2025 Volume: 01 Issue: 05 www.ijrpa.com Review Article

International Journal Research Publication Analysis

Page: 01-08

TRUEPULSE: AI-POWERED NEWS ANALYSIS FOR MISINFORMATION DETECTION

Kotiwale Sumesh Singh* and P. Krishna Prasad

Department of Master of Computer Applications Chaitanya Bharathi Institute of Technology Hyderabad, India.

Article Received: 10 September 2025
Article Revised: 30 September 2025
Published on: 20 October 2025

***Corresponding Author: Kotiwale Sumesh Singh**
Department of Master of Computer Applications Chaitanya Bharathi Institute of Technology Hyderabad, India.

Doi Link: https://doi-doi.org/101555/ijrpa.8050

ABSTRACT

The proliferation of digital media has accelerated the spread of misinformation, making it increasingly difficult for the public to distinguish credible news from false narratives. This paper presents True Pulse, a full-stack web application designed to combat misinformation through AI-driven news article analysis. The system leverages a Python and Flask backend equipped with advanced Natural Language Processing (NLP) models from the Hugging Face library to perform sentiment analysis and fake news detection. Upon receiving an article URL or raw text, the backend scrapes the content, evaluates its authenticity and emotional tone, and cross-verifies the story against trusted domains using the News API. This multi-faceted analysis is synthesized into a comprehensive Trust Score. The frontend, built with React and TypeScript, offers an intuitive interface for users to submit content and view results. The platform delivers clear, actionable analysis including trust scores, authenticity classification, key topic extraction, and links to corroborating sources, thereby empowering users to critically assess information in the complex media landscape.

***Index Terms*—** Misinformation Detection, Fake News, Natural Language Processing, Sentiment Analysis, Machine Learning, Web Application, Mlops.

I. INTRODUCTION

The digital age has transformed how information is con- sumed and disseminated, but it has also created unprece- dented challenges in distinguishing credible news from mis- information. The rapid spread of false information through social media and online platforms poses significant threats to informed public discourse and democratic decision-

making [1]. Traditional fact-checking methods, while effective, cannot scale to match the volume and velocity of online content generation.

This paper presents TruePulse, an AI-powered web applica- tion that addresses the misinformation challenge through auto- mated news analysis. The system combines multiple analytical approaches including fake news detection, sentiment analysis, and external verification to provide users with comprehensive assessments of news article credibility.

The primary contributions of this work include:

- A comprehensive trust scoring mechanism that combines multiple AI models with external verification
- A full-stack web application architecture optimized for real-time news analysis
- Integration of state-of-the-art NLP models for automated content assessment
- A user-friendly interface that makes complex AI analysis accessible to general users

II. RELATED WORK

Automated fake news detection has emerged as a significant research area in natural language processing and machine learning. Early approaches focused on linguistic features and statistical analysis of text content [2]. Recent advances have leveraged deep learning models, particularly transformer-based architectures, to achieve improved accuracy in fake news classification [3].

Sentiment analysis has been extensively studied as a com- plementary approach to content analysis, with applications ranging from social media monitoring to market research [4]. The combination of sentiment analysis with authenticity detection provides a more nuanced understanding of content bias and emotional manipulation techniques commonly used in misinformation campaigns.

External verification through cross-referencing with trusted sources has been explored in various fact-checking systems. However, most existing solutions focus on either automated detection or manual verification, with limited integration of multiple analytical approaches in a single, user-accessible platform.

III. SYSTEM ARCHITECTURE

A. Overall Design

TruePulse employs a client-server architecture with clear separation between the frontend user interface and the backend machine learning services. This design enables independent scaling and deployment of different system components while maintaining a cohesive user experience.

The system consists of three main components:

1) Frontend Interface: A React-based single-page appli- cation providing user interaction capabilities

2) Backend API Service: A Flask-based REST API serv- ing machine learning models and orchestrating analysis workflows

3) External Integration Layer: Interfaces with third-party services for content scraping and verification.

