLP-based approaches focused on surface-level features such as word frequency, n-grams, and sentiment polarity. While these features provided initial discriminatory power, they were limited in their ability to capture deeper contextual meaning and narrative coherence. Advances in NLP introduced word embedding techniques that represent words as dense vectors in continuous semantic space, preserving contextual similarity and linguistic relationships. Embedding-based representations significantly improved fake news detection by enabling models to identify subtle semantic distortions, unusual word associations, and inconsistencies characteristic of deceptive content. Additionally, discourse-level analysis and sentiment modeling have been employed to assess narrative structure, coherence, and emotional intensity. Collectively, NLP techniques provide the essential pipeline for converting unstructured news text into meaningful inputs for machine learning and deep learning models.

IV. DEEP LEARNING APPROACHES FOR FAKE NEWS CLASSIFICATION

Deep learning has revolutionized fake news detection by enabling models to automatically learn hierarchical and contextual representations from textual data without extensive manual feature engineering. Among deep learning architectures, Recurrent Neural Networks—particularly Long Short-Term Memory networks—have demonstrated strong performance due to their ability to model sequential dependencies and long-range context in text. Fake news often embeds deception across multiple sentences or paragraphs, making sequential modeling critical for effective detection.

LSTM-based models, when combined with word embeddings, capture both semantic meaning and narrative flow, allowing detection of inconsistencies and deceptive cues distributed throughout an article. Hybrid architectures incorporating Convolutional Neural Networks and attention mechanisms further enhance performance by capturing local linguistic patterns and highlighting informative textual segments. Although transformer-based models have achieved state-of-the-art results, their computational complexity limits practical deployment, reinforcing the relevance of efficient LSTM-based frameworks for real-world applications.

V. EVALUATION METRICS AND MODEL RELIABILITY

The evaluation of fake news detection systems demands a comprehensive and multidimensional assessment strategy that extends well beyond overall classification accuracy. While accuracy provides a general measure of how often a model produces correct predictions, it can be misleading when used in isolation, particularly in misinformation detection tasks where class distributions may be imbalanced or the societal consequences of misclassification are substantial. In real-world digital media environments, even a small proportion of undetected fake news can lead to widespread misinformation, while incorrect labeling of legitimate journalism can undermine public trust. Consequently, rigorous evaluation frameworks incorporating multiple performance metrics are essential for assessing both the technical effectiveness and practical reliability of fake news detection models. Precision and recall are among the most critical metrics in this context, as they provide complementary perspectives on model behavior. Precision measures the proportion of news articles classified as fake that are genuinely fake, reflecting the reliability and trustworthiness of positive predictions. High precision is essential to avoid false positives, where legitimate news is incorrectly flagged as fake, potentially leading to censorship concerns and erosion of journalistic credibility.

Recall, on the other hand, measures the proportion of actual fake news articles that are correctly identified by the model. In fake news detection, recall is often considered particularly important because false negatives—cases where fake news is misclassified as real—allow deceptive content to spread unchecked, posing risks to public discourse and decision-making. An effective detection system must therefore achieve a careful balance between precision and recall, avoiding excessive bias toward either metric. The F1-score, which represents the harmonic mean of precision and recall, is widely used to capture this balance in a single measure. In review studies of fake news detection, the F1-score is often regarded as a more informative metric than accuracy alone, as it accounts for both types of classification errors. Models that achieve high F1-scores are generally considered robust, as they demonstrate consistency in detecting fake news while minimizing incorrect labeling of real news. However, even the F1-score should be interpreted in conjunction with class-specific metrics, as overall values may obscure performance disparities between fake and real news categories.

Confusion matrix analysis provides deeper insight into model reliability by explicitly illustrating the distribution of true positives, true negatives, false positives, and false negatives. This analysis enables researchers to examine error patterns and identify systematic weaknesses in detection systems. For example, a high number of false negatives may indicate that the model struggles with sophisticated fake news that closely resembles legitimate journalism, while excessive false positives may suggest over-sensitivity to sensational language. Such insights are critical for refining model architecture, feature representation, and decision thresholds to better align with real-world requirements. In addition to classification metrics, training–validation performance analysis plays a vital role in assessing model reliability and generalization capability. Monitoring training and validation accuracy and loss curves helps identify overfitting, underfitting, and unstable learning behavior. A reliable fake news detection model should exhibit stable convergence, with validation performance closely tracking training performance over time. Significant divergence between training and validation metrics may indicate that the model has memorized training data rather than learned generalizable patterns. Techniques such as dropout regularization, early stopping, and stratified validation are commonly employed to mitigate overfitting and enhance robustness. Overall, comprehensive evaluation using multiple metrics and performance analyses is essential for establishing the reliability of fake news detection systems. Such rigorous assessment ensures that proposed models are not only technically sound but also suitable for deployment in dynamic, high-stakes digital information environments.

