International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 2, Mar-Apr-2026, ISSN (Online): 2395-566X

Fake News Detection Using Deep Learning Techniques

Adlin Jebakumari[1], Uzefa Begum[2], Kathi Harshitha Reddy[3], Mohammed Rameez[4]
[1]Professor, [2]Student, School of Computer Science and Information Technology, Jain (Deemed to be University), Jayanagar, Bengaluru, Karnataka, India.

Abstract- The growth of digital media in recent years has created a major public issue. This is evident in the increase of false information, often called fake news. Fake news refers to any news item that contains false information for the audience. This research project combines traditional machine learning methods with modern deep learning techniques to detect fake news using a hybrid detection system. The news articles will undergo several preprocessing steps: text cleaning, tokenization, stop word removal, and text data normalization for analysis. The team will preprocess the textual data, which will then be converted into numeric data for machine learning and deep learning models. This will use feature extraction methods like tokenization and word embeddings. The project will apply traditional machine learning models to create training data that captures the unique features of fake news and real news articles. The study will also use various deep learning models, including LSTM Networks and BERT. These models will help identify sequential and contextual relationships in articles by understanding complex language patterns and the connections among different types of text data.

Keywords – Fake News Detection, Natural Language Processing (NLP), Machine Learning, Deep Learning, LSTM, BERT, Text Classification.

I. INTRODUCTION

Deciding on a career is one of the most heavy-hitting choices a student will ever make. But today’s job market is a fast-moving target. Students aren't just looking at traditional roles anymore; they’re staring down a future filled with AI, Data Science, and Cyber security—fields that evolve almost weekly. While the opportunities are endless, so is the uncertainty.

Far too often, students make these life-altering decisions based on what their parents expect, what their friends are doing, or whatever incomplete info they can find online. When we choose a career without looking at our actual skills or passions, it leads to a predictable path of stress, burnout, and career shifts down the road.

Standard career counseling usually involves manual tests and a few sessions with a counselor. While helpful, these methods are slow, can be unintentionally biased, and often can’t keep pace with the demands of modern industry. This is where Artificial Intelligence and Machine Learning come in. These systems don't just guess; they analyze massive amounts of data to find patterns and match a student’s "profile"—their strengths and academic history—with the right professional fit. Unlike a static test, an AI model learns and gets smarter over time.

In the context of education in India, where the pressure to succeed is intense and awareness of diverse jobs is often limited, AI can be a great equalizer. It offers every student, regardless of their background, access to scientifically-backed advice. This study looks at how students actually feel about these tools—whether they trust them, find them easy to use, and if they’re ready to let an algorithm help shape their future.

analysed by processing [15] the words that come before and after the text. The model offers the best results for text classification tasks, as the model can detect fake news based on the difference between real news and fake news.

The research needs to include a comparison between the traditional machine learning approach and the advanced deep learning approach to see how accurately the machine learning model can perform the fake news detection task. In the research, a hybrid approach has been introduced that can help organizations use the

Machine Learning classifiers, such as the LSTM model and the transformer model, for fake news classification tasks. In the model, the accuracy measurement, precision measurement, recall measurement, and F1 score measurement are used to assess the model's performance.

The writing style continues to relate to the genuine news sources while making it more challenging to distinguish from their writing patterns. The two parties involved in this process encounter difficulties in proving their points using evidence. Today’s fake news has two major characteristics: it reveals linguistic complexity and lacks any grammatical [7] or

International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 2, Mar-Apr-2026, ISSN (Online): 2395-566X

linguistic errors. The emotional appeal associated with fake news emanates from the blending of genuine and false information. The process of detecting fake news becomes challenging due to its complexity.

The process of detecting fake news faces another challenge in that the materials used in the contents continue to transform. The topics used in the news reported by different media houses undergo a state of transformation in different periods of time. Different media houses have unique writing processes that continue to transform over time to give their writing style a unique touch. The process of detecting fake news becomes challenging due to the complexity [10] it presents to computer systems.

The methods of machine learning are old school, and they require people to develop special features. These are things like how a word is used and how words are used together. These are also things like how people feel about something and how words are put together in a sentence. These are good for numbers and statistics. These methods [4] are not good for seeing the whole picture of how words are connected in a sentence. Sometimes, the words in a sentence change because of something like sarcasm or using the opposite of what is meant. These methods cannot pick that up.

The sentence needs special ways of looking at things that can understand how words are connected. Deep learning methods are great at fixing the problems that the old methods have. One of these methods is called [12] LSTM. This is a good method, as it makes a map of how things fit together as it learns. This method can look at news articles. Follow the information to see where it comes from, and it can look at things that are hidden and things that are not true. This method, however, can only look at information one way.

