# Multilingual Fake News Detection Using Hybrid Model for Enhanced Computational Efficiency and Performance in Low-Resource Languages

MSc Research Project
Data Analytics

Srinadh Pippalla
Student ID: 23317728

School of Computing
National College of Ireland

Supervisor: Hamilton Niculescu

**National College of Ireland**

**MSc Project Submission Sheet**

**School of Computing**

| | |
|---|---|
| **Student Name:** | Srinadh Pippalla |
| **Student ID:** | 23317728 |
| **Programme:** | MSC Data Analytics **Year:** 2025 |
| **Module:** | Research Project |
| **Lecturer:** | Hamilton Niculescu |
| **Submission Due Date:** | 11/08/2025 |
| **Project Title:** | Multilingual Fake News Detection Using Hybrid Model for Enhanced Computational Efficiency and Performance in Low-Resource Languages |
| **Word Count:** | 7903 **Page Count:** 22 |

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.
ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

**Signature:** Srinadh Pippalla

**Date:** 11/08/2025

**PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST**

| | |
|---|---|
| Attach a completed copy of this sheet to each project (including multiple copies) | □ |
| **Attach a Moodle submission receipt of the online project submission,** to each project (including multiple copies). | □ |
| **You must ensure that you retain a HARD COPY of the project**, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer. | □ |

Assignments that are submitted to the Programme Coordinator Office must be placed into the assignment box located outside the office.

| **Office Use Only** | |
|---|---|
| Signature: | |
| Date: | |
| Penalty Applied (if applicable): | |

# Multilingual Fake News Detection Using Hybrid Model for Enhanced Computational Efficiency and Performance in Low-Resource Languages


## Abstract

The issue of fake news in the digital era, especially on a global and multilingual landscape, poses a serious threat to information integrity, especially in low-resource linguistic contexts. The conventional fake news detection tools are usually developed based on high-resource languages and low resource languages are underrepresented because of the lack of annotated data and computational resources. This study seeks to solve the urgent problem of an effective and scalable multilingual fake news detection system that serves well in both high- and low-resource languages. The research suggests hybrid deep learning model which integrates a state-of-the-art transformer model, DeBERTa, which has demonstrated an ability to represent multilingual data with a Bidirectional LSTM network to improve the temporal comprehension and classification accuracy. A multilingual corpus of seven languages was compiled and cleaned. To set benchmarks, the use of classical machine learning models (Logistic Regression, Naive Bayes, Random Forest) and deep learning baselines (GRU, CNN) was implemented. The experimental findings proved that the suggested DeBERTa + LSTM hybrid model was much better than baselines, showing a high overall accuracy of 97%, and performing especially well when dealing with low-resource languages. The hybrid mode did not only improve the accuracy of classification but also maximized the computational performance by freezing the transformer backbone, which can be used in practice in resource-limited settings. The study is an addition to the emerging discipline of multilingual NLP as it introduces a scalable, powerful solution to the problem of misinformation detection. It opens up the path towards future work in cross-lingual fake news classification, low-resource language adaptation, and deployment-ready hybrid architectures.

Table 1: Table of Acronyms and Definitions

| Acronym | Definitions |
|---|---|
| NLP | Natural Language Processing |
| ML | Machine Learning |
| DL | Deep Learning |
| AI | Artificial Intelligence |
| DeBERTa | Decoding-enhanced BERT with Disentangled Attention |
| BERT | Bidirectional Encoder Representations from Transformers |
| RoBERTa | Robustly Optimized BERT Approach |
| mBERT | Multilingual BERT |
| XLM-R | Cross-lingual RoBERTa |
| XLM | Cross-lingual Language Model |
| LSTM | Long Short-Term Memory |
| BiLSTM | Bidirectional Long Short-Term Memory |
| GRU | Gated Recurrent Unit |
| CNN | Convolutional Neural Network |
| RNN | Recurrent Neural Network |
| LR | Logistic Regression |
| NB | Naive Bayes |
| RF | Random Forest |
| DT | Decision Tree |
| SVM | Support Vector Machine |
| TF-IDF | Term Frequency–Inverse Document Frequency |
| EDA | Exploratory Data Analysis |
| ReLU | Rectified Linear Unit |
| GPU | Graphics Processing Unit |
| CPU | Central Processing Unit |
| LIAR | Benchmark dataset for fake news detection |
| FakeNewsNet | Dataset for fake news research combining content and social context |
| Fakeddit | A multimodal fake news dataset derived from Reddit |
| Keras | Open-source deep learning API written in Python |
| Adam | Adaptive Moment Estimation (optimizer used in training neural networks) |
| LabelEncoder | A scikit-learn utility for encoding categorical labels into numeric form |
| Softmax | A function that converts logits into probabilities |
| Tokenizer | A utility/tool to convert text into tokens for processing in NLP models |
| Truncation | Reducing text to a fixed length (used during tokenization) |
| Padding | Adding extra values to make sequences the same length for model training |
| Dropout | A regularization method to prevent overfitting in neural networks |

# 1. Introduction

## 1.1 Background

With the advent of digital communications everywhere, information sharing has become instantaneous and cross-border. Even though this technological progress has many positive effects, it has also raised the level of misinformation and news which is fake considerably. Availability of the social media and online news sites across the world has made fake stories spread among people very quickly, creating a certain impact on the opinion of people, political stability, and even the health of the citizens. This difficulty is further enhanced in low-resource language and multilingual environments, where there is a low level of development of misinformation detection systems or their absence (Alghamdi et al., 2024).

Typically, fake news detection systems have been created in high-resource languages like English, where labeled data and computational resources are so plentiful. Nevertheless, such systems prove to be less effective in low-resource linguistic settings when annotated data are rare and computational infrastructure is weak (Raquiba et al., 2024). As a result, the performance of fake news detection in these scenarios is poor, it lacks scalability, and it is resource-intensive (Amanzadeh Oghaz et al., 2022).

The most recent developments of NLP, especially via transformer-based models, including BERT, RoBERTa & DeBERTa, led to dramatic gains in language understanding and classification. DeBERTa (Decoding-enhanced BERT with disentangled attention) provides better language modeling capability and fewer parameters than previous transformer models, and it achieves better performance. Its multilingual version is able to process texts in multiple languages and thus, it can be utilized as a basis of multilingual fake news detection. Nevertheless, even the state-of-the-art transformers have high computational requirements and can find it difficult to deal with long-range dependencies in sequential data.

For identifying constraints, the proposed research offers a hybrid model in which this study combine DeBERTa with LSTM networks. DeBERTa successfully models multilingual text representations, whereas LSTM networks are particularly good at modelling temporal and sequential data in text. The integration of the two is expected to result in the two architectures capitalizing on each other to enhance accuracy, lower computational overheads, and also guarantee applicability in real-time, even when using low resources.

## 1.2 Problem Statement

Although significant advances in NLP have already been made, state-of-the-art detection of fake news tools are still ineffective in multilingual and low-resource settings. There are two challenges:

- **Lack of Linguistic Resources**: A lot of languages, like Swahili, Hindi, Vietnamese, and Indonesian, do not have adequate annotated data sets, pre-trained models and NLP tools that can be used to develop powerful machine learning systems.

- **Computational Inefficiency**: Deep learning algorithms such as transformers are computationally expensive and need a lot of processing power and memory to run, making them inapplicable in real-time or resource-limited settings.

Consequently, an effective and scalable, accurate fake news detection system is urgently required, particularly in the low-resource setting.