VI. ETHICAL CONSIDERATIONS AND HUMAN-CENTRIC DEPLOYMENT

Ethical responsibility occupies a central position in the design, development, and deployment of fake news detection systems, particularly as such technologies increasingly influence public discourse and information access. Unlike many technical applications, fake news detection operates in socially, politically, and culturally sensitive contexts, where algorithmic decisions can have far-reaching consequences. Automated systems must therefore address critical ethical concerns related to data privacy, algorithmic bias, transparency, and potential misuse. Failure to consider these dimensions risks undermining public trust, infringing on freedom of expression, and reinforcing existing inequalities in information ecosystems. Data privacy is one of the foremost ethical challenges in fake news detection research. Many detection systems rely on large-scale textual datasets collected from online news platforms and social media, which may contain sensitive or personally identifiable information. Ethical research practice requires that such data be sourced responsibly, anonymized where necessary, and processed in compliance with data protection regulations. Even when publicly available data is used, researchers must ensure that models do not inadvertently expose private information or enable invasive surveillance practices. Privacy-preserving data handling and transparent disclosure of data sources are therefore essential components of ethical fake news detection frameworks. Algorithmic bias represents another significant concern, as detection models are inherently influenced by the data on which they are trained. Datasets constructed from politically or culturally skewed sources may encode implicit biases that lead models to disproportionately flag certain viewpoints, topics, or communities as deceptive. Such bias can result in unfair treatment of minority voices or the suppression of legitimate dissent. Ethical deployment of fake news detection systems requires careful dataset curation, bias evaluation, and ongoing monitoring to ensure equitable performance across diverse content and user groups. Without these safeguards, automated systems risk amplifying existing power imbalances rather than mitigating misinformation.

Transparency and interpretability are closely linked to ethical accountability. Deep learning models, despite their strong performance, are often criticized for operating as “black boxes” that provide little insight into how decisions are made. In the context of fake news detection, opaque decision-making can erode trust among journalists, policymakers, and the public, particularly when content is flagged or suppressed without clear justification. As a result, explainable artificial intelligence techniques have gained increasing attention in the literature. Methods such as attention visualization, feature attribution, and post-hoc explanation frameworks aim to provide human-interpretable insights into model behavior.

While full interpretability remains challenging, these techniques represent important steps toward transparent and accountable detection systems. Equally important is the recognition that fake news detection systems should not function as autonomous arbiters of truth. Human-centric deployment emphasizes the integration of automated models within decision-support frameworks that preserve human judgment and contextual understanding. Journalists, content moderators, and policymakers possess domain knowledge, ethical reasoning, and situational awareness that cannot be fully replicated by algorithms. Automated systems are best positioned to perform large-scale screening and prioritization of potentially deceptive content, while final decisions and interventions remain under human control. This human-in-the-loop approach reduces the risk of erroneous or ethically problematic outcomes and aligns with broader principles of responsible artificial intelligence. In summary, ethical considerations are not peripheral but foundational to fake news detection research. Addressing privacy, bias, transparency, and human oversight is essential for ensuring that detection systems are deployed responsibly and gain public acceptance. A human-centric approach that combines computational efficiency with ethical accountability offers the most promising path toward combating misinformation while safeguarding democratic values and public trust.