The issues that I discussed earlier can be addressed through the use of transformer-based models like BERT. This is because the self-attention mechanism used in the transformer [7] model allows the model to look at the sentence as a whole. This means that the model is able to understand the context of the

words. For example, the model is able to see the words that are before a word. After a word. This makes the contextual embeddings of the transformer model suitable for use in natural language processing tasks.

Even though we are aware that classical models and deep models can be used to detect news, we need to compare their effectiveness. Most researchers have concentrated on using a single model at a time. In addition, we need to conduct research to find out the effectiveness of combining sequential learning models and the transformer model, such as BERT.

This paper is trying to fill a hole by testing and evaluating different ways to model things like traditional machine learning classifiers, models that use LSTM and models that use BERT on a dataset of fake news. The goal of this comparison is to find the way to get good results work efficiently and make sure the models can be used in many situations for detecting fake news.

The main things that this paper does can be listed like this:

- It tries out machine learning classifiers to see how well they work.
- It creates a learning model that uses LSTM to look for patterns in fake news.
- It takes a model that uses BERT and adjusts it to work for classifying fake news.
- It compares how well each model works using tests.
- It looks at the bad things about each model when they are used in real life to detect fake news.

By looking at these models in a way we can learn a lot about how well the latest natural language processing techniques work to stop fake news. The results of this paper help make systems that can automatically detect news and reduce its impact on the internet. The paper is about fake news detection and how to make it better. The fake news detection models that are tested in this paper are very important, for news detection.

The detection of news is not just about looking at the words. It also depends on the intent and how fake news spreads. The psychological aspects of news are really important here. Researchers have found that fake news that gets people is more likely to spread fast. So a good fake news detection system needs to be able to see when the content is trying to manipulate peoples emotions.

Another thing that matters when detecting news is the class imbalance problem. In the world there are a lot more real news articles than fake ones. This means that a machine learning model might favor the type of news it sees often. As a result the model does not work well as it should.

When we look at how the model is working we use things like precision and recall and the F1-score [13]. These metrics help us understand if the model is good at detecting news. The detection of news and the class imbalance problem are closely related. The effectiveness of a news detection system depends on how well it can handle the class imbalance problem and recognize fake news.

The detection of news is a tough task and it requires a system that can look at many different things, including the intent and the psychological aspects of fake news. Fake news detection systems need to be able to recognize when content is trying to manipulate peoples emotions and when it is a normal news article. The detection of news is important and it requires a lot of work to make sure that the systems we use are good at it.

- The detection of news depends on many things, including the intent and the psychological aspects of fake news.
- The class imbalance problem is an issue when it comes to detecting fake news.
- A good fake news detection system needs to be able to recognize manipulation and handle the class imbalance problem.

The detection of news is something that we need to be good, at if we want to stop the spread of fake news. Fake news detection systems are important. We need to make sure they are working well.

II LITERATURE SURVEY

Other than what the words say, finding news also depends on why it is being spread and how it is being shared. The psychological aspects of news are really important here. Researchers have found that fake news that gets people is more likely to spread fast. So a good fake news detection system needs to be able to tell when the content is trying to manipulate peoples emotions.

Another big problem in finding news is that there are a lot more real news articles than fake ones. This means that a machine learning model might favor the news articles just because there are more of them. As a result the model is not as good at finding news. To see how well the model is working we need to look at things like precision, recall and F1-score.

People have been studying news detection for about ten years now especially since social media became so popular. At first researchers thought of news detection as a simple problem of classifying text as real or fake based on how it was written. They used techniques to tell real news articles from fake ones.

One of the ways researchers tried to detect fake news was by using old machine learning methods like Naïve Bayes, Decision Trees and Support Vector Machines. These models looked at things like how words were used what part of speech they were and how positive or negative the text sounded. These models were pretty good at detecting [15] news but they only worked well if the features were carefully chosen and the datasets were specific to the topic.

On researchers tried using models that looked at patterns of words and how important each word was. These models were good at finding patterns of words that were typical of fake news. They were also pretty good, at detecting news but they did not work as well when they were tested on new datasets that they had not seen before. Fake news detection is still a problem and fake news is still a big issue.

Researchers started using networks to find fake news when deep learning got better. They used something called Convolutional Neural Networks [15] to look at news articles and find phrases. These Convolutional Neural Networks were better than what people used but they had a problem with understanding things that happen over a long time in a story.

