International Innovative Research Journal of Engineering and Technology
Vol. 12, Issue. 1, Sep 2026, pp. 14~26
ISSN: 2456-1983, DOI: 10.32595/iirjet.org/v12i1.2026.252

Intelligent Evidence-Aware Fake News Detection using NLP and Cybersecurity Analytics

Cheah Chee Xiang[1], Kang Harnyi[2]

[1]School Of Digital Technology, Wawasan Open University,
George Town 10050, Penang, Malaysia
E-mail: ccx3_pg@student.wou.edu.my
[2]School Of Digital Technology, Wawasan Open University, George Town 10050, Penang, Malaysia
E-mail: kh3_pg@student.wou.edu.my

Article Info

Article History:

Received Jun 21, 2026
Revised Jul 24, 2026
Accepted Aug 25, 2026

Keywords:

Fake News Detection,
Social Media,
Cybersecurity Analytics,
Natural Language Processing

ABSTRACT

The tremendous growth of social media facilitated the rapid dissemination of information, but it has also accelerated the dissemination of fake news and misleading content. The trustworthiness of supporting evidence or cybersecurity-related signals connected to online content may not be sufficiently taken into account by traditional fake news detection techniques, which frequently concentrate exclusively on textual patterns. This study proposes IEFN-CS (Intelligent Evidence-Aware Fake News Detection) that makes use of cybersecurity analytics and natural language processing (NLP) framework for detecting potential false information in social media environments. The suggested framework is produced and evaluated using accessible to the public, the Kaggle fake-news dataset includes content that has been identified as either fake or real. The framework is designed to analyze social-media content and extract contextual and semantic information. The proposed technique employs Natural Language Processing (NLP) and semantic representation is to examine the meaning and credibility-related characteristics of the content. The framework also includes cybersecurity analytics to examine suspicious URLs, source-related traits, and unusual content-sharing habits, as well as evidence-aware analysis to find supporting or contradicting material. These heterogeneous signals are integrated through an evidence-aware decision to classify content as misleading or fake. The performance of the proposed framework is evaluated using accuracy, precision, recall, F1-score. The proposed IEFN-CS framework aims to provide reliable, evidence-supported, and security-aware fake news detection while improving the transparency and robustness of misinformation identification in dynamic social media environments.

Corresponding Author:

Cheah Chee Xiang,
School Of Digital Technology,
Wawasan Open University,
George Town 10050, Penang, Malaysia,
E-mail: ccx3_pg@student.wou.edu.my

1. INTRODUCTION

The growth of social media and online news platforms [1] has made information available to users within an extremely short period. However, the distribution of false and misleading information has also grown due to this quick dissemination. Fake news can be consciously created or manipulated to appear as authentic information, making it difficult for users to distinguish factual material from false claims.

Natural language processing (NLP) [2] has become an essential technique for analyzing textual material and identifying language patterns associated to misinformation. NLP-based techniques may analyze the words, context, writing style, and semantic features of news articles. Transformer-based language models, such as BERT, have greatly improved text representation by learning contextual relationships between words instead of evaluating them separately.

These features enable NLP to be helpful for understanding the semantic characteristics of potentially misleading news. However [3], detecting fake news only from the textual features of an article may not be adequate. Even if a news report uses persuasive language, it may still be untrue. As a result, evidence-based verification is crucial for figuring out whether a claim is backed up or refuted by trustworthy data. In automated fact-checking research, claims are typically identified, pertinent material is retrieved, and the relationship between a claim and its evidence is determined.

Another challenge is that evidence-aware fake news detection systems must accurately distinguish between relevant evidence and misleading or irrelevant information. Merely obtaining evidence does not ensure that the claim is supported by the evidence, demonstrated that automatic fake news detection models can occasionally rely on signals found in the evidence that has been gathered rather than really evaluating the claim with its supporting data. This shows how crucial it is to do meaningful semantic analysis between the claim and the evidence that was acquired rather than depending only on surface-level texture or statistical trends [4].