VII. RESEARCH CHALLENGES AND EMERGING DIRECTIONS

Despite substantial progress in fake news detection using Natural Language Processing and deep learning techniques, several critical research challenges continue to limit the robustness, generalizability, and real-world applicability of existing systems. One of the most persistent challenges is dataset bias, which arises from the way fake news datasets are constructed and labeled. Many widely used datasets are curated from specific fact-checking websites, political contexts, or media ecosystems, which may not fully represent the diversity of real-world misinformation. As a result, models trained on such datasets often learn dataset-specific patterns rather than generalizable indicators of deception. This bias can lead to skewed predictions when detection systems are applied to news content from different regions, topics, or ideological perspectives, raising concerns about fairness and reliability. Generalization across domains and time presents another major challenge in fake news detection research. Fake news is a dynamic phenomenon that evolves in response to social events, technological changes, and detection mechanisms. Linguistic styles, narrative strategies, and thematic focus of misinformation continuously adapt, leading to concept drift over time. Models trained on historical data may therefore experience significant performance degradation when confronted with newly emerging misinformation patterns. While deep learning models exhibit strong learning capacity, most existing approaches rely on static training paradigms that assume stable data distributions. Addressing temporal drift requires adaptive learning strategies, such as continual learning, incremental model updating, and periodic retraining, while carefully mitigating risks such as catastrophic forgetting. Annotation subjectivity further complicates fake news detection research. The distinction between fake and real news is not always clear-cut, as news articles may contain partially true information, misleading framing, or context-dependent interpretations. Human annotators may disagree on labeling decisions, introducing noise and inconsistency into training data. Binary classification frameworks, although practical, often oversimplify the complex spectrum of misinformation. This subjectivity affects both model training and evaluation, as performance metrics may not fully capture nuanced misclassification scenarios. Future research may benefit from exploring multi-label or graded credibility classification schemes that better reflect real-world information ambiguity.

Another significant limitation highlighted in the literature is the lack of standardized evaluation protocols and benchmarking practices. Studies often employ different datasets, data splits, and performance metrics, making it difficult to compare results across models and research efforts. Some works emphasize accuracy, while others prioritize recall or F1-score, often without consistent justification. This inconsistency hampers reproducibility and obscures true progress in the field. Establishing standardized benchmarks, evaluation metrics, and reporting guidelines is essential for advancing fake news detection research in a systematic and transparent manner. From a practical deployment perspective, computational efficiency and scalability remain critical concerns. While advanced deep learning architectures, including transformer-based models, achieve impressive accuracy, their high computational cost and resource requirements limit real-time deployment, particularly in resource-constrained environments. This challenge underscores the need for efficient architectures that balance detection performance with scalability and responsiveness. Lightweight models, optimized training strategies, and hybrid frameworks that combine efficiency with contextual understanding represent promising directions for practical implementation. Emerging research directions increasingly emphasize explainability, multilingual detection, and human-centric design. Explainable artificial intelligence techniques aim to enhance transparency and trust by providing interpretable insights into model decisions. Multilingual and cross-cultural fake news detection is essential for global applicability, as misinformation transcends linguistic and geographic boundaries. Finally, integrating automated detection systems with human oversight through hybrid human–AI frameworks offers a pragmatic path forward. Such approaches combine the scalability of deep learning with human judgment and ethical reasoning, addressing many of the limitations identified in current research. In summary, while NLP and deep learning have significantly advanced fake news detection, unresolved challenges related to bias, generalization, adaptability, evaluation, and deployment remain. Addressing these challenges through adaptive, interpretable, and scalable frameworks will be crucial for developing effective and responsible fake news detection systems in the future.

VIII. CONCLUSION

This review paper has presented a comprehensive synthesis of dissertation-based insights and contemporary scholarly research on fake news detection using Natural Language Processing and deep learning techniques. The analysis highlights the growing significance of automated misinformation detection in the context of rapidly evolving digital media ecosystems, where the volume, velocity, and influence of online information far exceed the capacity of traditional verification mechanisms. By examining linguistic characteristics, methodological developments, and evaluation practices, this review demonstrates that NLP-driven deep learning models represent a robust and scalable approach to addressing the complex challenge of fake news detection. A central conclusion of this review is that deep learning architectures, particularly those designed for sequential modeling such as Long Short-Term Memory networks, are highly effective in capturing contextual and narrative dependencies present in news articles. Fake news often embeds deceptive cues across multiple sentences or paragraphs rather than relying on isolated false statements. LSTM-based models, supported by word embeddings and structured preprocessing pipelines, are well suited to identify such long-range dependencies and subtle semantic inconsistencies. The empirical findings reported across studies consistently indicate that these models outperform traditional machine learning approaches that rely on handcrafted linguistic features, especially in terms of generalization and adaptability. However, the review also emphasizes that achieving high classification accuracy alone is insufficient for real-world deployment in socially sensitive domains. Interpretability and transparency emerge as critical requirements for responsible fake news detection systems. Many deep learning models operate as black boxes, making it difficult for stakeholders to understand or trust their decisions. This lack of explainability can hinder adoption by journalists, policymakers, and content moderators who require clear justifications for classification outcomes. Consequently, future research must place greater emphasis on explainable artificial intelligence techniques that provide insight into model reasoning, such as attention visualization, feature importance analysis, or interpretable hybrid frameworks. Equally important is the integration of ethical safeguards to address concerns related to data privacy, algorithmic bias, and potential misuse of automated detection systems.

