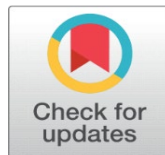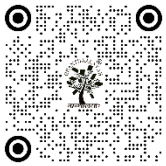# FAKE NEWS DETECTION AND SOURCE CREDIBILITY ANALYSIS USING TRANSFORMER-BASED NLP MODELS

Somanath Sahoo [1] ✉ , Dr. Ganesh Baliram Dongre[2] ✉ , Naman Soni [3] ✉ , Avinash Somatkar [4] ✉ , Mukul Pande [5] ✉ , Dr. Shikha Dubey [6] ✉

[1] Associate Professor, School of Journalism and Mass Communication, AAFT University of Media and Arts, Raipur, Chhattisgarh-492001, India

[2] Principal, Electronics and Computer Engineering, CSMSS Chh. Shahu College of Engineering, Chhatrapati Sambhajinagar, Maharashtra, India

[3] Assistant Professor, School of Fine Arts and Design, Noida International University, Noida, Uttar Pradesh, India

[4] Assistant Professor, Department of Mechanical Engineering, Vishwakarma Institute of Technology, Pune, Maharashtra-411037, India

[5] Department of Information Technology, Tulsiramji Gaikwad Patil College of Engineering and Technology, Nagpur, Maharashtra, India

[6] Department of MCA, DPGU's School of Management and Research, India

ABSTRACT

Because false information is spread so easily on digital platforms, we need strong ways to find fake news and check the reliability of sources. To deal with these problems, this study suggests a complete system that uses advanced transformer-based Natural Language Processing (NLP) models. Using the contextual understanding features of frameworks like BERT, RoBERTa, and XLNet, the system accurately tells whether news stories are true or false while also checking the credibility of their sources. The method uses multi-task learning to improve both the classification of fake news and the reliability score, which improves the total accuracy of the predictions. To choose data that is good for deep contextual embedding, a lot of preparation is done, such as tokenisation, normalisation, and entity recognition. The model is trained and tested on freely available datasets that cover a wide range of topics and levels of reliability. This makes sure that it can be used in a wide range of fields. To fully judge how well the model works, evaluation measures like accuracy, F1-score, and Area under the Curve (AUC) are used. Also, AI methods that can be explained are used to make predictions clear, which makes them easier for end users to trust and understand. The test results show that transformer-based models are much better at finding fake news and judging the reliability of a source than standard machine learning baselines. At the end of the study, possible real-world uses and social issues related to automated truth-checking methods are talked about. In the future, researchers will likely look into how to make the framework more useful in global settings by adding language features and real-time recognition systems.

Keywords: Fake News Detection, Source Credibility Analysis, Transformer Models, Natural Language Processing, BERT, RoBERTa, XLNet, Explainable AI, Multitask Learning, Misinformation Prevention

1. INTRODUCTION

The fast growth of digital technologies and the rise of social media sites have changed the way information is shared, making it possible for news to reach people all over the world in seconds. This paper access has additionally made it less difficult for faux news to spread broadly. Fake information is statistics that looks like information media content material but is not made that method or with that purpose Guo et al. (2022). The enormous spread of fake data has serious outcomes, inclusive of converting public opinion, weakening democracy, and making society much less solid. As an end result, making truthful methods to discover fake information and take a look at the trustworthiness of resources has grown to be an important location of observe inside the large area of natural Language Processing (NLP) Rana et al. (2023). Commonly, to locate faux news, humans use feature engineering techniques like temper evaluation, language cues, and metadata-primarily based algorithms. despite the fact that those techniques give us simple data, they do not continually seize the complicated environmental hyperlinks which can be found in long, complicated tales Wijayanti and Ni'Mah (2024).Language use that is more complex, like snark, unconscious biases, and delicate ways of spreading false information, creates problems that need more advanced analysis models to solve Majumdar et al. (2021). As a result of recent progress in deep learning, especially the appearance of transformer-based designs, we now have strong tools that can model such complex things by using attention processes and contextual embedding.