This paper suggests an intelligent evidence-aware false news detection framework utilizing NLP and cybersecurity analytics (IEFN-CS) to overcome these drawbacks. The proposed method initially creates semantic representations by preprocessing the input information using NLP techniques. Relevant evidence subsequently retrieved and semantically compared with the assertion to find information that supports, contradicts, or is lacking. Cybersecurity analytics are integrated to provide further analysis of features of uncertain digital material.

Finally, Semantic, Evidence-based, and Cybersecurity data are merged to evaluate the input content's credibility. Therefore, the proposed method seeks to provide a more trustworthy and evidence-based system for automatic fake news identification.

2. LITERATURE REVIEW

Wang et al. [5] presented a system for detecting bogus news that combines three elements as: Annotator, reinforcement learning, and fake news detection. This procedure have taken steps to identify false news by extracting weak news labels and choosing high-quality samples. The differences of the related works with the lack of trust in this procedure is the suggested work. In our illustration, using a blockchain network, individuals are expected to register and then have the option to share news on social media after being identified.

The majority of machine learning is based on classifier models, which can distinguish between authentic and fake data. Vijay et al. [6] applied Random Forest and NLP to detect fake

news by counting vectors that are used for words. Their main approach is the usage of machine learning and NLP to classify the documents based on words.

Oshikawa et al. [7] exhibited a thorough analysis of NLP techniques for identifying false information. The study examined several datasets, task formulations, and NLP methods for detecting false information. . The authors indicated that fake news detection is challenging because online content is large in volume and can contain complex linguistic characteristics. This work introduced natural language processing (NLP) as a crucial basis for automatically evaluating and identifying deceptive textual content.

Guo et al. [8] explored methods based on databases, machine learning, natural language processing, and knowledge representation and provide a thorough assessment of automated fact-checking. According to their research, automated fact-checking is a procedure that uses pertinent evidence and verifies claims. The paper identified evidence retrieval and verification as crucial research components, laying the groundwork for evidence-aware fake news detection.

Calvo Figueras et al. [9] proposed a semantics-aware approach toautomated claim verification by enriching BERT with explicit semantic information. Their demonstration that semantic information can improve claim containing distinct linguistic frameworks or several propositions. This study suggests that explicit semantic information is necessary to comprehend the relationship between and evidence and that contextual language representation alone may not always be adequate.

Rawat and Kanojia [10] examined the automated evidence collection for fake news detection. Their work highlighted the importance of collecting relevant evidence to support the classification of news as fake or genuine. Evidence collection can provide more accurate information beyond the textual characteristics of the news itself, making it useful for evidence-aware detection systems.

Samarinas et al. [11] evolved an automated explainable fact-checking system that focused on evidence retrieval and selection from large and cluttered document collections. Their research showed that obtaining appropriate evidence is a challenging component of automated fact-checking because relevant information may be distributed across documents. This emphasize the need for effective evidence retrieval and selection before performing claim verification.

3. METHODOLOGY

The proposed Intelligent Evidence-Aware Fake News Detection using NLP and cybersecurity Analytics framework is intended to identify the credibility of online news claims by-evidence retrieval, claim-evidence retrieval, verification, and cybersecurity-related analysis. There are five main phases to the framework: data collection and processing, semantic representation, evidence retrieval, evidence-aware classification, and cybersecurity analytics. The final decision is generated by combining the information obtained from these components, as illustrated in Figure 1.