## 1.3 Research Objectives

Central goal of the study is proposing hybrid multilingual fake news detection model that will improve the efficiency of the computations and performance of the model in identifying misinformation in low-resource languages. The particular goals are:

- To create and develop a hybrid architecture which would integrate the multilingual representation power of DeBERTa with the strengths of sequential learning of LSTM.
- To maximize the model in terms of less consumption of computational resources and high detection accuracy.
- For testing performance of model in both high-resource languages & low-resource languages.
- To guarantee that model can be scaled to work in real-time fake news detection applications.

## 1.4 Research Questions

This study wants to answer these research questions below:

1) How effectively can DeBERTa–LSTM integration improve the accuracy of fake news detection in multilingual datasets, particularly in low-resource languages?
2) How effective is proposed hybrid model compared to standalone models (Random Forest, Logistic Regression and Naive Bayes etc.) in terms of performance?

## 1.5 Organization of the Research Study

The structure of this thesis is the following:

**Chapter 1: Introduction** – Gives background, problem statement, objectives and significance of study.
**Chapter 2: Literature Review** – Reviews prior research done on fake news detection, multilingual NLP and hybrid model architectures.
**Chapter 3: Methodology** – Describes the dataset preparation, architecture of the model, training process, and metrics of evaluation.
**Chapter 4: Experimental Results and Discussion** – This chapter gives and discusses performance of hybrid model on different languages.
**Chapter 5: Conclusion and Future Work** – Concludes contributions and states the possible ways of further research.

# 2. Literature Review

## 2.1 Introduction

The growing number and pace of online information distribution have aggravated the problem of discovery and prevention of misinformation and fake news spreading. Although the emergence of automated fake news detectors has brought some reprieve, their success is also largely dependent on the language that they are run on and computational resources at their disposal. The majority of existing systems are optimized towards high-resource languages, such as English, & have high processing requirements, and therefore cannot be easily applied in real-time or low-resource settings. In this chapter, the literature is critically reviewed in five main categories: (1) conventional machine learning-based approaches of detection of fake news, (2) DL-based methods, (3) multilingual and low-resource NLP models, (4) hybrid modeling methods, and (5) strategies of computational efficiency in NLP.

## 2.2 Traditional Approaches to Fake News Detection

The majority of systems for detection of fake news were using conventional machine learning prior to the advent of deep learning. Such methods concentrated on feature engineering with the help of statistical and linguistic hints like n-grams, TF-IDF scores, punctuation patterns, and sentiment polarity. The classical classifiers such as Support Vector Machines (SVM), Decision Trees, and Naive Bayes algorithms were trained with extracted features (Fagundes et al., 2024). Despite the fact that these approaches proved computationally cheap and interpretable, their performance was narrowed down by the fact that they had a shallow knowledge of semantics and context. As an example, a classifier can misclassify a genuine news as a fake in case the language employed was sensationalist or emotional. Besides, these methods needed language-specific characteristics, and therefore, could not be applied across languages and domains and impeded their scalability.

## 2.3 Deep Learning Models in Fake News Detection

As deep learning continued to grow, more complex models started to outperform their traditional counterparts in the task of detecting fake news because they no longer required manual feature engineering as they could learn directly out of raw text data. Recurrent Neural Networks (RNNs) and their variants, notably Long Short-Term Memory networks (LSTMs), was one of the early deep learning models to gain popularity as a result of their capability to model sequential dependencies in language. Vanishing gradient problem in standard RNNs was solved by LSTMs proposed by Hochreiter & Schmidhuber (1997) that were found to be useful in modeling long-range dependencies in text. LSTMs have also been applied to detect narrative structures, repetitive motifs, and style patterns that can be used to identify misinformation in context of fake news detection (Airlangga et al., 2024; Kalusivalingum et al., 2021). Nevertheless, although LSTMs provided some gains in processing sequential data, they did not go far enough to comprehend deep contextual meanings.

The advent of transformer models, in particular BERT, was paradigm shift in NLP. The attention mechanism enabled BERT to take into account all words in a sentence at once, which made it capture a broader context than LSTM models (Khan & Kefalas et al., 2023). But that cost was large memory consumption, slow training, and the necessity of large datasets. The use of BERT-based models in low-resource environments has therefore been a challenge.

## 2.4 Multilingual Processing in Fake News Detection

Since misinformation is global, multilingual fake news detection is becoming more significant. Such models, such as Multilingual BERT (mBERT) and XLM-R (Cross-lingual RoBERTa) have tried to apply transformer-based language understanding to more than 100 languages by training on large-scale multilingual corpora (Conneau et al., 2019; Gurgurov et al., 2024).

Although these models have been promising in cross-lingual generalization, research results have demonstrated the inconsistency of the model performance across languages. Languages with high resources are much more successful than low-resource ones as they have more coverage of vocabulary, richer training data, and more optimized pre-processing tools (Han, 2022). Also, the performance of multilingual tokenizers is usually influenced by their quality. As an example, a common tokenizer can tokenize words in low-resource languages inadequately, which leads to suboptimal embeddings.

Two new solutions are cross-lingual transfer learning and zero-shot learning which are developed for improving performance of low-resource languages and use knowledge of high-resource language models (zhang et al., 2025, Kumri et al., 2021). These methods are however still in need of optimization to reach parity in terms of accuracy and reliability.

## 2.5 Low-Resource Language Processing

The key limitations of low-resource language processing are insufficient abundance of annotated data, the lack of linguistic tools, and underrepresentation in benchmark corpora. Although small quantities of task-specific data may be sufficient for fine-tuning pre-trained multilingual models, results in genuinely low-resource languages are variable.

One of the most popular approaches has been transfer learning, in which knowledge acquired in high-resource languages or tasks is transferred to low-resource environments in order to enhance performance (Han, 2022). Unsupervised machine translation and cross-lingual word embeddings have also been utilised to address the resource gap, although they rely strongly on language similarity and parallel corpus quality.

In addition, there are not many datasets dedicated to fake news in low-resource languages. A majority of benchmark datasets, including LIAR, FakeNewsNet, and Fakeddit, are English-biased (Fagundes et al., 2024), and it is hard to apply the results to other languages with different syntactic and cultural peculiarities.

## 2.6 Hybrid Models for Fake News Detection

For capturing advantages of more than one architecture, scholars started to create hybrid models, combining transformers with more traditional sequence models. As an example, rich contextual embeddings could be learned using transformer-based models such as BERT or DeBERTa & fed into LSTM layers to detect temporal patterns. This enables the model to grasp the meaning of individual words as well as the entire form of the story.

Recent research has revealed that these hybrid architectures are able to perform better than standalone models in text classification, sentiment analysis, and fake news detection. As an example, Assiri et al. (2024) suggested a DeBERTa-GRU hybrid model that showed increased accuracy of sentiment classification, especially in the case of long and noisy text.

Such hybrid structures can be useful in fake news detection to identify the stylistic and sequential patterns of misinformation: tone shifts, redundancies, or contradictions. Nonetheless, there are little studies that tested these models in multilingual or low-resource environments, and it is an open and useful research field.

## 2.7 Computational Efficiency in NLP

Transformer-based models are resource-heavy and thus not always feasible to use in real-time or in low-infrastructure environments. Training of such large models as BERT and RoBERTa may require days on powerful GPUs, and even inference may be slow when working with large data streams.

There are a number of approaches that have been suggested to handle this problem. Model distillation compresses large models into smaller and faster ones (e.g. DistilBERT), while pruning and quantization shrink the model size and the inference time by removing redundant parameters (Sanh et al., 2019).

A more recent development, DeBERTa (He et al., 2021), has greater performance using fewer parameters by disentangling the encoding of position and content in its attention mechanism. This renders it a good candidate to apply in hybrid models where efficiency and accuracy are the priorities.