Then people started using something called Recurrent Neural Networks, Long Short-Term Memory models to deal with this problem. Long Short-Term Memory models were good at remembering what happened before and getting results. Some people also used something called word embedding like Word2Vec and GloVe with Long Short-Term Memory models to get better results [1]. Looking at things one after the other was a problem.

When people made something called attention mechanisms it made learning much better at finding fake news. Models that used attention mechanisms got results and were easier to understand. This made news detection using deep learning techniques, like neural networks and attention mechanisms a lot better. Fake news detection, with learning is getting better because of these new techniques, especially attention mechanisms and neural networks.

The field of news detection saw a lot of improvement when people started using transformer-based architectures. This is when the transformer model was introduced. It allowed us to use self-attention mechanisms on all the words in a sequence at the time. Models like [4] BERT were trained beforehand.

They did a lot better on many tasks related to natural language processing including detecting fake news. For example BERT can look at the context between two pieces of text. Understand how they are related. This means it can model contexts that go in both directions. The transformer model can capture the meaning behind the relationship between two pieces of text which's something that traditional models might not be able to do.

Lately people have been trying out approaches that combine different methods. They are using machine learning classifiers and deep contextual embeddings together. Some studies have shown that this approach works well. When we combine architectures we can use the strengths of each model to get better results. Other new ideas include looking at news detection in a more complex way by including things like social context and user behavior. We are also looking at information to help us detect fake news.

Even though we have made a lot of progress in detecting news there are still some things that we do not know. Most studies focus on ways of modeling [5] rather than comparing the different approaches. We still have some problems to solve like when the data's not balanced or when we are dealing with different domains or when the systems are not efficient enough. These are all challenges that we need to overcome if we want to use news detection systems, in real life. The fake news detection systems are still not perfect. We need to work on them. Fake news detection is a task and we need to keep working on it to make it better.

Besides looking at what people write some researchers have also tried to figure out if a news story is fake by looking at who's talking about it and how it is spreading. They think that by looking at how people interact with each other on media and how they share news stories they can tell if a story is fake or not. Some studies have shown that fake news stories spread fast and that people who share them often talk to each other a lot on social media. This means that fake news stories can form a kind of group on social media. Some people have also suggested using computer programs to look at these patterns and figure out what is going on.

Besides trying to detect fake news some researchers are also working on something called stance detection and claim verification. This is like a helper task that can make it easier to figure out if a news story is fake or not. Stance detection is when you try to figure out what someone thinks about a statement or claim. You can tell if they are positive, negative[8] or neutral about it. For example if someone writes about a news story you can try to figure out if they think it is a thing or a bad thing. Fake news detection is, like trying to solve a puzzle and stance detection's one of the clues that can help you figure it out.

Domain adaptation and cross-domain learning have gotten a lot of attention lately. The thing is, models do not work well when the test set is different from the training set. This is because the words used and the context are different. People have found that transfer learning works better when it comes to using models in situations. This is especially true when we use -trained models like the transformer model [11]. These models are good at working with types of data.

There is a problem with fake news datasets. The problem is that there are not examples of some types of news. Some researchers have tried to fix this problem when it comes to finding news. They have used things like SMOTE [9] and cost-sensitive learning. These things help a bit.. If we do not fix this problem our models can become too good at guessing the data they were trained on. This is called overfitting.

People are also talking a lot, about explainability when it comes to news detection. This means we want to know why our models are making decisions. This is a deal because our models are getting more and more complicated. We have found some ways to make our models more explainable. For example we can use attention visualization, SHAP and LIME [3]. These things help us understand what our models are doing.

Explainability is really important when it comes to news detection models. We need to know why our fake news detection models are making decisions.

Recent advances in the field of news detection include using many learners to make predictions together. This helps make the final output more robust as seen in some studies. Using models, like stacking and voting classifiers makes the results more stable and less varied compared to using just one model. Fake news detection is a big area of research and people are looking into something called robustness [6]. This means that some people might try to trick the detection systems by changing the text or adding information that is actually true. They might also use ways of saying the same thing to get around the systems as some researchers have found.

There are ways to detect fake news but not many studies have compared the old ways of doing it with the new ways like using special models such as LSTM and BERT. Many studies have shown that using these models is better, than the old ways but they do not talk about the downsides like how long it takes to compute and how well it works in the real world. Fake news detection is an area of research and people need to look at all the options and see what works best for detecting fake news.

The fake news detection research that is there shows that people have made a lot of progress in figuring out what is real and what is not.. At the same time there are still a lot of problems that need to be fixed. If we compare the ways that people are trying to solve this problem we can get a better idea of what works well and what does not work well for detecting fake news.