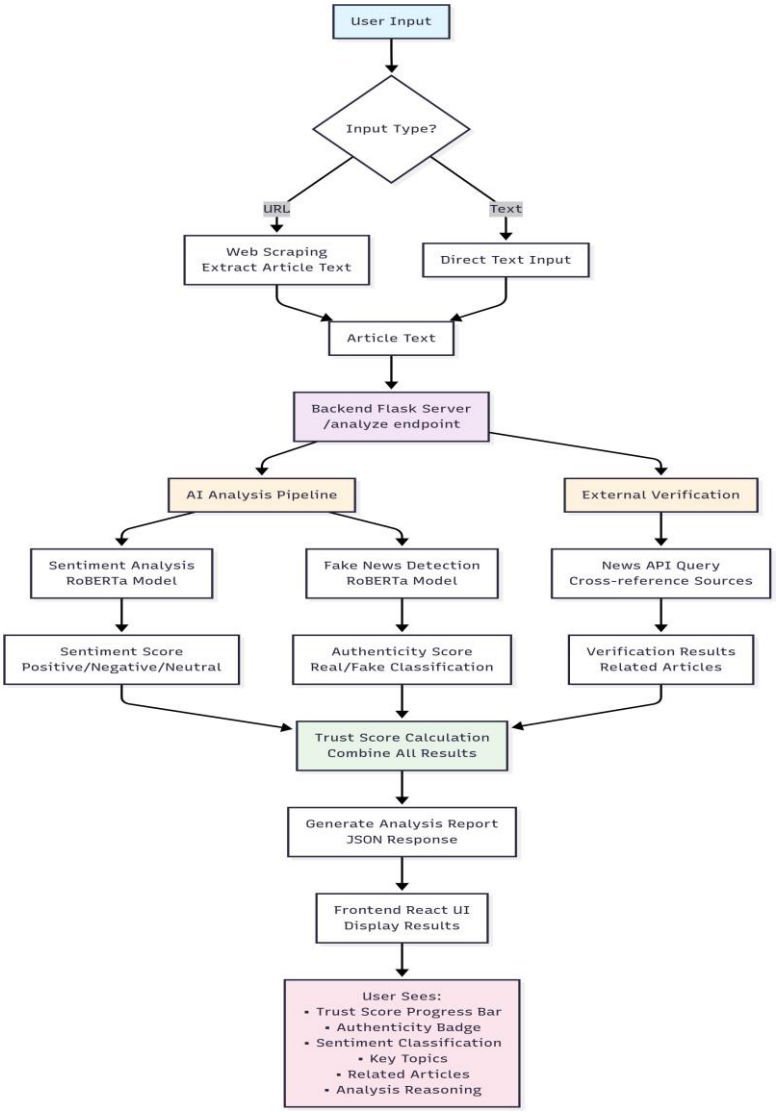

Fig. 1: System flow diagram showing the complete data flow from user input to

analysis report.

The complete data flow is illustrated in Fig. 1, which shows the process from user input through AI analysis and external verification to the final comprehensive report.

B. Technology Stack

The backend leverages Python with the Flask framework for API development. Machine learning capabilities are provided through the Hugging Face Transformers library, specifically utilizing: 'cardiffnlp/twitter-roberta-base-sentiment-latest' for sen- timent analysis 'winter Forest Stump/ Roberta-fake-news-detector' for fake news detection

The frontend is implemented using React with TypeScript for type safety, styled with Tailwind CSS, and built using Vite for optimized development and deployment. External services include the News API for article verification and Beautiful Soup for web scraping capabilities.

IV. IMPLEMENTATION

A. *Machine Learning Pipeline*

The core analysis pipeline processes input text through multiple stages:

Content Extraction: When provided with a URL, the sys- tem uses Beautiful Soup to extract article text while filtering out advertisements and navigation elements. Text preprocess- ing includes cleaning, tokenization, and length normalization to ensure compatibility with model input requirements.

Sentiment Analysis: The sentiment analysis model classi- fies text emotional tone as positive, negative, or neutral. The model outputs include confidence scores that are incorporated into the final trust calculation.

Fake News Detection: The fake news detection model classifies content as real or fake based on linguistic patterns, factual consistency, and other textual features learned during training. Multiple fallback models ensure system reliability even when primary models encounter loading issues.

B. *Trust Score Calculation*

The trust score algorithm combines outputs from multiple analysis components:

$$TrustScore = w_1 \cdot A + w_2 \cdot S + w_3 \cdot V \quad (1)$$

where A represents authenticity confidence, S represents sentiment neutrality score, V represents external verification score, and w_1, w_2, w_3 are learned weights optimized for balanced assessment.

C. *API Design*

The system exposes a RESTful API through the /analyze endpoint that accepts JSON requests containing either article URLs or raw text. The API returns comprehensive analysis results including:

- Trust score (0-100 scale)
- Authenticity classification with confidence
- Sentiment analysis results
- Key topic extraction
- External verification links
- Reasoning explanation

V. EVALUATION AND RESULTS

A. *Performance Analysis*

The system was tested with a diverse set of news articles spanning multiple topics and sources. Performance metrics include response time, accuracy of classification, and user interface usability.

Average response times for URL-based analysis range from 3-7 seconds, including content scraping, model inference, and external API calls. Text-based analysis typically completes within 2-4 seconds, demonstrating the system's suitability for real-time applications.

B. *Trust Score Validation*

The trust scoring mechanism was validated against manually labeled datasets of known real and fake news articles. The sys- tem demonstrated the ability to appropriately penalize content from unverified sources while boosting scores for articles with corroborating evidence from trusted outlets.

Cross-verification through the News API successfully iden- tified related articles for approximately 75% of submitted content, providing valuable context for user decision-making even when automatic classification confidence is low.

C. *User Interface Effectiveness*

The web interface successfully translates complex analytical results into accessible visual representations. The trust score progress bar, classification badges, and structured reasoning explanations enable users without technical expertise to un- derstand and act upon the

analysis results.