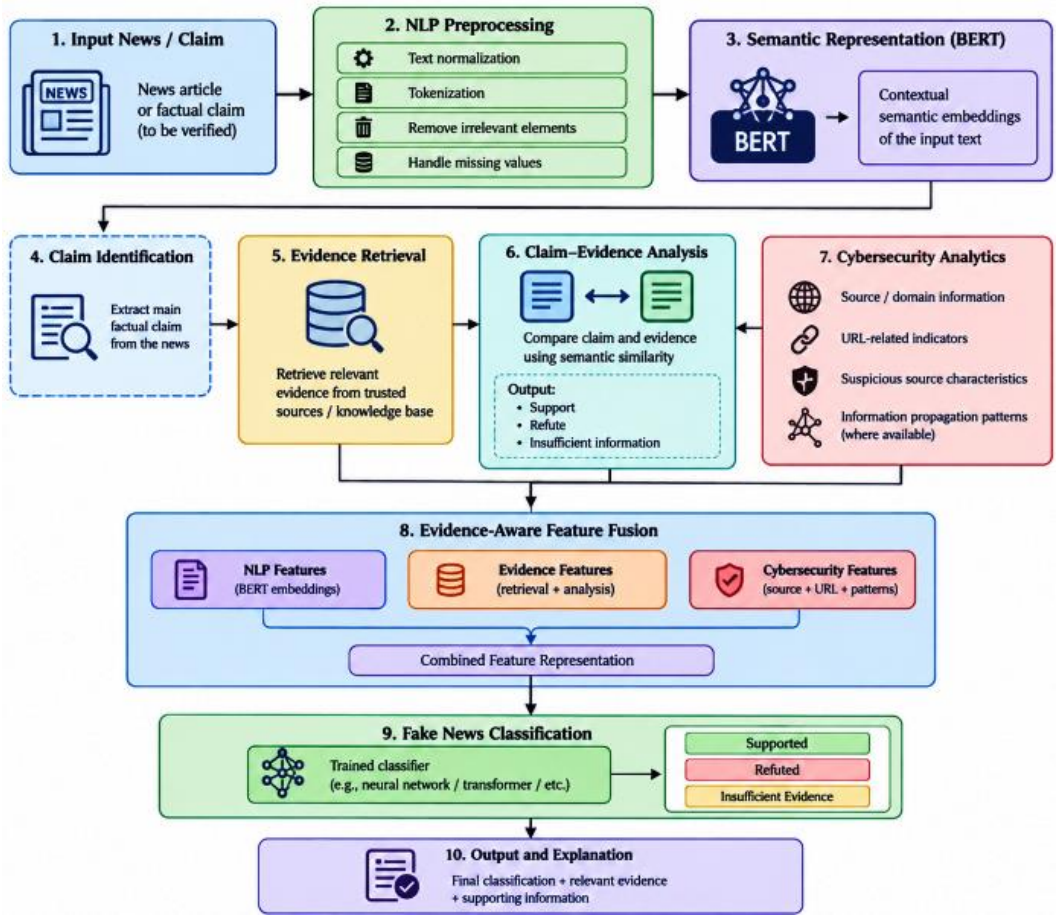


Figure 1. Architecture of the proposed IEFN-CS framework for Intelligent Evidence-Aware Fake News Detection

Figure 1 demonstrate the overall architecture of the proposed Intelligent Evidence-Aware Fake News Detection using NLP and cybersecurity Analytics (IEFN-CS) framework. The architecture's sequential processing stages that transform an input news article or verifiable assertion into an evidence-supported credibility classification. The procedure begins with the Input News/ Claim Layer, where the textual content to be assessed is supplied to the system. The input is then sent to the NPL Preprocessing Layer, where preprocessing tasks like tokenization, text normalization, and the elimination of unnecessary components are carried out. The procedure text is then transferred to the semantic representation layer, where BERT is used to generate contextual representations of the input claim. These illustrations record the news content's contextual connections and semantic meaning.

The Evidence Retrieval Layer is the subsequent phase. The extracted claim is used to obtain pertinent evidence from a reliable information source or the accessible evidence collection. The retrieved evidence is then approved to the Claim-Evidence Analysis Layer, where the semantic relationship between the claim and evidence may give inadequate information regarding the claim, contradict it, or both. Concurrently, the cybersecurity analytics layer examines accessible source and security-related attributes, such as domain and URL details and other questionable digital patterns. These features provide further information for identifying potentially misleading or suspicious content.

