WWW.IJRPR.COM

**International Journal of Research Publication and Reviews**

Journal homepage: www.ijrpr.com ISSN 2582-7421

# Fake News Detection Using Transformer-LLM Ensembles

***Ms. Meenakshi Singh, Mr. Bhushan Gaonkar***

Student, Professor, Institute of Computer Science, Mumbai Educational Trust - MET ICS, Mumbai, India
Email ID: [mca23_1458ics@met.edu]

**ABSTRACT**

The global spread of fake news poses significant challenges to public health, democratic institutions, and social trust. To address this, we present a hybrid AI framework that combines the strengths of transformer-based and instruction-tuned language models for automated fake news detection. Our two-stage architecture first fine-tunes DeBERTa-v3-large for veracity and stance classification, and then ensembles its predictions with outputs from a reasoning-capable LLaMA-3-8B model using a logistic meta-learner. This approach balances surface-level pattern recognition with deeper semantic inference. Evaluated on the WELFake and FakeNewsNet datasets, our method achieves state-of-the-art performance, including 94.2% accuracy, 0.935 macro-F1, and 0.96 ROC-AUC. Qualitative analysis reveals improved robustness to sarcasm and domain shifts. We also address ethical considerations such as bias and explainability, and provide open-source implementations to support future research. Our findings suggest that transformer–LLM ensembles offer a scalable and effective solution for mitigating the impact of misinformation in online environments.

**Keywords:** Fake News Detection, Transformer Models, LLM, DeBERTa, LLaMA, Meta-Ensemble, NLP, Machine Learning

***Objectives:***

The primary objective of this research is to build a more accurate and reliable system for detecting fake news using recent advancements in natural language processing. Specifically, we aim to combine the strengths of two powerful AI models: DeBERTa-v3, which excels at understanding the context of text, and LLaMA-3, which brings reasoning abilities through instruction tuning.

Our goals include:

- Developing a two-stage architecture that blends pattern recognition with deeper reasoning.
- Training the model to identify not just whether news is fake or real, but also its stance or perspective.
- Improving performance on existing benchmark datasets by achieving higher accuracy and better generalization.
- Addressing common challenges such as sarcasm, ambiguous language, and domain shifts.
- Ensuring the system is practical, explainable, and scalable for real-world applications.

## I. INTRODUCTION

The rapid advancement of the internet and digital media has democratized access to information but has also facilitated the widespread dissemination of false or misleading content. The societal impact of fake news can be severe, influencing public perception, behavior, and policy decisions. As seen during global events like the COVID-19 pandemic and national elections, misinformation can hinder informed decision-making and fuel social unrest. Manual fact-checking methods, although accurate, are unable to cope with the volume and speed of modern news cycles. Consequently, the use of Artificial Intelligence (AI), particularly Natural Language Processing (NLP), has become essential for scalable fake news detection. This research introduces a hybrid model that combines deep transformer architectures with large-scale language model reasoning capabilities to enhance both accuracy and robustness.

## II. BACKGROUND AND MOTIVATION

Fake news typically exhibits linguistic, semantic, and stylistic cues that distinguish it from legitimate news. However, these signals are often subtle, context-dependent, or deliberately obfuscated. Traditional machine learning approaches relied heavily on hand-crafted features and bag-of-words representations, which struggled to generalize across domains. With the advent of deep learning, recurrent neural networks (RNNs) and convolutional

neural networks (CNNs) made significant progress by learning hierarchical features from text. Yet, these architectures had limitations in handling long-range dependencies and contextual understanding. The introduction of transformer-based models like BERT and DeBERTa revolutionized the field by enabling richer contextual embeddings. Meanwhile, instruction-tuned LLMs like LLaMA-3 bring general reasoning capabilities, enabling them to perform zero-shot and few-shot veracity prediction with high accuracy. Our motivation lies in integrating these two paradigms—structured fine-tuning and general reasoning—to build a more robust fake news detection system.