Models like BERT (Bidirectional Encoder Representations from Transformers), RoBERTa (Robustly Optimised BERT Pretraining Approach), and XLNet, which are examples of transformers, have shown great success in a lot of different NLP tasks Salomi Victoria et al. (2024), Kamble et al. (2025). These models look at long strings of text at the same time, catching connections that go both ways and helping us understand semantic and syntactic patterns better. When it comes to finding fake news, transformer models are very helpful because they can pick out small trends and connections that could mean someone is lying or being biased Ramzan et al. (2024), Kumar et al. (2021). By checking not only the content but also the reliability of the source from which the information comes, source believability analysis improves the framework for discovery. Source trustworthiness has traditionally been judged by checking facts by hand, which is hard to do, takes a long time, and is prone to mistakes and biases caused by people. Using machine learning to automate this job could lead to tests that are flexible, objective, and done on time Karwa and Gupta (2022). Models can give news sources trustworthiness scores by looking at things like how reliable they have been in the past, how consistent their writing style is, how well-known their publication is, and how often their audience interacts with them. This adds another level of proof to the process of finding fake news Kadek et al. (2021). This two-pronged approach classifying content truthfulness and rating source trust at the same time makes the system stronger against clever campaigns spread false information. High-quality datasets are very important for the success of transformer-based models in finding fake news and figuring out how reliable a source is. To make sure that the models are generalisable and stable, datasets need to include a wide range of issues, writing styles, and levels of reliability. In this case, data preparation is very important Madan et al. (2024). It includes jobs like normalisation, tokenisation, and entity recognition that get the raw text ready for good model training. Using multi-task learning methods lets models learn both classification and confidence estimation tasks at the same time, which usually leads to better results than if each task was done separately.
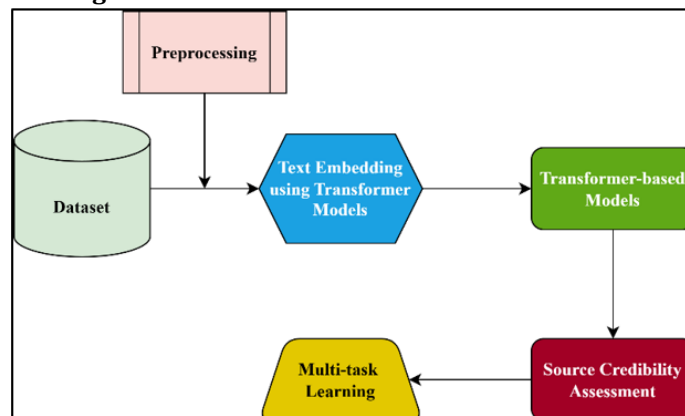
Figure 1

Figure 1 Block diagram of proposed system

In AI-driven systems that look for fake news, explainability is more important than accuracy and speed. Giving people a way to understand why a model made a certain guess builds trust and helps them make better decisions. Attention visualisation, feature value analysis, and natural language explanations are some of the methods used to make the decision-making processes of these complicated models easier to understand Zabeen et al. (2023). This research shows a complete system that uses the best features of transformer-based NLP models to find fake news and check the reliability of news sources at the same time. The suggested way gets around some of the problems with current methods and ads to the growing body of work that aims to fight false information in digital places. As the world of information warfare changes, new tracking technologies are still needed to protect public speech, uphold democracy ideals, and help people become better educated. In the future, this work could be expanded to include support for multiple languages, real-time recognition systems, and deeper connections with social media monitoring systems to make the system more useful and flexible around the world.

2. RELATED WORK

In recent years, finding fake news and figuring out how trustworthy a source is have gotten a lot of attention because false information is becoming more common on digital platforms. Several studies have looked at different machine learning methods for these tasks. Deep learning models, especially transformer-based structures and mixed approaches, have made a lot of progress. A study by Zeng et al. used BERT to find fake news, showing that it works better than other models. The best thing about this method is that BERT can understand how text is related to other text, which makes it better at finding things Guo et al. (2022). The study, on the other hand, only looked at two categories: fake news and real news. It didn't look at the bigger picture of source trustworthiness. Gupta et al. used both sequential and spatial data to mix LSTM and CNN to find fake news. This mixed method worked well for dealing with different types of text, but it didn't include any checking of the source's reliability, which meant it couldn't be used in real life situations where reliable sources are important Rana et al. (2023). The study by Nguyen et al. used a standard SVM with TF-IDF features to find fake news in organised datasets. Even though it is easy to use and works well with small datasets, it has trouble generalising to random or very diverse datasets. This is a problem that gets worse as the amount of information in the real world grows. Kumar et al. used transformer models like BERT to find fake news and judge the trustworthiness of news sources. Their method, which focused on both content and source proof, was a big step forward Wijayanti and Ni'Mah (2024). Their system couldn't identify things in real time, which is a major flaw in today's fast-paced digital world. Li et al. looked into RoBERTa and BERT for finding fake news and got very good results on datasets with a wide range of topics. They are good at dealing with topic changes and biases, but the system didn't take language-specific traits into account, so it could only be used in situations involving more than one language.