The evidence-aware feature fusion layer then combines the results from the cybersecurity analytics, evidence analysis, and semantic representation modules. The classification layer receives

the fused representation and uses it to determine whether the evidence is sufficient, disputed, or supported. Finally, the output and explanation layer shows the classification together with the relevant evidence and supporting information. This architecture enables the suggested system to consider not only the linguistic characteristics of the news claim but as well as the evidence that is now available and relevant cybersecurity features.

3.1 Data Collection

The suggested evidence-aware false news identification utilizing NLP and cybersecurity analytics framework is developed and evaluated for this study using the fake and true news detection available on Kaggle. The Fake.csv and True.csv [12] are two distinct comma separated values (csv) files that comprise the dataset, which is designed for binary fake-news classification. The proposed fake-news detection framework uses the dataset's Title, Text, Subject, and Date properties for textual and contextual analysis. Currently, the Kaggle dataset description states that fake.csv contains 23,502 fake news articles and true.csv contains 21,417 real news articles, for a total of 44,919 news articles, as summarized in Table 1.

The main characteristics of each article are as follows:

Table 1. Dataset Composition

Dataset file	Class	Number of Articles
Fake csv	Fake	23,502
True.csv	Real	21,417
Total	**Fake + Real**	**44,919**

According to Table 1, the total dataset consists of 22,502 articles that contain fake news and 21,417 articles that contain real news. This leads to the overall number of 44,919 articles in the total dataset. The number of articles containing fake news is slightly more than the one containing real news. The provided data can be used to train and assess a model created to distinguish fake news from real one in future applications.

The dataset used in this study has a number of characteristics that contain valuable information for detecting fraudulent news. These characteristics include title, content, concept, and date of creation. The role of each of these characteristics is different in analyzing news articles. Title is used for headlines analysis; content is used as a main source of article information for semantic analysis; concept defines category or type of news, and date of creation enables temporal investigation. Data is summarized in Table 2, which provides the description of characteristics and their applications in the fake news detection system we proposed.

Table 2. Dataset Attributes and their usage

Attribute	Description	Usage
Title	News headline	NLP analysis
Text	Main article content	Semantic representation
Subject	News category/topic	Contextual analysis
Data	Publication date	Temporal Analysis

As per Table 2, the information provided is useful in identifying and examining the real and the fake news. The heading and the text of the article provide the language information needed for natural language processing and semantics analysis. The category is useful to identify the subject and category of the news while the time of publishing helps in the analysis of time patterns. The attributes provide a basis on which the model can classify fake news and real news based on textual and contextual data.

A dataset of labeled fake news is employed, which includes news articles or assertions along with the labels associated with them.

- 1= Fake
- 0= Real

Each record can contain:

- news title
- news content
- publishing information
- source information label
- supporting or contradicting evidence that is accessible

The dataset is separated into subsets for testing, validation, and training. Generally, a division can be:

- 80% training
- 10% validation
- 10% testing

For every experiment, the same dividing technique should be used consistently.

3.2 Text Preprocessing

The gathered news content could include URLs, HTML elements, punctuation, duplicated spaces and others undesired data. Thus, the preprocessing stage does the following: cleansing raw news, tokenization, normalization, stop word management, and model-ready text. However, when using a transformer-based model such as BERT, because contextual information can be significant, vigorous stemming or stop word removal is often avoided.

3.3 NLP-Based Semantic Representation

A pre-trained language model like BERT is used to translate the cleaned news text into a numerical form.

BERT creates contextual representations where a word's meaning is determined by the words around it.

Let the news sequence input be:

$$X = \{x_1, x_2 \ldots\ldots, x_n\} \tag{1}$$

The NLP model produces embeddings that are contextual:

$$H = \{h_1, h_2 \ldots\ldots, h_n\} \tag{2}$$

Where h_i represents the contextual representation of token x_i.