This paper is trying to help people understand more about news detection by using old school machine learning methods, deep learning models that use the LSTM approach and models that use the BERT approach to see what works best for detecting fake news. The paper is looking at news detection and trying to find the best way to do it. Fake news detection is a problem and this paper is trying to help solve it by looking at different fake news detection methods.

International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 2, Mar-Apr-2026, ISSN (Online): 2395-566X

III. METHODOLOGY

The current study is about finding a way to detect news. It does this by using kinds of machine learning models and comparing them. The study looks at machine learning models models that use something called LSTM and a model called BERT that uses a transformer.

The overall workflow of the proposed system can be outlined as follows:

Dataset Description
The dataset used for the study has news articles that are labeled as real or fake. Each piece of data has the text of the article and a label that says if it is real or fake. The dataset is split into two parts: one for training and one for testing. This is done so that the models are not biased.

The dataset is mixed up and split into two parts: 80 percent for training and 20 percent for testing. This is done so that the models have data to learn from and can be used in general.

Data Preprocessing
Data preprocessing is an important step in the proposed system. This step has parts:

Text Cleaning: This is where punctuation marks, special characters, website addresses and numbers that are not needed are removed from the dataset.
Lowercasing: All the text is changed to lowercase so that everything is the same.
Stopword Removal: Common words like "the" and "and" are removed from the dataset because they do not add much to what the text's saying.
Tokenization: This is where sentences are broken down into words.

The study uses news and real news to train the models. The goal is to make a system that can detect news. The dataset is important, for this. The system has to be able to look at a piece of news. Say if it is real or fake. The current study is trying to find the way to do this by using different models and comparing them.

- The study is about finding news. It uses computer models for this.
- The study compares the models, special deep learning models, and a BERT model.

This is a summary of the system's functions:

Dataset Description
The dataset contains news articles. These are labeled either real or fake.
Each news article has a set of words and a label.
We have divided the dataset into two sets: training and testing. This ensures the models perform well
We have randomly divided the dataset into 80% for training and 20% for testing.
This way, the models can use the data for training.

Data Preprocessing
This step is very crucial for the success of the project. This includes the following steps:

Text Cleaning: We remove punctuation, special characters, web links, and numbers from the articles.
Lowercasing: We change all the characters of the text to lowercase.
Stopword Removal: We remove the words that do not add much meaning.
Tokenization: We split the sentences into individual words.

BERT-Based Transformer Model
The BERT model is fine-tuned to perform binary text classification. The pre-trained BERT-base model is used and fine-tuned to the fake news text classification problem. The architecture of the model is as follows:

- BERT Encoder Layer Dropout Layer
- Fully Connected Classification Layer

The output of the [CLS] token [13] is taken, and the output of the last layer is passed through the classification layer to predict whether the news is fake or real.

Model Training
All the models are trained using the training data set, and validation is performed using the testing data set. The parameters taken into consideration for the model's performance are:

- Batch size Learning rate Epochs
- Optimizer – Adam optimizer for deep learning models

Early stopping is used to prevent overfitting while training the deep learning model.

Performance Evaluation Metrics
In order to evaluate the performance of the models, various metrics are taken into consideration:
Accuracy – The accuracy of the model.
Precision – The proportion of actual fake news predicted by the model to the total number of fake news predicted.
Recall – The ability of the model to predict actual fake news.
F1 score – The harmonic mean of precision and recall.
Confusion matrices are used to evaluate the performance of the model in more detail.

Comparative Analysis Framework
The last step is to see how well the traditional machine learning models do compared to the LSTM model and the BERT model. We do this in the environment so it is fair. This step helps us understand which of these models is really good at learning from context and getting the job done. We want to know which one is the best, at learning and which one works well. The traditional machine learning models and the LSTM model and the BERT model are all being compared.

IV. CONCLUSION

This paper presents a framework for detecting fake news using machine learning and deep learning techniques.

- The authors compare machine learning techniques with LSTM and BERT-based deep learning techniques to detect fake news.
- The main goal is to evaluate how statistical feature-based and advanced learning techniques work for fake news detection.
- From the results we can see that traditional machine learning techniques give better results with less complexity.
- They have limitations in understanding the relationships between the input data.
- The LSTM-based model does a job in understanding the relationships between news articles and classifies the news more accurately.
- It provides results than traditional machine learning techniques in detecting fake news.
- Machine learning and deep learning techniques are essential, for news detection.

The results show that LSTM and BERT-based models can improve news detection.

The BERT-based transformer model is really good at giving us results because it can look at information in a way that considers the context from all sides. This model is called BERT.