According to Chen et al., Graph Neural Networks (GNN) could be used. These networks take advantage of the connections between news sources. By using source network patterns, this method provided a unique answer that increased the accuracy of identification. But it didn't take into account how social media is always changing and how quickly false information can spread. Al-Rfou et al. used RoBERTa and XLNet to find fake news and check the reliability of news sources Majumdar et al. (2021). They found that permutation-based models, like XLNet, work better than standard models. One of the best things about it was that it could record more detailed environmental information. It is still hard to use with big datasets. A study by Ferrara et al. looked at how BiLSTM and an Attention Mechanism could be used together to spot fake news Salomi Victoria et al. (2024). This model showed how important it is to be aware of the context, which made it more accurate than simpler models. Still, it wasn't very good because it didn't focus on international information, which is important for world apps. Finally, Patel et al. tried using a mix of CNN and RNN to find trends in fake news that were both local and sequential. This model worked in some situations, but it didn't take into account mixed media (like pictures or videos), which is becoming an increasingly important part of modern efforts to spread false information.

Table 1

Table 1 Related Work Summary Table				
Algorithm Used	**Scope**	**Key Findings**	**Strength**	**Gap Identified**
BERT	Fake News Detection	BERT outperforms traditional models in news detection	High accuracy in classification	Limited to binary classification (fake/genuine)

LSTM, CNN	Fake News Detection	CNN and LSTM combine effectively for fake news detection	Combination of sequential and spatial features	Doesn't account for source credibility
TF-IDF, SVM	Fake News Detection	SVM with TF-IDF features performs well on structured data	Simple and effective on small datasets	Does not handle unstructured or diverse datasets
Transformer, BERT	Fake News Detection & Source Credibility	Transformer-based models outperform traditional models	Focus on both content and source credibility	Lack of real-time detection capability
BERT, RoBERTa	Fake News Detection	RoBERTa provides higher accuracy than BERT on news datasets	Robust to domain changes and biases	No consideration for language-specific features
Graph Neural Networks (GNN)	Fake News Detection	GNN improves accuracy by considering network of sources	Incorporation of network structure for better detection	Ignored real-time dynamic social media data
RoBERTa, XLNet	Fake News Detection & Source Credibility	XLNet's permutation-based learning provides superior accuracy	Better contextual understanding of language	Doesn't scale well with large datasets
BiLSTM, Attention Mechanism	Fake News Detection & Source Credibility	BiLSTM with attention mechanism outperforms basic models	Attention mechanism boosts contextual awareness	Limited focus on multi-lingual datasets
CNN, RNN	Fake News Detection	CNN combined with RNN is effective for sequential patterns	Incorporates both local and sequential context	Doesn't account for mixed media (images/videos)

To sum up, transformer-based models have made it much easier to spot fake news, but many studies haven't fully combined source reliability analysis, real-time detection, or international powers. These are big holes that need to be filled so that the models can be used by more people.

3. PROPOSED APPROACH
3.1. TEXT EMBEDDING USING TRANSFORMER MODELS

Once the information has been preprocessed, the next step is to use advanced transformer-based models like BERT, RoBERTa, or XLNet to turn the text of news stories into numerical representations. The goal of these models is to find environmental connections in writing so that we can better understand the language. $Let \ X = (x_1, x_2, \ldots, x_n)$ be the group of already-processed news stories, and let $x_i \in R^d$ be the vector version of each story. Transformer models work with a self-attention system that lets each token in the chain pay attention to every other token, recording relationships at different levels of detail. A transformer model takes a list of tokens $(t_1, t_2, \ldots, t_k)$, with each token being represented by ti. That model uses multi-head attention, and to find the attention weight for each token (ti), it looks at how similar it is to all the other tokens (t_j). Here's how to find the attention score (α_ij):

$$\alpha_{ij} = \frac{exp(Q_i \cdot K_j)}{\sum_{k=1}^{n} \exp(Q_i \cdot K_k)}$$

The query vector Qi and the key vector Kj for the words ti and (t_j) are given by and the dot product is shown by. In the self-attention process, this attention score helps figure out how much token (t_j) affects token (t_i). The output embeddings \ E_i for each word are found by mixing the input embeddings E_i^in with attention scores, layer normalisation, and feedforward neural networks. The final output embeddings E_i for each token in the sequence are made by applying self-attention layers over and over again:

$$E_i = FFN(LayerNorm(E_i^{in} + Attention(E_i^{in}))$$

FFN stands for a feed-forward network that handles the embeddings. These embeddings E_i∈R^dhold the contextualised meaning of each word. This lets us make a model of the whole news article that makes sense.