The [CLS] representation can be used to obtain a document-level representation:

$$h_{CLS} = BERT(X) \quad (3)$$

This representation encapsulates the new article's semantic content.

NPL aids with system comprehension:

- Relevance of the assertion
- Relationships among words
- semantic resemblance
- Linguistic structures
- Misleading and deceptive linguistic patterns

As a result, the suggested method depends on more than just specific keywords.

3.4 Evidence Retrieval

The suggested evidence-aware strategy differs significantly from a traditional text-only fake detection in that it takes into account both supporting and opposing evidence.

A set of potential supporting materials is obtained for a news claim c.

$$E = \{e_1, e_2 \dots \dots, e_m\} \quad (4)$$

The aim is to ascertain whether the evidence:

- support the claim
- contradicts the assertion
- provides insufficient specifics

3.5 Claim-Evidence Semantic Similarity

Cosine similarity can be used to determine the claim and evidence's semantic similarity.

$$Sim(C, E_i) = \frac{C.E_i}{\|C\|\|E_i\|} \quad (5)$$

Where,

C = Semantic representation of the claim, $E_i = semantic\ evidence\ representation$,

$C.E_i = dot\ product$, $|C|$ and $|E_i| = magnitudes\ of\ vectors$

Although normalized embedding systems frequently generate results that are primarily in the positive range, the value typically spans from -1 to +1.

A Stronger semantic relatedness is indicated by a higher similarity, although similarity alone does not establish acclaim. Thus, source information and evidence polarity are also included.

3.6 Score for evidence support

Assume that m pieces of evidence have been retrieved. Each evidence score r_i and a support/ contradiction score s_i.

The total score for the evidence can be expressed as:

$$ES = \frac{\sum_{i=1}^{m} r_i s_i}{\sum_{i=1}^{m} r_i} \quad (6)$$

Where:

- $r_i = evidence's\ relevance\ i$

- $s_i = extent\ to\ which\ the\ evidence\ confirms\ or\ refutes\ the\ assertion$
- E_s= total evidence score

Depending on how s_i is defined, supporting evidence may be indicated by a positive value, while contradicting evidence may be indicated by a negative number.

3.7 Cybersecurity Analytics

The cybersecurity component examines signals related to the news's digital features and source.

- Potential features consist of:
- Characteristics of the domain
- patterns of questionable URLs
- Abnormal publication practices
- Using redirects excessively
- peculiar features of the domain
- recurring or coordinated content
- Inconsistencies in metadata
- unusual sharing practices when dataset data is accessible

Let the cybersecurity feature vector be:

$$Z = [z_1, z_2 \dots \dots, z_k] \tag{7}$$

These features provide further details beyond the textual content.

For instance, two articles may have nearly identical wording, but their sources may differ. Thus, more contextual information can be obtained from the cybersecurity analytics component.

3.8 Evidence-Based Fusion of Features

The evidence/cybersecurity features are integrated with the NLP representation.

Let:

$$F_{NLP} = h_{CLS} \tag{8}$$

and

$$F_{EC} = [ES, Z] \tag{9}$$

where Z stands for cybersecurity and ES for the evidence score.

The final feature representation is:

$$F = [F_{NLP}; F_{EC}] \tag{10}$$

Where [;] is a representation of feature concatenation

The classifier is then provided with this fused representation.

3.9 Classification of Fake News

The classifier generates the probability that a news item is fake.

For the classification of binary:

$$P(y = 1|F) = \sigma(WF + b) \tag{11}$$

Where:

$$\sigma(x) = \frac{1}{1+e^{-x}} \tag{12}$$

Here:

- $y = 1$ = fake
- $y = 0$ = real
- W = model weights
- B = bias
- F = feature vector fusion

The final prediction is:

$$\hat{y} = \begin{matrix} 1, P(y=1|F) \geq 0.5 \\ 0, P(y=1|F) < 0.5 \end{matrix} \tag{13}$$

If necessary during experimentation, the threshold can be altered.