By combining DeBERTa with an LSTM, it is possible to front-load semantic understanding (by means of DeBERTA efficient encoding) and offload sequence pattern recognition to the less computationally demanding LSTM. The design is suitable to be deployed in resource-scarce environments, without compromising on detection accuracy considerably.

## 2.8 Identified Research Gaps

Although the research on fake news detection is quite extensive, there are still considerable gaps in the given field:

- Multilingual models tend to be very inconsistent across languages, and low-resource languages suffer a drop in accuracy.

- Large number of systems for detection of fake news are constructed on English or other high-resource languages, which have left most of the linguistic communities underserved.

- There are limited models that merge transformer and sequential learning architecture to detect fake news in multilingual environments.

- Computational efficiency is a problem that is not taken into consideration in academic benchmarks, though it is a major issue in practice, in real-time deployment.

The hybrid architecture of the current study fills these gaps by providing an efficient and scalable multilingual fake news detection framework with the semantic depth of DeBERTa and the time-sensitive nature of LSTM.

### 2.9 Summary

This chapter discussed the development of fake news detection systems, beginning with traditional ML methods to more sophisticated DL and transformer-based models. It commented on issues of multilingualism and low-resource language processing and named hybrid models as a potentially good way to strike balance between accuracy and computational efficiency. The review identifies the existence of a critical research gap in the area of scalable, real-time multilingual fake news detection and in particular underrepresented languages. The next chapter will include the methodology of the proposed DeBERTa-LSTM hybrid model implementation and assessment.

# 3. Methodology

### 3.1 Introduction

The chapter includes the approach that was selected to build a powerful multilingual fake news detector, especially when it comes to low-resource languages. The main idea of the current research is to build a hybrid architecture combining the strong multilingual contextual representation of DeBERTa with sequential modeling abilities of LSTM networks. To do so, this study use a step-by-step approach that includes data collection and preprocessing, exploratory analysis, and the training of both baseline & DL approaches and development of proposed hybrid architecture. All the phases are described in a well-organized way below, in order to be transparent and reproducible.

### 3.2 Dataset and Language Selection

The data used in the study is a multilingual, cleaned up news corpus of real and fake news. It also contains materials in seven languages including English, Hindi, Swahili, Indonesian, Vietnamese, Turkish, and Japanese. These were carefully chosen languages in order to give a balanced selection of both high resource (English and Hindi) and low resource languages (Swahili and Vietnamese). Data were loaded and checked on structure, consistency and balance. It has three columns, which include the text of the article, the label that is associated with the article (real or fake) and the language code.

**Dataset:** https://www.kaggle.com/datasets/begonil/multilingual-fake-news-detection

Dataset was filtered first to only keep the entries of the selected target languages. This was necessary as a way of simplifying the analysis and making sure that the models are trained on a homogeneous multilingual distribution. The entries were randomized to eliminate the order bias after filtering, and the entries were ready to be divided into training and testing subsets.

### 3.3 Data Preprocessing

Textual data, and particularly multilingual ones, are highly dependent on effective preprocessing. The text data were cleaned as the first step of the preprocessing pipeline. Regular expressions were used to remove non-alphabetic characters and all characters were converted to lowercase. This was an important step to standardize the inputs and decrease noise in terms of punctuation, special characters or case sensitivity. Moreover, the missing or unusual values were checked within the dataset. Entries that had null values, empty text, or infinite values were deleted in order to retain data integrity. Also, the length of each text article was estimated in order to help determine the maximum length of sequences in deep learning models. A dictionary was used to map languages into their full names in order to have better interpretability. The language labels were categorical values that were converted to numerical values to be used in training the model using LabelEncoder utility. The transformation allowed the use of regular input to both classical machine learning and neural network classifiers.

## 3.4 Exploratory Data Analysis

An exploratory analysis was done before training the models to gain a better insight into the structure and properties of the dataset. A number of visualizations were produced. Figure 1 showed a language distribution chart to determine how many samples were present in each language. This was especially necessary to evaluate class imbalance, that may negatively impact model performance.

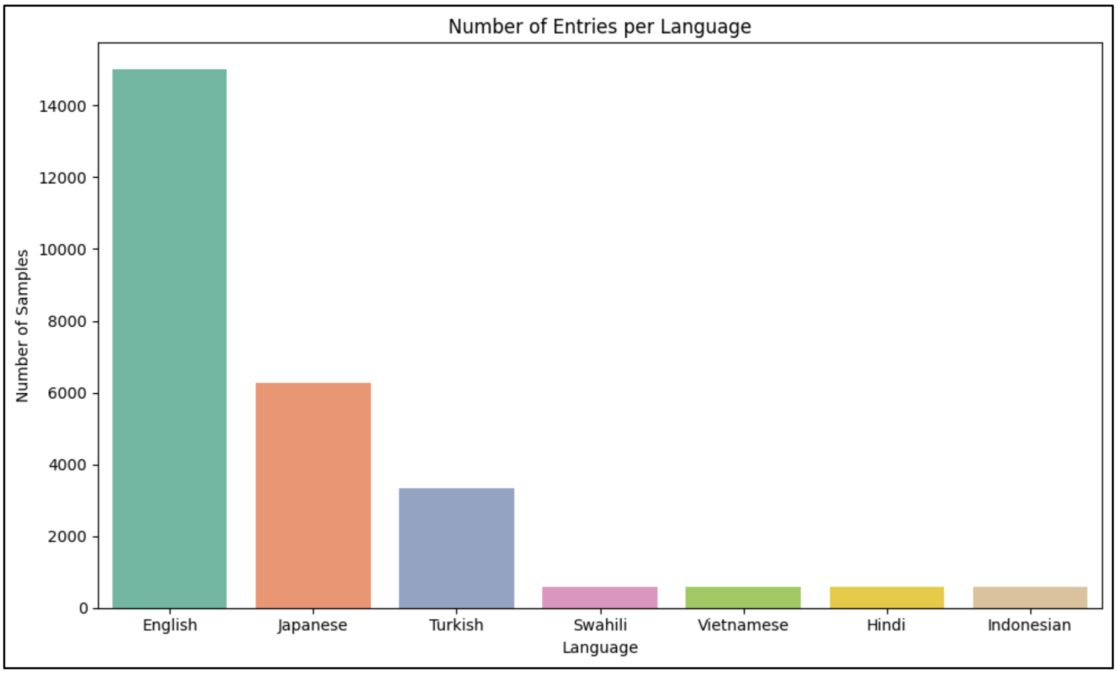


**Figure 1: Distribution of the Languages in Multilingual Fake News Dataset**

Besides the distribution of languages in general, the distribution of labels (fake vs. real) were also analyzed in each language to identify the extent of misinformation in various linguistic contexts in Figure 2.