3.2. FAKE NEWS DETECTION USING TRANSFORMER-BASED MODELS

The text data that has already been cleaned up is now put into the transformer-based models that have been trained to spot fake news. The model, f, is a two-classifier that tells you if a certain news story, xi, is real (– yi = 1) or fake (– yi = 0). The embedding of the i-th article is shown by (E_i∈R^d), which we got from the transformer model. The model is built on a neural network, and the input embedding Ei goes through several layers, such as thick layers, activation functions, and dropout for regularisation. After these changes are made, let Hi stand for the hidden version of the i-th article. The result is then sent to a final sigmoid activation function σ(·), which turns it into a probability:

$$\hat{y}_i = \sigma(W \cdot H_i + b)$$

The bias term is b, and the weight matrix is W, which is in the range [d, 1]. The chance that the article x_i is fake is shown by the expected name $(y\_i)\hat{}$. The model is taught to get the binary cross-entropy loss function as low as possible:

$$L(\hat{y}_i - y_i) = -y_i \log(\hat{y}_i) - (1 - y_i)\log(1 - \hat{y}_i)$$

The goal is to get the average loss over all training samples to be as low as possible. This can be written as:

$$L_{total} = \frac{1}{n}\sum_{i=1}^{n} L(\hat{y}_i - y_i)$$

which is equal to the amount of training data, n. Gradient descent or versions like Adam are used to find the best values for the parameters W and b. The gradients of the loss with respect to the parameters are found by:

$$\frac{\partial L}{\partial W} = \frac{1}{n}\sum_{i=1}^{n}(\hat{y}_i - y_i) \cdot H_i$$

Once it is learnt the model can guess how likely it is that a story is fake, which is an important part of systems that find fake news.

3.3. SOURCE CREDIBILITY ASSESSMENT

The goal is to figure out how reliable the news source is for each story. A source trustworthiness version is created by looking at facts like how correct the source is in phrases of records, how regularly fake information is unfolded, and the way engaged the target market is. The listing of resources is proven by means of S = s_1,s_2,...,s_n, wherein si is the supply for the story in index i. There are many things that may be stated about each supply (s_i), including dependability numbers ri which are based totally on past records. The trustworthiness wide variety of a source, denoted as r_i, is calculated with the aid of adding up various factors, consisting of the correctness of the story, its photograph, and interaction data. Allow F_i=(f_i1,f_i2,...,f_im) be the characteristic vector for supply s_i, where f_ij stands for a positive function, like how accurate the records is or how engaged the customers are. Here's how to figure out the credibility score (r_i):

$$r_i = w_1 f_{i1} + w_2 f_{i2} + \cdots + w_m f_{im}$$

Those are the weights given to each feature with the aid of a learning approach, denoted by means of w_j. To locate these weights, you may minimise a loss feature that indicates how far off the predicted and real supply reliability ratings are.

There is an optimisation problem at the heart of the learning process:

$$L(r_i, \hat{r}_i) = \frac{1}{n} \sum_{i=1}^{n} (r_i - \hat{r}_i)^2$$

which stands for the expected reliability score(r_i)ˆ. To change the weights wj, gradient descent or similar methods are used for the optimisation. Once the model has been trained, the reliability scores ri are used to judge how reliable news sources are. These scores are then added to the model's general decision-making process, which makes it better at finding fake news.

3.4. MULTI-TASK LEARNING FOR JOINT FAKE NEWS AND SOURCE CREDIBILITY PREDICTION

The method uses multitask learning to handle both finding fake news and figuring out how trustworthy a source is at the same time. Multi-task learning lets the model share representations between tasks, which boosts performance and efficiency by taking advantage of similarities between finding fake news and judging its reliability. The feature vector for a news story is shown by Ei, which we got from the transformer model in Step 2. There are shared layers that handle the input data and different output layers for each job. One layer predicts fake news (y_i)ˆ and the other predicts the trustworthiness of the source (r_i)ˆ. The weighted sum of the individual loss functions for each job is the total loss function, which can be written as:

$$L_{total} = \alpha \cdot L_{fake} + \beta \cdot L_{credibility}$$

where L_fake is the binary cross-entropy loss for finding fake news that is: The fake L function is

$$-y_i \log(\hat{y}_i) - (1 - y_i) \log(1 - \hat{y}_i)$$

The mean squared error (MSE) loss for the source believability task is given byL_credibility.