It uses something called self- attention and knowledge it learned beforehand in the field of natural language processing. This helps BERT figure out the differences between real news and fake news. When we compare the BERT-based transformer model to models we see that it does a better job than traditional machine learning and deep learning models at figuring out what is fake news.

Even though we have gotten some results from using machine learning and deep learning models to classify fake news there are still some things we need to think about for the future. For example deep learning models need a lot of computer power to work which's more than what traditional models need. We also want our model to be able to explain what it is doing and not be affected by people trying to trick it.

This research shows us how important it is to understand the context of what people are saying when we are trying to classify fake news. The BERT- based transformer model and other models like it can help us make fake news classification models. These models can be used to find fake news, on the internet and other digital platforms.

REFERENCES

1. N. J. Conroy, V. L. Rubin, and Y. Chen, “Automatic deception detection: Methods for finding fake news,” 2015.GoogleScholar:https://scholar.google.com/scholar?q=Automatic+deception+detection+Methods+for+finding+fake+news+C onroy+Rubin+Chen
2. H. Ahmed, I. Traore, and S. Saad, “Detection of online fake news using n-gram analysis and machine learning techniques,”2017.Google Scholar:https://scholar.google.com/scholar?q=Detection+of+online+fake+news+using+ngram+analysis+and+machine+learning+techniques+A hmed+Traore+Saad
3. W. Y. Wang, “Liar, liar pants on fire: A new benchmark dataset for fake news detection,” 2017Google Scholar:https://scholar.google.com/scholar?q=Liar+liar+pants+on+fire+A+new+benchmark+dataset+for+fake+news+detection+W ang
4. M. Granik and V. Mesyura, “Fake news detection using naiveBayesclassifier,”2017.GoogleScholar:https://scholar.google.com/scholar?q=Fake+news+detection+using+naive+Bayes+classifier+Granik+Mesyura
5. J. Ramos, “Using TF-IDF to determine word relevance in documentqueries,”2003.GoogleScholar:https:/scholar.google.com/scholar?q=Using+TFIDF+to+determine+word+relevance+in+document+queries+Ramos
6. S. Hochreiter and J. Schmidhuber, “Long Short-Term Memory,”1997.GoogleScholar:https://scholar.google.com/scholar?q=Long+ShortTerm+Memory+Hochreiter+Schmidhuber
7. J. Devlin et al., “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” 2019.GoogleScholar:https://scholar.google.com/scholar?q=BERT+Pretraining+of+Deep+Bidirectional+Transformers+for+Lanuag e+Understanding+Devlin
8. A. Vaswani et al., “Attention Is All You Need,” 2017. GoogleScholar:https: //scholar.google.com/scholar?q=Attention+Is+All+You+Need+Vaswani
9. T. Mikolov et al., “Efficient Estimation of Word Representations in Vector Space,” 2013.Google Scholar:https://scholar.google.com/scholar?q= Efficient+Estim ation+of+Word+Representations+in+Vector+Space+ Mikolov

10. D. Bahdanau, K. Cho, and Y. Bengio, "Neural Machine Translation by Jointly Learning to Align and Translate," 2015.GoogleScholar:https:// scholar.google.com/scholar?q=Neural+Machin e+Translation+by+Jointly+Learning+to+Align+and+T ranslate+Bahdanau
11. K. Shu et al., "Fake News Detection on Social Media: A Data Mining Perspective," 2017.Google Scholar:https:/ /scholar.google.com/scholar?q=Fake+News+De tection+on+Social+Media+A+Data+Mining+Perspect ive+Shu
12. S. Vosoughi, D. Roy, and S. Aral, "The Spread of True and False News Online," 2018.Google Scholar:https: //scholar.google.com/scholar?q=The+Spread+of+True+ and+False+News+Online+Vosoughi
13. S. Ruchansky, S. Seo, and Y. Liu, "CSI: A Hybrid Deep Model for Fake News Detection," 2017.Google Scholar:https// scholar.google.com/scholar?q=CSI+A+Hybrid+Deep+M odel+for+Fake+News+Detection+Ruchansky
14. M. T. Ribeiro, S. Singh, and C. Guestrin, "Why Should I Trust You? Explaining the Predictions of Any Classifier," 2016.GoogleScholar:https:/ /scholar.google.com/scholar?q=Why+Should+I+Trust+Yo u+Explaining+the+Predictions+of+Any+Classifier+Ribei r o
15. N. V. Chawla et al., "SMOTE: Synthetic Minority Over-sampling Technique," 2002.Google Scholar:https:// scholar.google.com/scholar?q=SMOTE+Synthetic+Mi nority+Over-sampling+Technique+Chawla