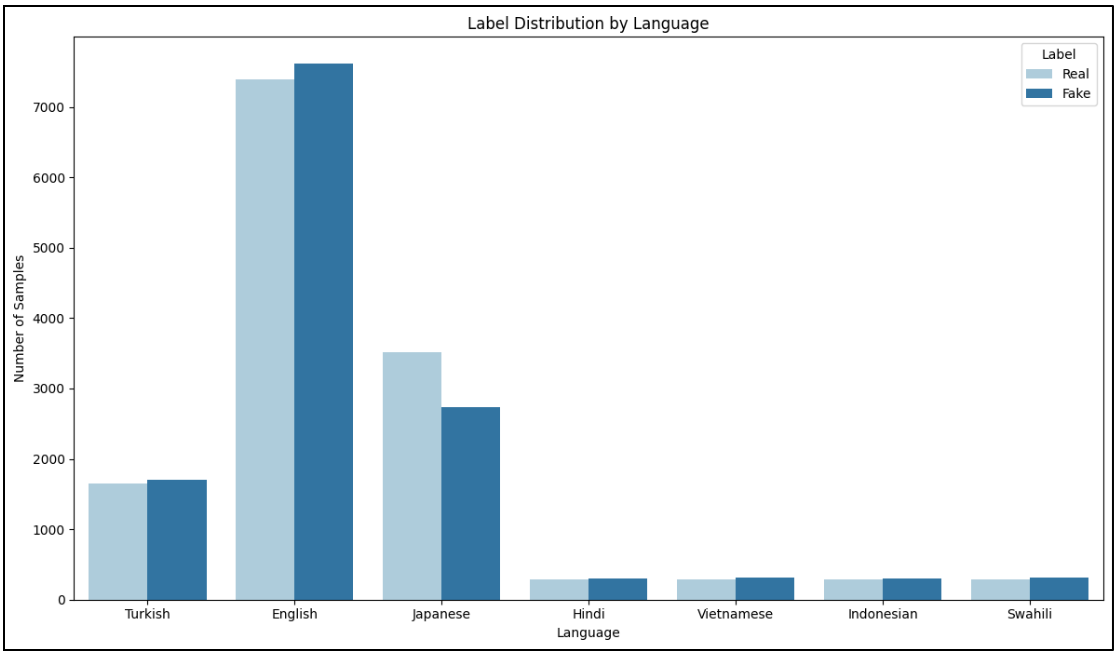


**Figure 2: Label Distribution for each Languages**

The creation of word clouds was also helpful in the analysis since it gave a clear visual overview of the most common words in the English portion of the data. This assisted in finding key words and lexical patterns that are prominent to each class, and gave an intuitive idea on the thematic focus and distinguishing language features, as illustrated in Figure 3.


**Figure 3: Word Cloud for English News Text**

Moreover, the analysis of the text length provided the average and maximum number of tokens per entry that affected the choice of the length of sequence padding of specific models, such as GRU, CNN, and LSTM.

## 3.5 Implementation of Classical Machine Learning Models

For setting up a baseline in the performance, some classical ML approaches were applied, like Logistic Regression, Naive Bayes, and Random Forest classifiers. Textual data was cleaned and pre-processed and vectorized utilizing Term Frequency-

Inverse Document Frequency (TF-IDF) representation where textual data was converted to numerical vectors but the importance and the frequency of words in the corpus is maintained.

Data was, thereafter, divided into sets for training & testing. All models have been trained on TF-IDF features of the news articles and predictions have been made on test set. Evaluation metrics, like accuracy, precision, recall and F1-score, were calculated using classification reports. These measures gave a clue of how well the performance was at the beginning and were used as comparative benchmarks by more sophisticated models.

## 3.6 Implementation of Deep Learning Models

The effectiveness of classical models was tested on two deep learning architectures, namely the Gated Recurrent Unit (GRU) model & Convolutional Neural Network (CNN) model, to evaluate their performance in dealing with multilingual text sequences.

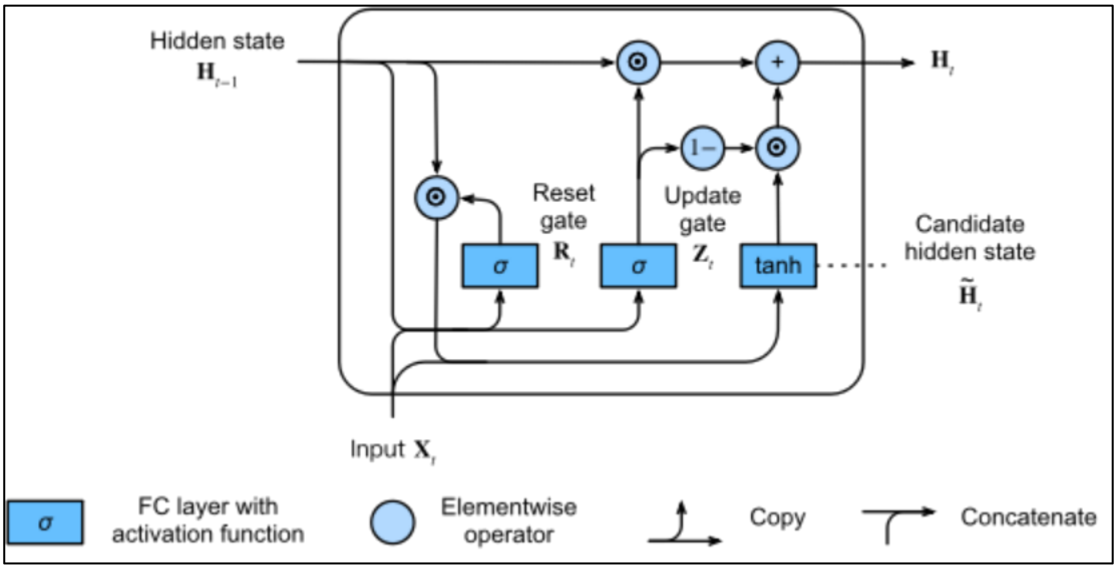


**Figure 4: GRU Model Architecture**

In these models, the textual data was tokenized and transformed to preserve only the 10,000 most common words with Keras Tokenizer was to meet the goal of vocabulary richness and computational efficiency, to capture the most important lexical information and to minimize noise introduced by rare words. Padding to a constant length of 250 tokens provided a fixed input size that is necessary to train neural networks in batches. The value 250 was chosen after an exploration of the distribution of the length of the dataset sequences and represents a compromise between truncation and padding the vast majority of texts but not too much.

The GRU model included an embedding layer, a GRU layer with 64 units, and dense output layer with softmax activation to deal with multi-class classification over the seven language categories. Equally, CNN model had embedding layer, 1D convolutional layer with 128 filters, global max pooling layer, and dense output layer.

The Adam optimizer was used to compile both models and the sparse categorical crossentropy loss was used to train the models. They were trained on a single epoch with batch size of 512 & performance was checked on a validation set that was held out. The outcomes of these models gave a better insight into the performance of sequential and convolutional architectures on multilingual fake news data.

### 3.7 Proposed Hybrid Model: DeBERTa + LSTM

The major novelty of the research presented is the creation of a hybrid model that integrates the power of DeBERTa with a bidirectional LSTM. DeBERTa was chosen because it demonstrates success in predicting deep contextual interactions in various languages, and the LSTM block was added to provide long-term dependencies in the textual sequence that transformers could overlook.

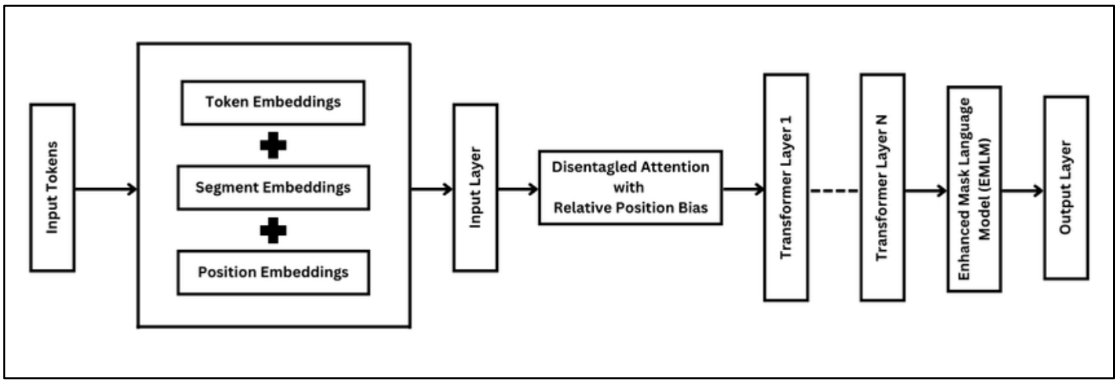


**Figure 5: DeBERTa Model Architecture**

The implementation started with the loading of the pre-trained model of microsoft/deberta-base and tokenizer. The transformer model was frozen and only its contextual embeddings were taken out for the input text in order to decrease computational complexity. Truncation and padding to 128 tokens (maximum sequence length) were used for carrying out tokenization.