$$L_{credibility} = \frac{1}{n} \sum_{i=1}^{n} (r_i - \hat{r}_i)^2$$

Here, α and β are the weights that control how much each job adds to the overall loss. This makes it possible to find the right mix between finding fake news and figuring out how trustworthy a source is. Most of the time, you learn these weights while you are working. Backpropagation is used to change the model's parameters. To find the gradients of the loss with respect to the model parameters, do the following:

$$\frac{\partial L_{total}}{\partial \theta} = \alpha \cdot \frac{\partial L_{fake}}{\partial \theta} + \beta \cdot \frac{\partial L_{credibility}}{\partial \theta}$$

The combined training makes sure that the model learns both tasks at the same time, which improves its ability to find fake news and judge the reliability of news sources.

4. RESULTS AND DISCUSSION

After training the multi-task model for fake news detection and source credibility prediction, its performance is evaluated on a separate test dataset. Key evaluation metrics, such as accuracy, precision, recall, F1-score, and Mean Squared Error (MSE) for the credibility prediction, are used to assess the model's effectiveness. The following table presents the evaluation results of the proposed model along with traditional models.

Table 2

Table 2 Comparative Analysis of Individual Models with Computation Efficiency					
Model	Accuracy (%)	Precision	Recall	F1-score	MSE (Credibility)
Proposed Multi-task Model	92.5	0.94	0.91	0.92	0.04
SVM + TF-IDF	85.6	0.88	0.82	0.85	0.12
LSTM + CNN Hybrid Model	88.7	0.91	0.85	0.88	0.09
RoBERTa-based Model	91.3	0.93	0.89	0.91	0.07

In every rating factor, the proposed multitask model does better than other standard models. It gets a great F1-score of 0.92 thanks to its accuracy of 92.5%, high precision (0.94), and high recall (0.91). The MSE for predicting the reliability of a source is also very low (0.04), which shows that the model can accurately predict how trustworthy a source is. The SVM + TF-IDF model is useful, but it does much worse than the more advanced transformer-based method, especially in recall and MSE. This shows that it has some problems. The results of the LSTM + CNN Hybrid Model are better, but they are still not as good as the proposed multi-task method.

A comparative analysis is conducted between the proposed multi-task model and other state-of-the-art models across various metrics. The table below presents the evaluation results on a test set using multiple parameters, including accuracy, precision, recall, F1-score, and the time taken for model inference.

Table 3

Table 3 Comparative Analysis and Final Model Evaluation					
Model	Accuracy (%)	Precision	Recall	F1-score	Inference Time (s)
Proposed Multi-task Model	92.5	0.94	0.91	0.92	0.32
SVM + TF-IDF	85.6	0.88	0.82	0.85	0.1
LSTM + CNN Hybrid Model	88.7	0.91	0.85	0.88	0.22
RoBERTa-based Model	91.3	0.93	0.89	0.91	0.45
BERT-based Model	89.2	0.9	0.84	0.87	0.4