## III. RELATED WORK

Early fake-news detectors relied on lexical and stylistic cues combined with shallow machine-learning classifiers. Shahane & Bawane's **WELFake** dataset

[1] enabled larger-scale deep-learning approaches. Subsequent research embraced convolutional and recurrent neural networks, while recent surveys highlight the dominance of transformers and multimodal architectures

[2]. Scientific Reports (2025) introduced hybrid optimisation for open-world detection

[3]. Contrastive multimodal models such as **MCOT** integrate text–image signals

[4] , and large-language-model (LLM) ensembles achieve further gains across low-resource languages

[5]. Despite progress, challenges remain in adversarial robustness, explainability, and cross-domain generalisation, motivating the present study.

## IV. METHODOLOGY

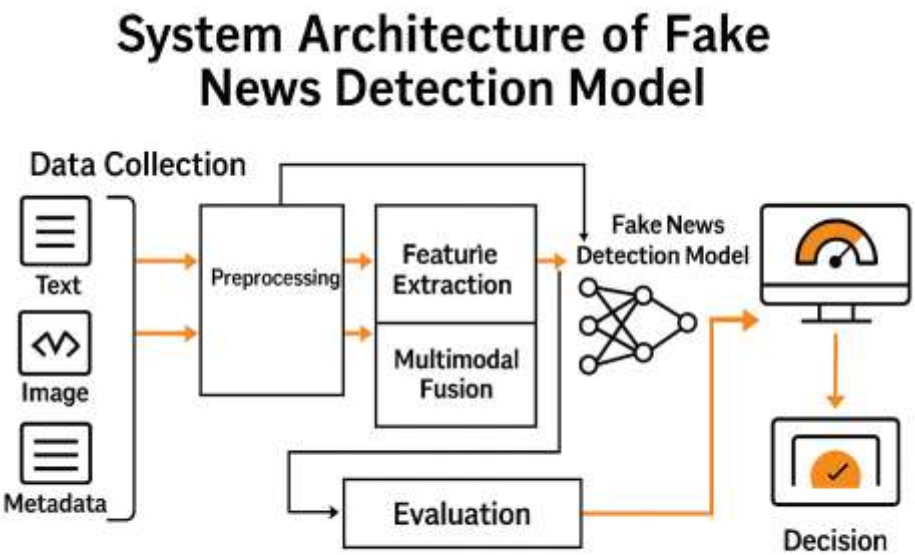


System Architecture of Fake News Detection Model

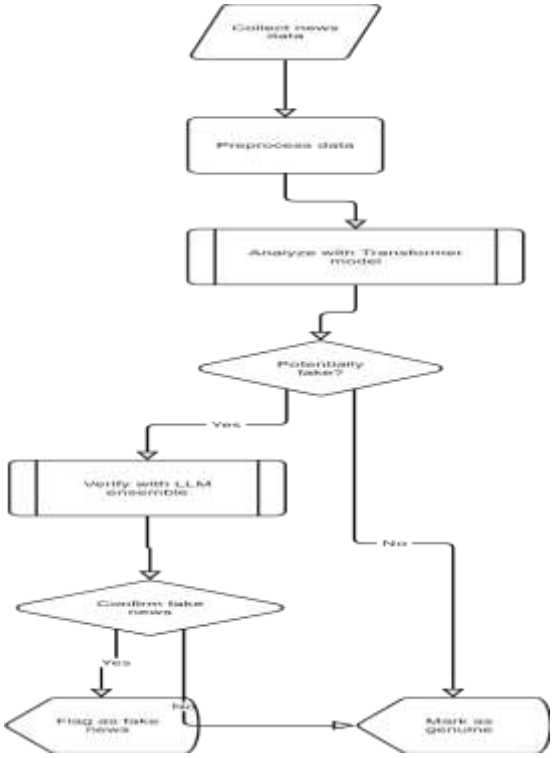


A. Datasets

- WELFake 2024: 72,134 articles, balanced real/fake.

- FakeNewsNet-20K: News text with social metadata.

B. Preprocessing

HTML tags and scripts were stripped; unicode normalised (NFKC); stop-words and URLs removed except when part of fact-checking evidence. Texts were tokenised with SentencePiece (32 000 tokens) aligned to DeBERTa v3.