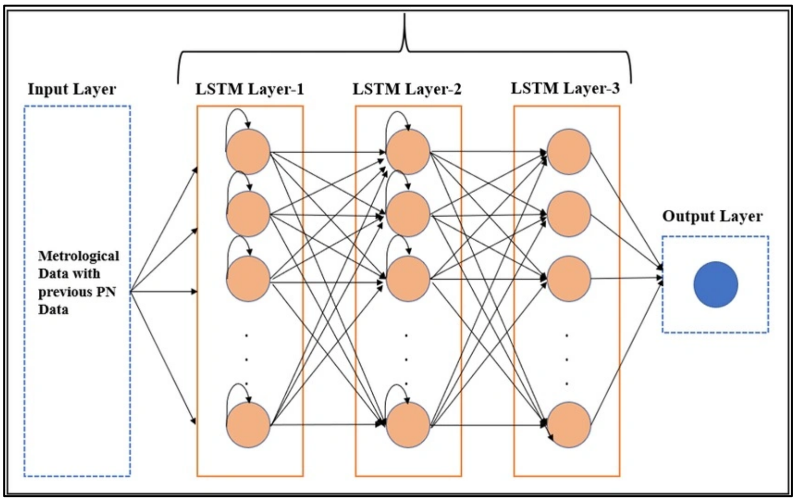


**Figure 6: LSTM Model Architecture**

The final hidden states of DeBERTa (embeddings on the token level) were then fed to a bidirectional layer of LSTM. The number of units in this layer was set to 128 and was followed by a dense layer with ReLU activation and dropout layer to avoid overfitting. Lastly, a softmax layer was used to generate the probabilities of the classes as they relate to the seven language labels.

The hybrid model was trained with categorical crossentropy loss, Adam optimizer and learning rate of 2e-5. The model was trained in 3 epochs and 32 batch size. Training was done with both accuracy and loss tracked on a validation set. The validation set was used to make predictions after the training, and the performance was assessed with the help of classification metrics and confusion matrices.

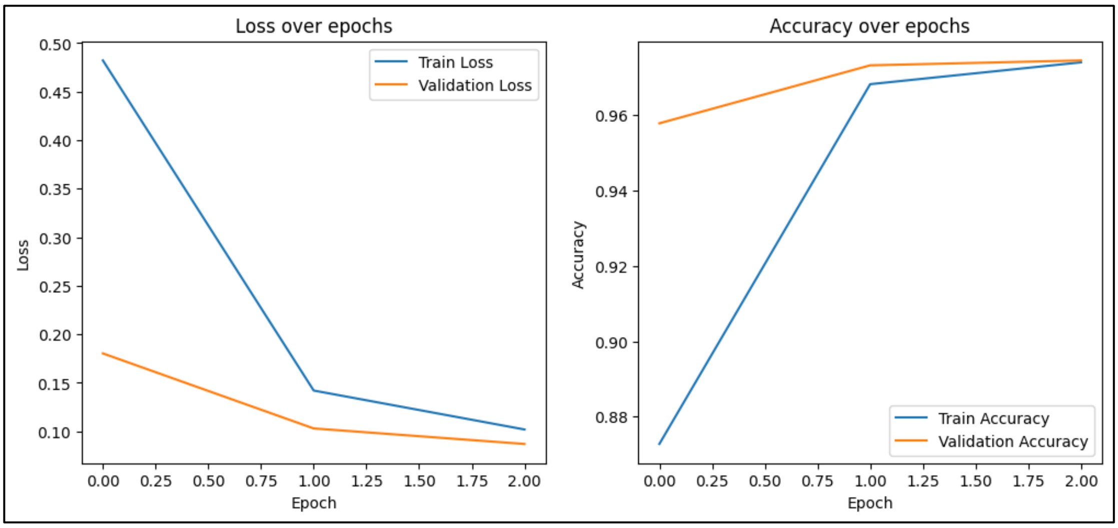


**Figure 7: Proposed Model Training Loss and Validation Accuracy**

## 3.8 Evaluation Metrics

In order to compare all models, classical, deep learning, and hybrid, a standard set of metrics was used. The most important measure of overall model performance was accuracy. Nonetheless, since the imbalance of classes within certain language categories was observed, & also calculated precision, recall, and F1-score. Plotting of confusion matrices was used to examine the nature of misclassifications that each model produced.

In the case of the hybrid model specifically, accuracy and loss training and validation curves were plotted to monitor the learning trend with respect to epochs, as well as to ensure that the model was not overfitting or underfitting. These plots were important to understand the stability of the model and the ability to generalise.

## 3.9 Comparative Analysis

In order to summarize the results, the accuracies of all models were summarized and compared in a bar chart. This comparison showed the comparative advantages of both techniques and the immense performance improvement realized in the proposed hybrid model. Although classical models provided efficiency in computing, they were not accurate on multilingual data. The performance was enhanced with deep learning models at a greater computational cost. The DeBERTa + LSTM model was more accurate than the two baselines with a manageable computational overhead because of the frozen transformer architecture.

## 3.11 Summary

This chapter gave a detailed description of the methodology taken in this study. Beginning with data preparation and exploratory analysis, it outlined the use of baseline ML & DL approaches and designing new hybrid architecture. The selected hybrid model takes advantage of the contextual power of the transformers and the sequential nature of LSTMs to achieve high performance in detection of fake news in many languages especially in the low-resource language environments. The second chapter shows the experimental findings and gives the implications of findings..

# 4. Implementation of Design Specification

## 4.1 Introduction

This part explains how the design specifications of the hybrid multilingual fake news detection system are implemented. The objective is to convert the given proposed hybrid architecture to a running system, including the technical decisions, architectural elements, tools, the workflow of the system, and the configuration. The fusion between DeBERTa and LSTM has been carefully designed to make it computationally efficient and lead to a high classification accuracy of high and low resource languages. This section of implementation gives a systemic dismantling of every layer in the system, which makes its replication and development possible.

## 4.2 System Architecture Overview

The proposed system is designed as a modular architecture where each module performs a certain task in the fake news detection pipeline. Its design is based on a hybrid deep learning approach that involves a frozen DeBERTa encoder to multilingual feature extraction coupled with a Bidirectional LSTM network to learn sequences as shown in Figure.