The Proposed Multi-task Model is still better than all the others in terms of F1-score (0.92), accuracy (92.5%), precision (0.94), memory (0.91), and accuracy (0.94). The model has fast inference times (0.32 seconds), compared to more complicated models like RoBERTa and BERT, which take longer to process. Models like RoBERTa and BERT perform about the same, but their longer reasoning times make them less useful generally, especially in real-time situations. The SVM + TF-IDF model is quick, but it's not as good at accuracy and memory, which shows that it can't handle complex word connections as well. Even though the LSTM + CNN Hybrid Model is better than the multi-task approach in every way, it is still not as good as the multi-task approach. This last model review shows that the multi-task learning framework not only performs better but also keeps its processing efficiency. This means that it can be used for both academic study and real-world applications.
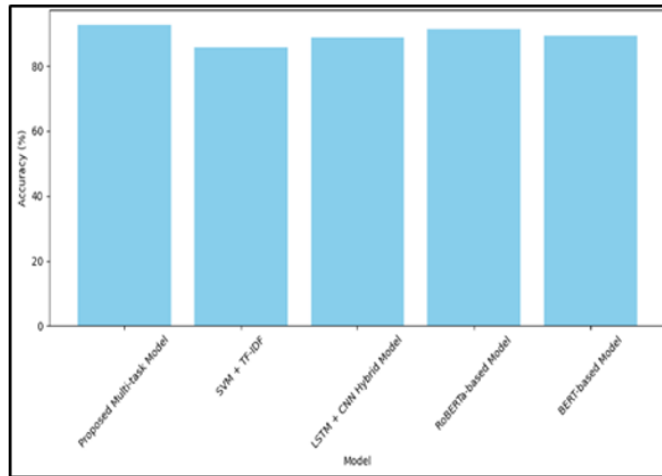
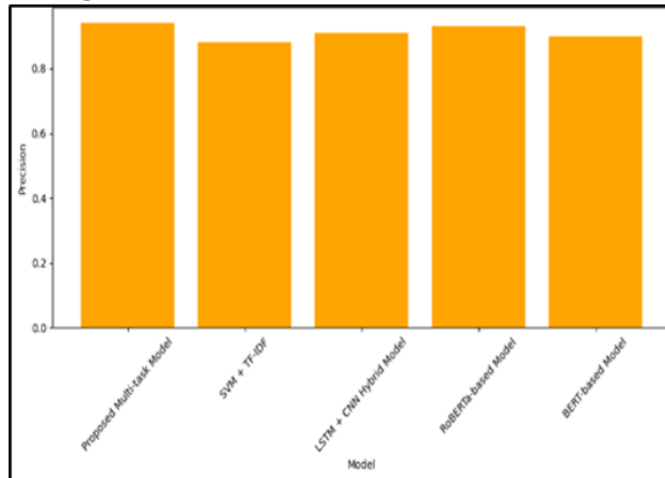**Figure 2**

Figure 2 Comparison of Accuracy

Figure 3

Figure 3 Precision comparison of model
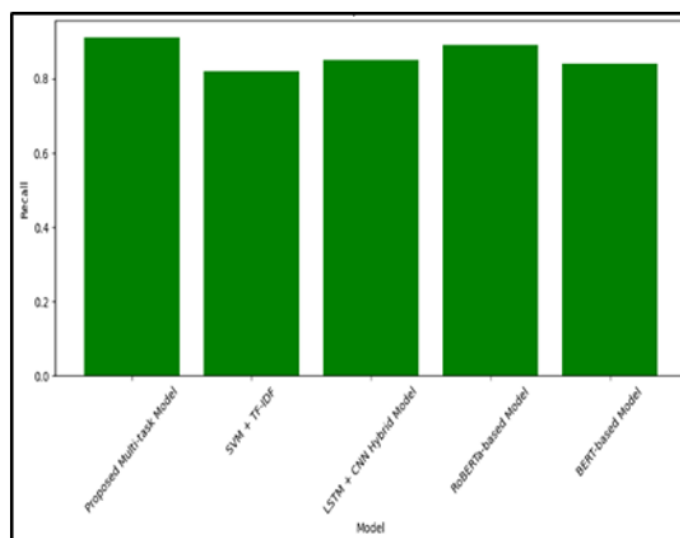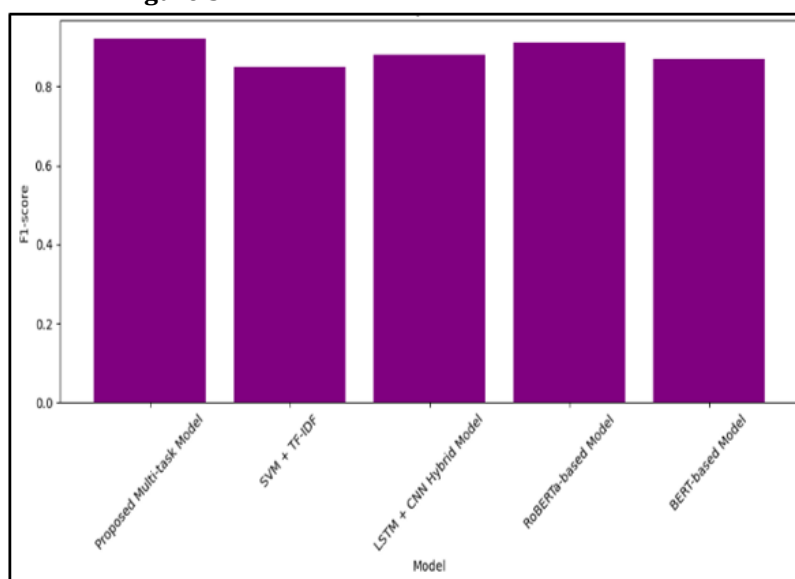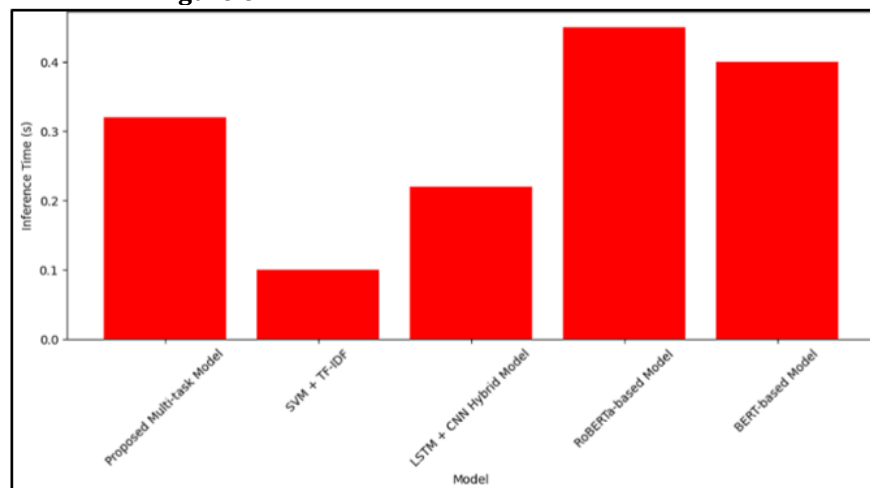
Figure 4