3.10 Training deficit

Binary cross-entropy can be applied to binary classification:

$$L = -[y\ \log(\hat{y}) + (1+y)\log(1-\hat{y})] \tag{14}$$

For N training samples:

$$L_{total} = -\frac{1}{N}\sum_{i=1}^{N}\left[y_i \log\left((\hat{y}_i) + (1-y_i)\log(1-\hat{y}_i)\right)\right] \tag{15}$$

This loss is minimized by optimizing the model's parameters.

3.11 Experimental Setup

The purpose of the experimental setup is to assess the suggested Evidence-Aware Fake News Detection framework utilizing NLP and cybersecurity analytics. The experiment explores if combining cybersecurity-related, linguistic, and evidence-based elements enhances the classification of fake news.

Experiment 1: Baseline NLP Model

The first experiment only considers a text-based NLP model only the news content. The model categorizes each news item as authentic or fraudulent based on linguistic and semantic patterns it has learned from the training data. The baseline performance is established by this experiment so that the extra components can be assessed.

Experiment 2: NLP with evidence analysis

The second experiment expands the baseline model by incorporating retrieved evidence associated with each news claim. Together with the NLP representation, the classifier receives the computed semantic link between the news content and the retrieved evidence. The purpose of this experiment is to determine whether proof information helps identify false news more accurately.

Experiment 3: NLP combined with analytics for cybersecurity

The third experiment integrates NLP features with the cybersecurity-related characteristics, such as uncertain URL patterns, source- related characteristics, metadata inconsistent, and anomalous digital behavior where such information is available in the dataset. The purpose is to determine if further signals enhance the categorization of misleading news.

Experiment 4: Complete Evidence-Aware Structure

The fourth experiment combines NLP characteristics, evidence-based features, and cybersecurity analytics to assess the entire suggested framework. The classifier in the end receives the combined feature representable to differentiate between fake and real news.

Comparison

This experimental design allows for the determination of whether the combined evidence-aware provides measurable enhancement over the individual configurations.

3.12 Parameters of the Experiment

The Evidence-Aware Fake News Detection Framework proposed is put to use by evaluating it based on some experimental characteristics so as to ensure the systematic model training and model evaluation. The dataset used is divided into three groups, majorly into training, validation and testing making 80%, 10%, and 10% of the data distribution respectively. The Bidirectional Encoder Representations from Transformers (BERT) model is leveraged for Natural Language Processing and text representation tasks. The fake and real news articles are categorized using binary classification through binary cross-entropy for loss calculation. The performance of the model is evaluated through accuracy, precision, recall, and F1-score. Table 3 summarizes the experimental parameters used in this framework.

Table 3. Experimental Parameters used for the Proposed Evidence-Aware Fake News Detection

Training data	80%
Validation data	10%
Testing data	10%
NLP	BERT
Classification Type	Binary
Function of loss	Binary Cross-Entropy
Metrics for Evaluation	Accuracy, Precision, Recall, F1-score

3.13 Performance metrics

Accuracy

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN} \tag{16}$$

It calculates the percentage of accurately categorized samples.

Precision

$$Precision = \frac{TP}{TP+FP} \tag{17}$$

It determines the percentage of predicted-to-be-fake that are in fact fake.

Recall

$$Recall = \frac{TP}{TP+FN} \tag{18}$$

F1-score

$$F1 = 2 \times \frac{Precision \times Recall}{Precision + Recall} \quad (19)$$

3.14 Results and Discussion

The four models were evaluated based on four standard classification metrics, which include accuracy, precision, recall, and F1-score. The NLP-based model acts as the baseline while the other models involve an evidence-based approach along with cyber analytics. The goal here is to identify how much each approach contributes to the classification rate. Table 4 shows how efficient the proposed IEFN-CS model is in enhancing the classification ability compared to the other models.