The review further underscores the necessity of human oversight in fake news detection. Automated systems should not be positioned as definitive arbiters of truth but rather as decision-support tools that assist human experts in managing large volumes of content. Human–AI collaboration allows computational models to perform rapid screening and prioritization, while human judgement provides contextual understanding, ethical reasoning, and accountability. Such hybrid frameworks align with broader principles of responsible artificial intelligence and are particularly important in politically and socially sensitive contexts where misclassification can have serious consequences. Looking ahead, future research should prioritize the development of adaptive and resilient fake news detection frameworks capable of addressing concept drift and evolving misinformation strategies. Fake news creators continuously modify linguistic styles, narrative structures, and dissemination tactics to evade detection, necessitating models that can learn and update over time without sacrificing previously acquired knowledge. Expanding detection systems to support multilingual and cross-domain analysis is also essential for enhancing global applicability. Additionally, standardized datasets, evaluation protocols, and benchmarking practices are needed to ensure reproducibility and meaningful comparison across studies. In conclusion, this review affirms that NLP and deep learning-based fake news detection systems hold significant promise for mitigating the societal impact of misinformation. When designed with methodological rigor, interpretability, ethical responsibility, and human-centric principles, these systems can play a vital role in safeguarding public trust, supporting informed decision-making, and promoting a healthier digital information ecosystem.

REFERENCES

[1] Alshuwaier, F. A., & Alsulaiman, F. A. (2025). Fake News Detection Using Machine Learning and Deep Learning Algorithms: A Comprehensive Review and Future Perspectives. Computers, 14(9), 394.
[2] Lv, J., Gao, Y., & Li, L. (2025). Multi-modal fake news detection: A comprehensive survey on deep learning technology, advances, and challenges. Journal of King Saud University – Computer and Information Sciences.
[3] Roumeliotis, K. I., Tselikas, N. D., & Nasiopoulos, D. K. (2025). Fake News Detection and Classification: A Comparative Study of CNNs, LLMs, and NLP Models. Future Internet, 17(1), 28.
[4] Yakkundi, S. (2025). Exploring Machine Learning for Fake News Detection. SpringerPlus.
[5] Ayyasamy, R. K. (2025). A hybrid deep learning framework for fake news detection using metaheuristic algorithms. Scientific Reports.

[6] Embarak, O. (2025). Deep Learning for Fake News Detection: Analysing Facebook's Misinformation Networks. Procedia Computer Science.
[7] Jadhav, T. (2025). Fake News Detection on Social Media Using NLP. International Journal of Scientific Research and Science & Technology (IJSRST).
[8] Alkudah, N. M., & Holst, M. (2025). Fake news detection in Arabic media: comparative analysis of ML and DL algorithms. PeerJ Computer Science.
[9] Ahmad, K. S. F., & co-authors. (2025). Hybrid optimization driven fake news detection using modified transformer models. Scientific Reports.
[10] Salve, R. R. (2025). Fake News Detection Using Natural Language Processing: System Design and Evaluation. Thesis/Report.
[11] S Kumari (2024). A Deep Learning Multimodal Framework for Fake News Detection. ETASR.
[12] Singh, J., Liu, F., Xu, H., Ng, B. C., & Zhang, W. (2024). LingML: Linguistic-Informed Machine Learning for Enhanced Fake News Detection. arXiv.
[13] Hoy, N., & Koulouri, T. (2025). Improving Generalisability of Fake News Detection Models. arXiv.
[14] Sittar, A., Golob, L., & Smiljanic, M. (2025). Synthetic News Generation for Fake News Classification. arXiv.
[15] Data Mining in Personalized Education (2024). A review of data mining advancements for personalized learning. arXiv.
[16] Lin, Y., Chen, H., Xia, W. et al. (2025). Deep Learning Techniques in Educational Data Mining: A Comprehensive Survey. Data Science and Engineering.
[17] de la Paz, A. J. & co-authors. (2025). Educational Data Mining and Predictive Modeling in the Age of AI. Computers, 14(2), 68.
[18] Tuanaya, R. et al. (2025). Machine Learning in Educational Data Mining: Trends & Gaps. Journal of Technological Pedagogy and Educational Development.
[19] IJIRT Review Article (2025). Review on Fake News Detection Using Natural Language Processing. IJIRT.
[20] Jadhav et al. (2025). Automated Fake News Detection: A Critical Evaluation of NLP Approaches. IJSRST.