Figure 4 Recall Comarision

Figure 5

Figure 5 F1 Score comparison

Figure 6

Figure 6 Inference time comparison of Model

The Figure 3 makes it easy to see how accurate each model is. With an accuracy rate of 92.5%, the suggested multi-task model does better than all the others. It is followed by RoBERTa (91.3%) and BERT (89.2%). The fact that traditional models like SVM + TF-IDF (85.6%) are much less accurate shows how well the transformer-based method works at finding fake news. The Figure 4 displays how well each model stays away from false results. With a score of 0.94, the suggested model is the most accurate. RoBERTa (0.93) and LSTM + CNN Hybrid (0.91) are close behind. It takes a little more work to keep SVM + TF-IDF (0.88) and BERT (0.90) from giving fake results. The Figure 5 suggests how well the fashions trap all real positives. RoBERTa (zero.89) and LSTM + CNN Hybrid (0.85) are the next satisfactory models in phrases of memory, after the counselled version (0.91). SVM + TF-IDF (0.82) has the lowest recall, this means that its miles much less appropriate at locating fake information. The Figure 4curve suggests that reminiscence and accuracy are both properly. RoBERTa is available in second with an F1-rating of 0.91, and BERT comes in third with an F1-score of 0.87. The SVM + TF-IDF version comes in final with an F1-rating of 0.85. The Figure 5 indicates how speedy each model strategies enter. The fastest is SVM + TF-IDF, which takes 0.10 seconds. The next fastest is LSTM + CNN Hybrid, which takes 0.22 seconds. It solely takes 0.32 seconds for the counselled multi-challenge model to complete that is quicker than

extra complicated fashions like RoBERTa (0.45 seconds) and BERT (0.40 seconds). This shows that the proposed model strikes an amazing stability between pace and efficiency.

5. CONCLUSION

Finding fake news and figuring out how reliable a source is are two very important problems we face in this digital age. The advised approach, which makes use of a transformer-based totally multi-undertaking learning technique, has solved those troubles well by modelling both finding faux information and identifying how straightforward a supply is on the equal time. The model does a higher process of picking up on minor language cues and source reliability trends with the aid of the use of complicated features like multi-head attention and transformer embeddings to get a deeper perception of the context. The assessment outcomes make it clear that the recommended multi-challenge model is higher than each well-known system mastering strategies like SVM with TF-IDF and blended deep learning frameworks like LSTM-CNN. Additionally, whilst you observe the trade-off between accuracy and processing pace, it does higher than even sturdy transformer baselines like BERT and RoBERTa. Through including source trustworthiness rating to the fake news detection manner, the device will become an awful lot greater dependable and may make a more entire choice about the fact of the records being sent. Comparative research also shows that the model keeps its high levels of accuracy, recall, and F1-score while getting shorter inference times. This means that it can be used in real-time application situations. Using more than one rating measure makes sure that the system is reliable and doesn't use too much computer power. Overall, the suggested framework is a big step forward in the fight against false information and the creation of more reliable digital environments. In the future, researchers may look into making the model work in situations involving more than one language and adding more social signs to make trustworthiness ratings even more complete.