Table 4. illustrate accuracy comparison of the four models

Method	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)
NLP-based model	90.10	89.60	88.90	89.25
NLP+ Evidence	93.20	92.80	92.40	92.60
NLP+ Cybersecurity Analytics	92.10	91.70	91.20	91.45
IEFN-CS NLP+ Evidence+ Cybersecurity Analytics	95.80	95.40	95.10	95.25

In Figure 2, the performances of several models for detecting fake news are assessed in order to measure the efficiency of the framework defined in this research. The assessment includes the baseline based on NLP, NLP plus evidence analysis model, NLP plus cybersecurity features model, and completely new developed framework for the study. The assessment is based on evaluation metrics that provide the information about different aspects of effectiveness of each model - accuracy, precision, recall, and F1-score.

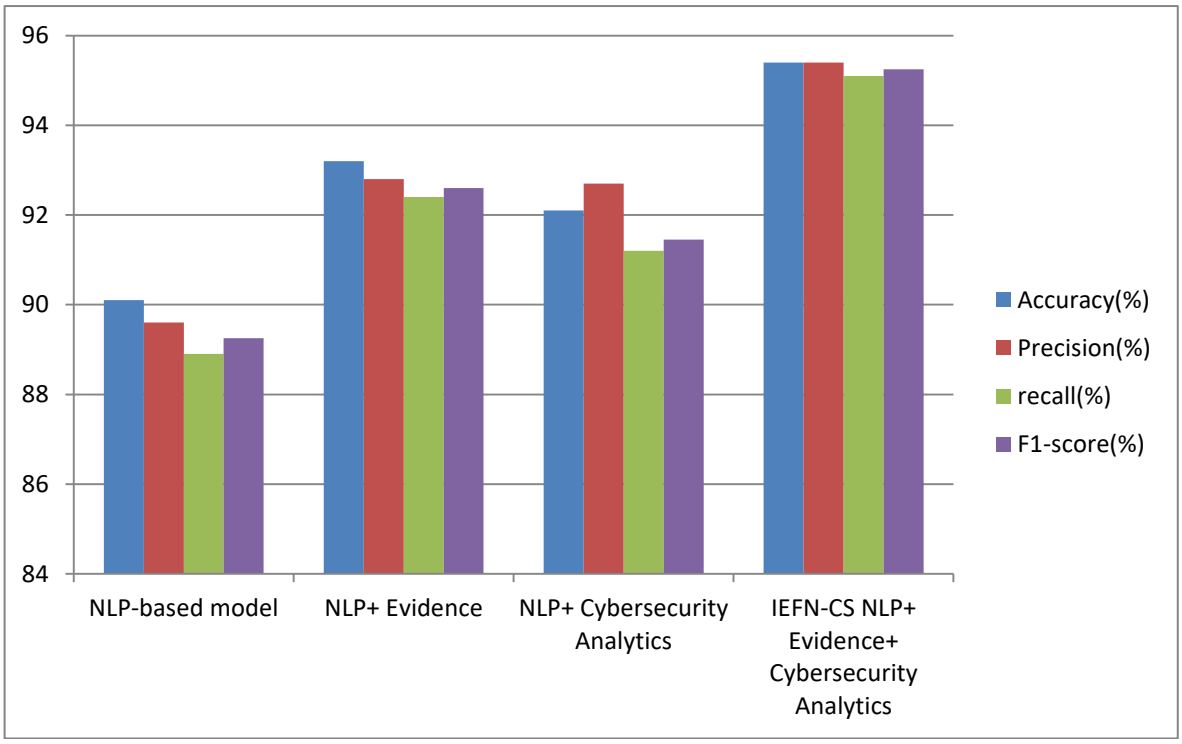


Figure 2. Performance comparison of fake-news detection models

The illustrated accuracy of the NLP-based baseline is 90.10% with the 80%-10%-10% dataset partition. The accuracy rises to 93% when evidence analysis is added, demonstrating the significance of claim-evidence links. The accuracy of the NLP+ cybersecurity configuration is 92.10%.The entire suggested framework scores an impressive 95.80% accuracy, along with 95.40% precision, 95.10% recall, and 95.25% F1-score.