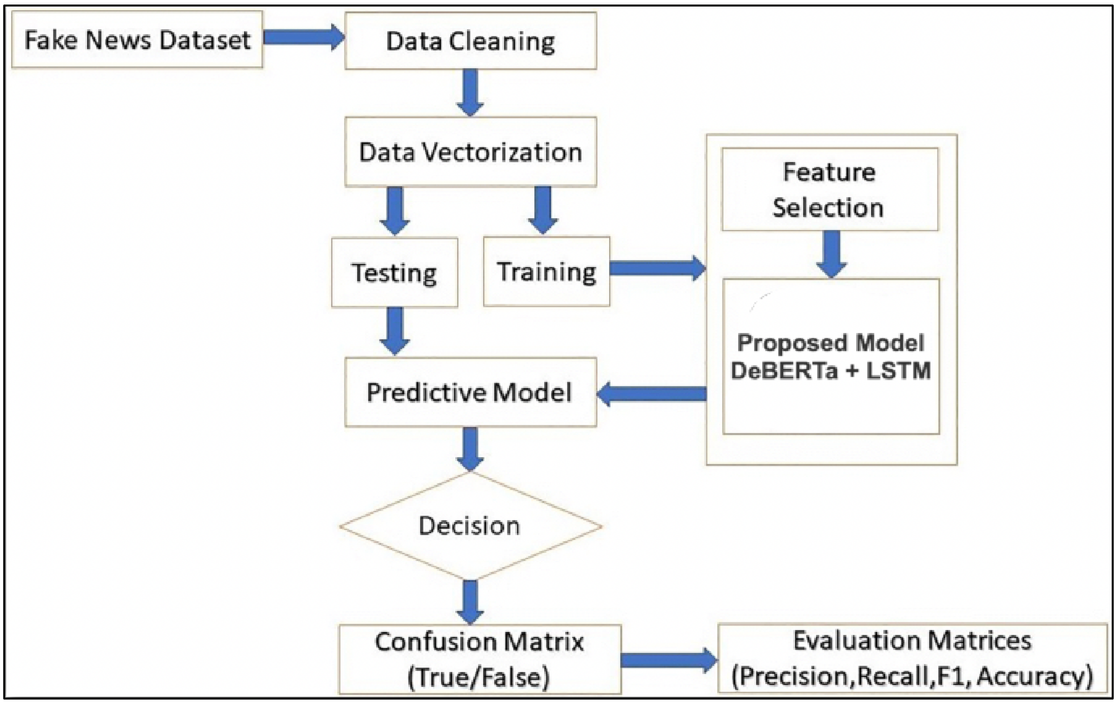


Figure 8: Architecture Design for detecting Multilingual Fake News

## 4.3 Summary

This section explained the deployment of the suggested hybrid framework of multilingual fake news detection and focused on its design rationale, technical implementation, and overall data flow. The combination of the frozen DeBERTa encoder and Bidirectional LSTM layer was aimed at achieving both the deep contextual understanding and pattern recognition in the sequence with the computational efficiency of the model. Every step, including preprocessing and tokenization, embedding extraction, and classification, was designed to accommodate multilingual data and to suit low-resource language limitations.

Transformer weights that were frozen during training greatly decreased time and resources needed for training & LSTM component provided a temporal sensitivity that is often absent in transformer-only models. The implementation is implemented using modular design and efficient training strategies, which allow its use not only in the case of batch processing of large datasets but also in the case of potential real-time usage. The system is a viable and scalable mechanism of dealing with the emerging menace of fake news within a linguistically diverse setting.

# 5. Experimental Results and Discussion

The chapter gives experimental results of the given multilingual fake news detecting system. The comparison involves classical ML & baseline DL approaches and hybrid architecture based on the DeBERTa transformer and an LSTM classifier. The aim of this experimental setting is to determine the ability of each of the models to work with multilingual and low-resource languages data in regards to classification accuracy, macro level performance statistics, and computational efficiency.

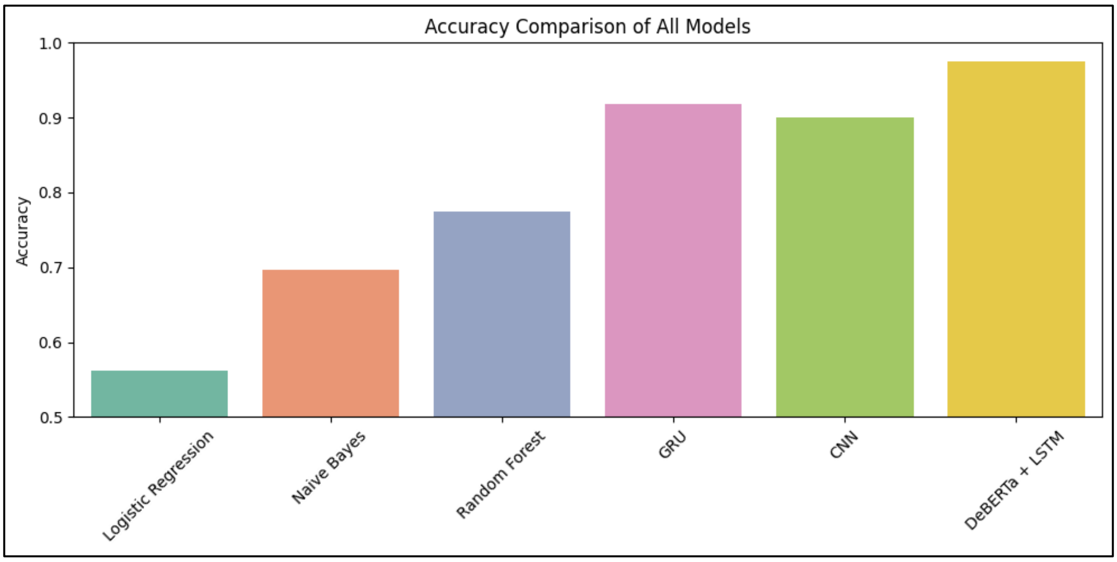


**Figure 9: Comparison of Experiment Results**

## 5.1 Performance of Classical Machine Learning Models

First baseline consisted of 3 classical ML approaches: Logistic Regression, Naive Bayes, and Random Forest. TF-IDF features on the preprocessed text data was used to train these models.

Logistic Regression produced a general accuracy of 56.18 %. Although this performance seemed to be relatively effective, additional examination showed that there was a significant pitfall, namely, that the model favored the majority group considerably more than the minority, which in this case was the English group. Languages that have less training samples like Swahili, Vietnamese and Hindi were labeled with close to zero precision, recall and F1-score. This means that model is incapable of generalizing well when faced with extreme problems of class imbalance and a lack of features to learn the semantic distinctions.

Naive Bayes classifier did better on this baseline with a total accuracy of 69.71%. It showed comparatively greater performance in some of the more represented languages such as English and Turkish. But it did not work at all in the case of underrepresented languages like Swahili and Vietnamese. This was also indicated by the macro-averaged values of precision and recall, which were very biased towards the high-resource languages.

Random Forest produced the highest results out of the classical models with an accuracy of 77.37%. It has managed to deal with some of the multilingual classes like Japanese and Turkish better than the earlier models. Nevertheless, it did not provide generalizability of all language classes. Most notable, the model performed poorly on low resource languages, with Swahili, Vietnamese and Hindi having precision and recall values of near zero. These results show the inadequacies of classical models in dealing with the complexities of multilingual text and especially in the absence of data.

## 5.2 Performance of Deep Learning Models (GRU and CNN)

The next step of the experiment was to test deep learning-based models: Gated Recurrent Unit (GRU) and Convolutional Neural Networks (CNN). These models were trained on padded sequences prepared based on tokenized text and tested based on categorical cross-entropy loss as multiclass classification.

The GRU model showed a vast improvement over the classical methods with a total accuracy of 91.75%. It showed high capacity to label English and Japanese text accurately. Other underrepresented languages that also experienced the temporal learning capabilities of the model were Indonesian and Turkish. Nevertheless, such extremely low-resource classes as Hindi or Vietnamese remained challenging to the GRU since the support was too low or the linguistic peculiarities were not enough to make the proper generalization. However, it had a macro-averaged F1-score of 0.65, which showed that it had a reasonable balance between the classes, a big jump in comparison to the conventional models.