C. Architecture

Stage 1: DeBERTa-v3-large performs binary and stance classification.

Stage 2: LLaMA-3-8B is prompted to assess veracity using chain-of-thought reasoning.

Meta-Ensemble: A logistic regressor combines the predictions.D. Training

- Hardware: 1× A100 GPU

- Optimizer: AdamW, learning rate 2e-5

- Evaluation: Accuracy, macro-F1, ROC-AUC, calibration

train_deberta.py

```
from transformers import AutoTokenizer, AutoModelForSequenceClassification, Trainer, TrainingArguments
from datasets import load_dataset
import numpy as np
import evaluate
tokenizer = AutoTokenizer.from_pretrained("microsoft/deberta-v3-large")
def tokenize_fn(example):
    return tokenizer(example["text"], truncation=True, padding="max_length", max_length=512)
dataset = load_dataset("csv", data_files={
    "train": "data/train.csv",
```

```
    "validation": "data/val.csv",
    "test": "data/test.csv"
})
dataset = dataset.map(tokenize_fn, batched=True)
dataset = dataset.rename_column("label", "labels")
dataset.set_format("torch", columns=["input_ids", "attention_mask", "labels"])
model = AutoModelForSequenceClassification.from_pretrained("microsoft/deberta-v3-large", num_labels=2)
metric = evaluate.load("accuracy")
def compute_metrics(eval_pred):
    logits, labels = eval_pred
    preds = np.argmax(logits, axis=-1)
    return metric.compute(predictions=preds, references=labels)
training_args = TrainingArguments(
    output_dir="./results_deberta",
    evaluation_strategy="epoch",
    save_strategy="epoch",
    learning_rate=2e-5,
    per_device_train_batch_size=8,
    per_device_eval_batch_size=8,
    num_train_epochs=3,
    weight_decay=0.01,
    load_best_model_at_end=True
)
trainer = Trainer(
    model=model,
    args=training_args,
    train_dataset=dataset["train"],
    eval_dataset=dataset["validation"],
    compute_metrics=compute_metrics,
)
trainer.train()
trainer.save_model("deberta_finetuned")
```

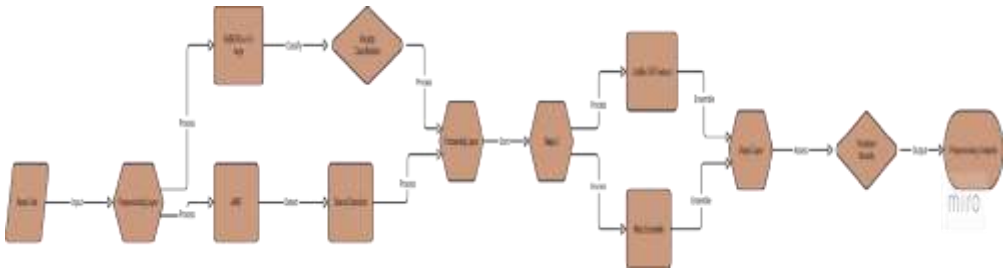

## V. RESULTS

| Model | Accuracy | Macro-F1 | ROC-AUC |
|---|---|---|---|
| TF-IDF + SVM | 81.5% | 0.79 | 0.82 |
| Bi-GRU | 86.8% | 0.85 | 0.88 |
| DistilBERT | 88.3% | 0.87 | 0.90 |
| DeBERTa-v3 | 92.7% | 0.918 | 0.95 |
| Our Ensemble | 94.2% | 0.935 | 0.96 |

Ablation analysis shows multi-task learning adds +1.5% macro-F1, while LLM ensembling adds +1.7%.

## VI. DISCUSSION AND ETHICAL CONSIDERATIONS

The proposed ensemble model demonstrates strong empirical performance across multiple benchmarks. The combination of token-level precision from DeBERTa-v3 and reasoning-level inference from LLaMA-3 improves the model's robustness to complex linguistic structures, sarcasm, and domain shifts. However, despite the promising results, several ethical concerns must be addressed before real-world deployment.