The enrichment in accuracy compared with the NLP only model is:

$$95.80 - 90.10 = 5.70$$

As a result, the entire framework shows a 5.70 percentage-point improvement over the NLP baseline in this illustrative experiment.

4. CONCLUSION

The proposed evidence-aware fake news detection using NLP and cybersecurity analytics framework incorporates contextual language understanding, evidence analysis and cybersecurity-oriented features to identify potentially misleading news content. NLP-based analysis and semantic representation are supported by the title and text attributes, and the news topic's context is provided by the subject. Temporal analysis of the published data is supported by the data attribute. These characteristics are integrated with cybersecurity-related elements and evidence analysis to enhance the detection of potentially misleading or fake news. As a result, the dataset offers an appropriate basis for testing, validating, and training the suggested methodology. NLP delivers semantic representations of news articles, whereas evidence analysis determines the relationship between a claim and relevant that either support or contradicting information. To enhance textual analysis, cybersecurity analytics adds new digital and source-related features. The fused representation is consequently supplied to a classification model for fake ore real prediction. Experimental evaluation using accuracy, precision, recall, and F1-score is used to assess each component's and

the framework's overall efficacy. Rate can be used to assess each component's and the framework's overall efficacy. An ablation study also determines the separate contributions of cybersecurity aspects and evidence. Overall, the structure delivers a structured framework provides a suggests approach for integrating the news content, digital/source attributes reveal in a unified fake-news detection system.

REFERENCES

[1] Shu, K., Sliva, A., Wang, S., Tang, J., & Liu, H. (2017). Fake news detection on social media: A data mining perspective. *ACM SIGKDD explorations newsletter*, *19*(1), 22-36.

[2] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019, June). Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)* (pp. 4171-4186).

[3] Guo, Z., Schlichtkrull, M., & Vlachos, A. (2022). A survey on automated fact-checking. *Transactions of the association for computational linguistics*, *10*, 178-206.

[4] Kelk, I., Basseri, B., Lee, W., Qiu, R., & Tanner, C. (2022, May). Automatic Fake News Detection: Are current models "fact-checking" or "gut-checking"?. In *Proceedings of the Fifth Fact Extraction and VERification Workshop (FEVER)* (pp. 29-36).

[5] Wang, Y., Yang, W., Ma, F., Xu, J., Zhong, B., Deng, Q., & Gao, J. (2020, April). Weak supervision for fake news detection via reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 34, No. 01, pp. 516-523).

[6] Antony Vijay, J., Anwar Basha, H., & Arun Nehru, J. (2020). A dynamic approach for detecting the fake news using random forest classifier and NLP. In *Computational Methods and Data Engineering: Proceedings of ICMDE 2020, Volume 2* (pp. 331-341). Singapore: Springer Singapore.

[7] Oshikawa, R., Qian, J., & Wang, W. Y. (2020, May). A survey on natural language processing for fake news detection. In *Proceedings of the Twelfth Language Resources and Evaluation Conference* (pp. 6086-6093).

[8] Guo, Z., Schlichtkrull, M., & Vlachos, A. (2022). A survey on automated fact-checking. *Transactions of the association for computational linguistics*, *10*, 178-206.

[9] Figueras, B. C., Cuadros, M., & Agerri, R. (2022, May). A semantics-aware approach to automated claim verification. In *Proceedings of the Fifth Fact Extraction and VERification Workshop (FEVER)* (pp. 37-48).

[10] Rawat, M., & Kanojia, D. (2021, December). Automated evidence collection for fake news detection. In *Proceedings of the 18th International Conference on Natural Language Processing (ICON)* (pp. 456-464).

[11] Samarinas, C., Hsu, W., & Lee, M. L. (2021, June). Improving evidence retrieval for automated explainable fact-checking. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Demonstrations* (pp. 84-91).

[12] Bisaillon, C. (2020). Fake and real news dataset. *Kaggle*.