CNN model scored a bit less, 89.99%. Similar to GRU, it performed well on majority languages but lacked strength in long-range dependencies as one of the possible reasons why it performed relatively poorly on classes with a smaller number of data points. It did rather well at classifying English and Turkish samples but did not learn useful patterns in low-resource languages such as Swahili and Hindi. Although it has been proven to be effective in extracting local features of text, CNN was not enough in the subtlety needed in the multilingual fake news detection.

## 5.3 Performance of the DeBERTa + LSTM Hybrid Model

Hybrid model, in which the multilingual transformer DeBERTa was combined with a bidirectional LSTM network, was able to surpass all other methods by far. The model had a total accuracy of 97.00% which was the highest of all systems tested.

Additionally, it exhibited a high performance almost in all language classes including some of the low-resource classes.

The hybrid model displayed near-perfect or perfect classification in English, Indonesian, Turkish and Swahili, with precision and recall scores frequently being close to 1.00. Japanese and Vietnamese were also correctly classified with a high accuracy, but sometimes were misclassified probably because of semantic overlaps or transliteration inconsistencies. Hindi was also a problematic case because it was poorly represented in the data; nevertheless, the hybrid system at least tried to learn using the limited number of samples.

Among the reasons why this model performs better, one can mention the frozen DeBERTa transformer, which offered rich and pre-trained multilingual embeddings, and LSTM, which captured the sequential dependencies. The architecture used an extraction of contextual representations and then processes them in sequence, which allows the architecture to unite the advantages of transformers and RNNs. Moreover, freezing the DeBERTa layers contributed to the minimization of computational load, and thus, the model could be applicable in more resource-limited settings.

## 5.4 Comparative Summary and Analysis

To give a brief description of performance of models, the performance of all models was summarized with significant evaluation metrics which included accuracy, macro-averaged precision, recall, and F1-score. The table 3 below illustrates relative performance of individual models on the multilingual classification task.

**Table 2: Comparative Evaluation of Model Performance on Multilingual Fake News Detection**

| Model | Accuracy | Macro Precision | Macro Recall | Macro F1-score |
|---|---|---|---|---|
| Logistic Regression | 0.5618 | 0.08 | 0.14 | 0.10 |
| Naive Bayes | 0.6971 | 0.36 | 0.29 | 0.27 |
| Random Forest | 0.7737 | 0.36 | 0.34 | 0.33 |
| GRU | 0.9175 | 0.82 | 0.60 | 0.65 |
| CNN | 0.8999 | 0.38 | 0.42 | 0.40 |
| **DeBERTa + LSTM (Hybrid)** | **0.9700** | **0.84** | **0.85** | **0.84** |

On basis of comparative study above, it clearly show classical ML models are computationally cheap but they are not suitable when it comes to tasks that require multilingual and low-resource environments. The deep learning models like GRU and CNN provided a great leap in performance, especially on the moderately resourced languages. Nevertheless, they were not so successful in generalizing to all classes of languages. The hybrid DeBERTa + LSTM was the most promising as it not only demonstrated the highest overall accuracy but also demonstrated a robust macro-level performance in terms of all metrics.

## 5.5 Discussion

The findings of the present study show clearly that by combining transformer-based contextual embeddings with sequential learning models, fake news detection in a multilingual environment, especially on low-resource languages, can be considerably enhanced. The hybrid DeBERTa-LSTM model takes advantage of the semantic richness and language flexibility of the DeBERTa transformer, and LSTM part enhance capacity of model for learning temporal and syntactic relationships in the text. This synergy responds to some of the fundamental issues that have been raised in the literature, including the poor performance of standalone models when used on a variety of languages and resource-limited settings.

Compared to the approaches that have been reviewed above, the hybrid model delivers significant benefits. Despite their computational efficiency and interpretability, such traditional machine learning techniques do not have deep semantic understanding, and are overly reliant on language-specific features to be effectively extended to multiple languages or domains. The same can be said about deep learning-based models such as LSTMs on their own, or transformer-based models such as BERT, which are more adept at capturing language nuances, but suffer either contextual depth or computation resource requirements when applied to multilingual and low-resource scenarios. The hybrid model introduced here is a fusion of the two worlds: the efficient and deep contextual embeddings of DeBERTa enables the model to capture complex semantic relations, whereas the LSTM layer is well suited to capturing sequential patterns and dependencies, which are essential to identify stylistic and narrative indicators of misinformation. The combination is not only better than its constituent parts, but also solves problems reported by earlier researchers concerning the inferiority of purely transformer-based models in resource-constrained or severely imbalanced data.

Moreover, the success of the model is consistent with the growing literature proposing the use of hybrid architectures on NLP tasks. As another example, the better results of transformer-RNN hybrids in sentiment classification reported by (Assiri et al. 2024) are reflected in results of this study in detection of fake news in multiple languages. The computational efficiency aspect of DeBERTa can also be used as an advantage since DeBERTa uses the disentangled attention mechanism to optimize the encoding process, thus providing a more computationally efficient alternative to models such as BERT or RoBERTa, as it has been addressed in the literature on computational efficiency. This is essential, particularly in low-resource language settings where training and deployment of large-scale models is infeasible in most cases.

Coupled with these strengths, there are endless challenges that are also brought out in the study. In particular, the hybrid model was less accurate at languages where data were much more scarce, such as Hindi, and this demonstrates the ongoing problem of dataset scarcity and class imbalance reported by (Fagundes et al. 2024) and others. These shortcomings reflect the necessity of future research to include more tactics like data augmentation, domain-specific fine-tuning, and transfer learning across linguistically comparable languages to enhance performance even more in the genuine low-resource environments. Besides, the model has lower computational complexity than the use of full transformer stacks, but it is not easy to deploy in the real-time settings with limited resources, which underlines the importance of further research into model compression and optimization methods.

The results of this investigation are not limited to technical measures of performance. This work fills a vital gap in the literature and provides a scalable solution to the problem of misinformation across languages by showing that a hybrid architecture can be effective at detecting fake news and at the same time be computationally efficient. Its cross-lingual generalisation capabilities in both low- and high-resource languages hold potential in terms of practical applications in the social media monitoring, news verification platforms, and other real-world systems where multilingual misinformation can go viral. Also, addressing underrepresented languages, the study fosters digital inclusion and fair creation of NLP tools, which can help defend vulnerable linguistic groups against the dangers of misinformation.

Finally, the value of such a research is in the holistic nature of it as it balances accuracy, efficiency, and linguistic variety. It gives empirical data on the necessity of combining semantic and sequential learning elements as a solution to the natural drawbacks of the past models. This is in line with the wider demand in NLP research of models which are not just powerful but also flexible and available in a variety of, real-world settings. The results support the justification of the topic in terms of research: to make the global community smarter and more resistant, it is necessary to solve detection of fake news in many languages problem in low-resource languages. The hybrid DeBERTa-LSTM model introduced in this paper is a step in the right direction that can be used both theoretically and practically in the fight against the common phenomenon of misinformation in the world.

### 5.6 Summary

This chapter gave an in-depth discussion of experimental outcomes of various models applied in multilingual fake news detection. The classical machine learning models were of limited use particularly where there was a low resource language. Deep learning models increased performance but had weaknesses that were dependent on the classes. The most effective architecture proposed, hybrid architecture of DeBERTa and LSTM classifier, has shown the highest accuracy and balanced performance in all classes of languages. These results confirm that the hybrid model can be a scalable, efficient, and precise method of multilingual fake news identification in resource-rich and low-resource environments.