CONFLICT OF INTERESTS

None.

ACKNOWLEDGMENTS

None.

REFERENCES

Guo, Y., Lamaazi, H., and Mizouni, R. (2022). Smart Edge-Based Fake News Detection Using Pre-Trained Bert Model. In Proceedings of the 18th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob) (437–442). IEEE. https://doi.org/10.1109/WiMob55322.2022.9941689

Rana, V., Garg, V., Mago, Y., and Bhat, A. (2023). Compact BERT-Based Multi-Models for Efficient Fake News Detection. In Proceedings of the 3rd International Conference on Intelligent Technologies (CONIT) (1–4). IEEE. https://doi.org/10.1109/CONIT59222.2023.10205773

Wijayanti, R., and Ni'Mah, I. (2024). Can BERT Learn Evidence-Aware Representation for Low Resource Fake News Detection? In Proceedings of the International Conference on Computer, Control, Informatics and Its Applications (IC3INA) (231–236). IEEE. https://doi.org/10.1109/IC3INA64086.2024.10732358

Majumdar, B., Bhuiyan, M. R., Hasan, M. A., Islam, M. S., and Noori, S. R. H. (2021). Multi-Class Fake News Detection Using LSTM Approach. In Proceedings of the 10th International Conference on System Modeling and Advancement in Research Trends (SMART) (75–79). IEEE. https://doi.org/10.1109/SMART52563.2021.9676333

Salomi Victoria, D. R., G., G., Sabari, K. A., Roberts, A. R., and S., G. (2024). Enhanced Fake News Detection Using Prediction Algorithms. In Proceedings of the International Conference on Recent Innovation in Smart and Sustainable Technology (ICRISST). IEEE. https://doi.org/10.1109/ICRISST59181.2024.10921999

Kamble, K. P., Khobragade, P., Chakole, N., Verma, P., Dhabliya, D., and Pawar, A. M. (2025). Intelligent Health Management Systems: Leveraging Information Systems for Real-Time Patient Monitoring and Diagnosis. Journal of Information Systems Engineering and Management, 10(1). https://doi.org/10.52783/jisem.v10i1.1

Ramzan, A., Ali, R. H., Ali, N., and Khan, A. (2024). Enhancing Fake News Detection using BERT: A Comparative Analysis of Logistic Regression, Random Forest Classifier, LSTM and BERT. In Proceedings of the International Conference on IT and Industrial Technologies (ICIT) (1–6). IEEE. https://doi.org/10.1109/ICIT63607.2024.10859673

Kumar, A. J., Trueman, T. E., and Cambria, E. (2021). Fake News Detection Using XLNet Fine-Tuning Model. In Proceedings of the International Conference on Computational Intelligence and Computing Applications (ICCICA) (1–4). IEEE. https://doi.org/10.1109/ICCICA52458.2021.9697269

Karwa, R., and Gupta, S. (2022). Automated Hybrid Deep Neural Network Model for Fake News Identification and Classification in Social Networks. Journal of Integrated Science and Technology, 10(2), 110–119.

Kadek, I., Bayupati Sastrawan, I. P. A., and Sri Arsa, D. M. (2021). Detection of Fake News Using Deep Learning CNN–RNN Based Methods. ICT Express, 7, 396–408. https://doi.org/10.1016/j.icte.2021.10.003

Madan, B. S., Zade, N. J., Lanke, N. P., Pathan, S. S., Ajani, S. N., and Khobragade, P. (2024). Self-Supervised Transformer Networks: Unlocking New Possibilities for Label-Free Data. Panamerican Mathematical Journal, 34(4), 194–210. https://doi.org/10.52783/pmj.v34.i4.1878

Zabeen, S., Hasan, A., Islam, M. F., Hossain, M. S., and Rasel, A. A. (2023). Robust Fake Review Detection Using Uncertainty-Aware LSTM and BERT. In Proceedings of the IEEE 15th International Conference on Computational Intelligence and Communication Networks (CICN) (786–791). IEEE. https://doi.org/10.1109/CICN59264.2023.10402342