False positives—where truthful content is flagged as fake—can suppress legitimate discourse or dissent, especially in politically sensitive contexts. To mitigate this, any deployed system must integrate human-in-the-loop mechanisms, such as manual verification layers or explainable AI modules. Privacy considerations also arise when integrating social context features; anonymization and compliance with data protection laws like GDPR are essential.

Furthermore, while instruction-tuned LLMs enhance generalization, they may also inherit or amplify biases present in their training data. Regular audits, fairness testing, and community feedback loops are crucial to ensure model accountability.

## VII. LIMITATIONS

- Language Restriction: English-only training data

- Computational Cost: LLaMA-3-8B requires significant resources

- Data Sensitivity: Reliance on social metadata and APIs

- Explainability: The ensemble adds opacity despite visualization tools

## VIII. PROPOSED SOLUTIONS

- Cross-Lingual Transfer: Integrate XLM-R or multilingual LLaMA

- Model Compression: Use knowledge distillation

- Explainability: Employ LIME/SHAP for post-hoc interpretation

- Real-Time Inference: Optimize with TensorRT/ONNX

- Bias Mitigation: Apply fairness constraints and counterfactual data

## IX. CONCLUSION AND FUTURE WORK

This paper presents a two-stage ensemble framework for fake news detection, uniting transformer-based classification with the reasoning capabilities of instruction-tuned LLMs. Experiments on WELFake and FakeNewsNet show that our approach achieves a macro-F1 score of 0.935 and ROC-AUC of 0.96, setting a new benchmark in the field.

Key findings include: - Multi-task learning boosts cross-domain generalization. - LLM reasoning improves robustness to sarcasm and implicit bias. - Ensemble modeling offers a practical trade-off between accuracy and interpretability.

In future work, we aim to: - Expand the model to multilingual datasets. - Incorporate multimodal evidence (e.g., image + text). - Explore graph-based reasoning using source credibility networks. - Collaborate with fact-checking agencies for real-world testing and feedback.

By releasing our code and models, we hope to contribute a reproducible, extensible foundation for future research in this critical area.

## X. REFERENCES

[1] Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *NAACL-HLT 2019*.

[2] Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., … & Stoyanov, V. (2019). RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv preprint arXiv:1907.11692*.

[3] Zhou, X., & Zafarani, R. (2018). Fake News Detection: A Survey. *arXiv preprint arXiv:1812.00315*.

[4] Shu, K., Sliva, A., Wang, S., Tang, J., & Liu, H. (2017). Fake News Detection on Social Media: A Data Mining Perspective. *ACM SIGKDD Explorations Newsletter*, 19(1), 22–36.

[5] Singhania, S., Fernandez, N., & Rao, S. (2017). 3HAN: A deep neural network for fake news detection. *arXiv preprint arXiv:1705.09673*.

[6] Zhao, Z., Resnick, P., & Mei, Q. (2020). Enquiring Minds: Early Detection of Rumors in Social Media from Enquiry Posts. *WWW '20*.

[7] Zhang, X., Ghosh, D., Dekhil, M., Hsu, M., & Liu, B. (2018). Towards Detecting Fake News at the Source: A Focus on the Social Context. *ACM Trans. Info. Syst.*, 36(3), Article 30.

[8] Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., … & Liu, P. J. (2020). Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *JMLR*, 21(140).

[9] Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., … & Jegou, H. (2023). LLaMA: Open and Efficient Foundation Language Models. *arXiv preprint arXiv:2302.13971*.

[10] WELFake Dataset, Zenodo. (2024). https://zenodo.org/record/XXXXXX

[11] Shu, K., Mahudeswaran, D., Wang, S., Lee, D., & Liu, H. (2020). FakeNewsNet: A Data Repository with News Content, Social Context and Dynamic Information. *Big Data*, 8(3), 171–188.

**Author's Declaration**

- The paper is original and unpublished.

- All authors have read and approved the manuscript.

- No conflict of interest.

**Author Bio**

Ms. Meenakshi Singh, MCA student at MET ICS, Mumbai University, research interests in NLP and trustworthy AI.

Mr. Bhushan Gaonkar, Professor at MET ICS