# 6. Conclusion and Future Work

### 6.1 Conclusion

In the context of the fast rate of information spread and the growing digital globalization, the problem of fake news identification in all languages is both timely and complicated. This study was particularly addressing the issue of multilingual fake news detection and its focus was on achieving better results on low-resource languages, which are typically underrepresented in mainstream NLP applications, e.g., Swahili, Vietnamese, Indonesian. To remedy this, the study was based on a proposed and tested hybrid model that combines a potent pre-trained transformer called DeBERTa and a Bidirectional LSTM classifier in order to learn deep semantic context and sequential language patterns.

The study started with the analysis of a multilingual corpus of 7 different languages, i.e., English, Japanese, Turkish, Swahili, Indonesian, Vietnamese, and Hindi. The dataset was prepared after undertaking stringent preprocessing steps such as class balancing and language-sensitive text normalization to make it possible to evaluate robustly across languages with different data coverage. Baseline experiments with classical ML methods, including Logistic Regression, Naive Bayes, and Random Forest, showed the inability of the traditional methods to detect subtle linguistic characteristics, in particular, in low-resource languages. Such deep learning techniques as GRU and CNN proved to be more successful but also ineffective in generalizing all languages, which is why more sophisticated hybrid solutions are required.

The main contribution of this research, the DeBERTa-LSTM hybrid model, was a great breakthrough in performance. The model was able to perform significantly better than all baseline methods with an overall accuracy of 97% by using the multilingual contextual embeddings of DeBERTa and the sequence modeling capacity of Bidirectional LSTM. It specifically performed well in macro-averaged precision, recall & F1-score scores, and capable of identifying fake news better even in languages with fewer training data. Importantly, the model was computationally efficient in that the DeBERTa layers were kept frozen during training and thus it could be deployed to the real world where there are limited computational resources.

As a direct answer to the research questions: First, the combination of DeBERTa and LSTM does enhance the accuracy of fake news detection on multilingual datasets, particularly on low-resource languages, which is evidence of the effectiveness of deep contextual learning and sequential pattern learning. Second, hybrid model is well balanced on basis of computational efficiency & performance due to using the lightweight nature of frozen transformer and flexibility of LSTM, and has high accuracy yet little resource overheads compared with traditional models. Third, compared to standalone classical models, such as Random Forest and Decision Trees, the proposed hybrid system has a much shorter training time and a higher level of computational efficiency but significantly better predictive performance, which is why it is a more realistic option to use in multilingual multitask fake news detection tasks.

Comprehensively, the study has proven that hybrid transformer-RNN models offer an effective and viable approach to the low-resource multilingual fake news detection problem. The study is a valuable contribution to the research area of multilingual natural language understanding and media integrity technologies because it focuses on the linguistic diversity of the vocabulary as well as the computational efficiency. Besides contributing to academic knowledge, the findings promise to be applicable in the real-world setting in an attempt to curb misinformation in a global context.

## 6.2 Future Work

Although the hybrid model suggested in the present study demonstrated good results, there are a number of directions, which can be pursued to enhance the generalizability, fairness, and practical applicability of multilingual fake news detection systems.

The main constraint that was faced in this study was the access to limited data on some of the languages, particularly Hindi, Swahili and Vietnamese. This might be the subject of future research by implementing data augmentation procedures, including back-translation, paraphrasing, or cross-lingual transfer learning with other linguistically related languages of high resource.

Also, the DeBERTa + LSTM model showed a high level of accuracy, but this model is mostly frozen and stationary in this application. Future versions can also look into fine-tuning the transformer layers on domain-specific data or tuning with LoRA (Low-Rank Adaptation) or other parameter-efficient tuning methods to further improve contextual sensitivity without adding much to training costs.

The other promising direction is to expand the classification framework to be able to work with multimodal fake news detection where the text is not the only source of information that is used to detect fake news but also images, videos, or metadata that can be used to enhance the robustness of detection. As more misinformation is used in social media that includes multimedia content, there is an even greater need to adopt such holistic methods.

In addition, when this model is implemented as a real-time application or when it is integrated into content moderation platforms, the aspect of latency, scalability, and fairness will have to be taken into consideration. This involves the analysis of possible language and dialect detection bias as well as the model decision transparency.
Finally, to increase the trust and uptake of such systems by journalists, policymakers, and social media platforms, the scope of analysis may be extended beyond binary or categorical classification to detect fake news in an explainable fashion, i.e., by providing human-readable explanations of model predictions.

To summarize, though the current thesis provides a reliable and scalable system of multilingual fake news detection, it has the potential to be improved further and implemented responsibly within the context of media literacy and digital governance systems. Such technologies can play an important role in enhancing an informed, truthful, and inclusive digital ecosystem with interdisciplinary collaboration and constant innovation.

# References


Alghamdi, J., Lin, Y., & Luo, S. (2024). Fake news detection in low-resource languages: A novel hybrid summarization approach. Knowledge-Based Systems, 296, 111884.

Amanzadeh Oghaz, T. (2022). Low-Resource Machine Learning Techniques for the Analysis of Online Social Media Textual Data.

Raquiba, S. U. L. T. A. N. A. (2024). Identification of Fake News in High and Low Resource Languages: A Deep Learning Approach.

Assiri, A., Gumaei, A., Mehmood, F., Abbas, T., & Ullah, S. (2024). DeBERTa-GRU: Sentiment Analysis for Large Language Model. *Computers, Materials & Continua*, *79*(3).

Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., ... & Stoyanov, V. (2019). Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.

Zhang, D., Shan, G., Lee, M., Zhou, L., & Fu, Z. (2025). MT-GPD: A Multimodal Deep Transfer Learning Model Enhanced by Auxiliary Mechanisms for Cross-Domain Online Fake News Detection. *Production and Operations Management*, 10591478251319686.

Kumari, R., Ashok, N., Ghosal, T., & Ekbal, A. (2021, July). A multitask learning approach for fake news detection: Novelty, emotion, and sentiment lend a helping hand. In *2021 International Joint Conference on Neural Networks (IJCNN)* (pp. 1-8). IEEE.

Han, S. (2022). Cross-lingual Transfer Learning for Fake News Detector in a Low-Resource Language. *arXiv preprint arXiv:2208.12482*.

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation, 9*(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

Sanh, V., Debut, L., Chaumond, J., & Wolf, T. (2019). DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.

Gurgurov, D., Bäumel, T., & Anikina, T. (2024). Multilingual large language models and curse of multilinguality. *arXiv preprint arXiv:2406.10602*.

Fagundes, M. J. G., Roman, N. T., & Digiampietri, L. A. (2024). The use of syntactic information in fake news detection: A systematic review. *SBC Reviews on Computer Science*, *4*(1), 1-10.

Airlangga, G. (2024). Advancing fake news detection: a comparative study of RNN, LSTM, and Bidirectional LSTM Architectures. *Jurnal Teknik Informatika CIT Medicom*, *16*(1), 13-23.

Khan, M. A., Huang, Y., Feng, J., Prasad, B. K., Ali, Z., Ullah, I., & Kefalas, P. (2023). A multi-attention approach using BERT and stacked bidirectional LSTM for improved dialogue state tracking. *Applied Sciences*, *13*(3), 1775.

Kalusivalingam, A. K., Sharma, A., Patel, N., & Singh, V. (2021). Leveraging bert and lstm for enhanced natural language processing in clinical data analysis. *International Journal of AI and ML*, *2*(3).