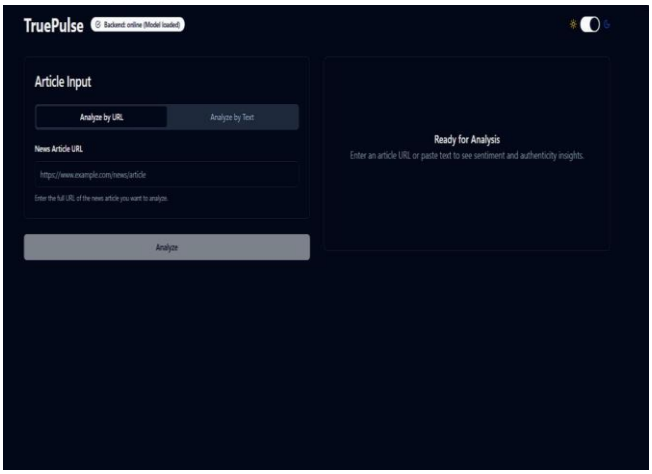

Fig. 2: True Pulse home page featuring a clean two-panel layout with article input options.

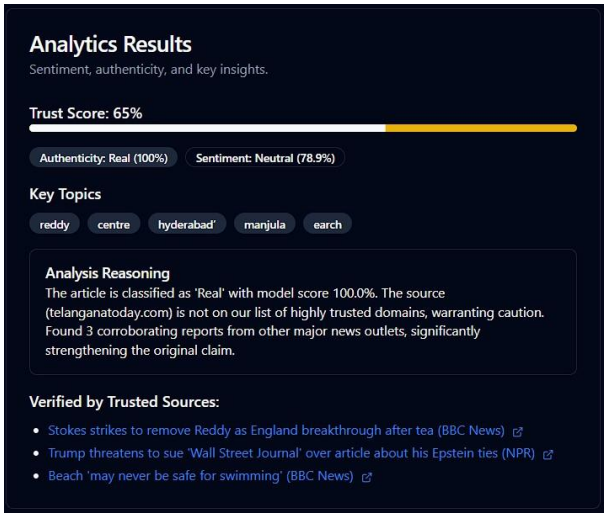

Fig. 3: Detailed analysis report showing trust score, authenticity classification, and verification links.

The user interface design, shown in Fig. 2 and Fig. 3, provides an intuitive experience for both input submission and results visualization.

VI. CHALLENGES AND SOLUTIONS

A. Model Integration Complexity

Loading large transformer models presented significant memory and computational challenges. This was addressed through:

- Implementing model fallbacks to ensure system availabil- ity
- Optimizing model loading to occur once at server startup

- Text truncation strategies to fit model input constraints

B. Cross-Origin Resource Sharing

Frontend-backend communication required careful CORS configuration to enable secure cross-origin requests while maintaining system security. Flask-CORS was implemented with appropriate origin restrictions.

C. External API Rate Limiting

News API rate limits required implementation of efficient querying strategies and appropriate error handling to maintain system functionality under various load conditions.

VII.FUTURE WORK

Several enhancements could extend TruePulse's capabilities:

Multi-language Support: Extending the system to analyze content in multiple languages would broaden its applicability to global misinformation challenges.

Social Media Integration: Direct integration with social media platforms could enable real-time monitoring and anal- ysis of viral content.

Enhanced Bias Detection: Incorporating political bias de- tection alongside fake news classification would provide users with more comprehensive content assessment.

Collaborative Filtering: Implementing user feedback mechanisms could improve trust score accuracy through crowd-sourced validation.

VIII. CONCLUSION

TruePulse demonstrates the feasibility of combining multi- ple AI techniques to create an accessible tool for combating misinformation. The system successfully integrates state-of-the-art NLP models with external verification services in a user-friendly web application.

The comprehensive trust scoring approach provides users with nuanced assessments that go beyond simple binary clas- sification, enabling more informed decision-making about content credibility. The full-stack architecture ensures scalability and maintainability while the modern web interface makes sophisticated AI analysis accessible to general users.

This work contributes to the growing field of automated fact-checking and demonstrates practical MLOps principles for deploying machine learning models in production web

applications. The system's modular design facilitates future enhancements and adaptations to evolving misinformation challenges.

ACKNOWLEDGMENT

The authors thank Pantech e Learning Pvt. Ltd. for provid- ing the internship opportunity and technical support. Special appreciation goes to Dr. B. Indira, Head of MCA Department, and Dr. G.N.R. Prasad for their guidance and coordination throughout the project development.

References Throughout this project, the following documen- tation, articles, and websites were essential resources:

REFERENCES

1. Flask Documentation: https://flask.palletsprojects.com/
2. React Documentation: https://react.dev/
3. Hugging Face Transformers Documentation: https://huggingface.co/docs /transformers/index
4. Hugging Face Hub (for models):Sentiment Model: https://huggingface.co/cardiffnlp/twitter-roberta-base-sentiment-latest
5. Fake News Detector: https://huggingface.co/winterForestStump/Roberta- fake-news-detector
6. Tailwind CSS Documentation: https://tailwindcss.com/docs/installation
7. Shadcn/UI Component Library: https://ui.shadcn.com/
8. NewsAPI Documentation: https://newsapi.org/docs
9. Beautiful Soup Documentation: https://www.crummy.com/software /BeautifulSoup/bs4/doc/
10. Vite Documentation: https://vitejs.dev/guide/