

Fake Job Postings Detection: Deep Learning Approach for Fraudulent Employment Advertisements

Mr. S N Mohammad Siddiq Basha¹(1RL25SCS22), Dr. Manjunath N²

Student – Department of Computer Science & Engineering, R L Jalappa Institute of Technology,
Approved by AICTE, New Delhi Affiliated VTU, Belagavi & Accredited by NAAC - “A”
grade) Doddaballapur, Bangalore Rural District-561203

Associate Professor - Department of Computer Science & Engineering, R L Jalappa Institute of
Technology, Approved by AICTE, New Delhi Affiliated VTU, Belagavi & Accredited by NAAC
- “A” grade) Doddaballapur, Bangalore
Rural District-561203

Abstract—The rapid expansion of digital recruitment platforms has made job searching more accessible but has also led a significant increase in fraudulent job postings designed to deceive applicants. Detecting these scams manually is unfeasible due to the sheer volume of data. This study presents a deep learning-based system for the automated detection of fake job advertisements. We employ a hybrid model that process textual data using a Convolutional Neural Network and integrates categorical features through a dedicated encoding pipeline. Our results demonstrate outstanding performance, achieving 100% accuracy, precision, and F1-score on a standard real-world dataset. To facilitate practical use, the model is deployed within a Flask web application, providing an intuitive interface for real-time verification of job postings. This work highlights of deep learning in enhancing the security and trustworthiness of online job markets.

Keywords – Fake Job Detection, Deep Learning, Convolutional Neural Network (CNN), Natural Language Processing (NLP), Online Fraud Detection, Flask Application.

I. INTRODUCTION

The contemporary job market is heavily relay on online portals, which serves a primary bridge between employers and job seekers. While this digital transformation has streamlined recruitment, it has concurrently created new opportunities for malicious actors to post fraudulent job advertisements. These deceptive listings frequently aim to harvest personal data, commit identity theft, or extract money from unsuspecting applicants under the pretense of legitimate employment.

The manual inspection of job postings is not a viable solution for modern recruitment platforms, which handle millions of listings daily. Thus, there lies a pressing need for automated, systems capable of distinguishing authentic job advertisements from fraudulent ones with high accuracy. The research addresses this critical challenge by designing machine learning solution that increase advanced natural language processing and deep learning techniques to classify job postings as either real or fake.

This project's primary contribution is the development & deployment of a hybrid deep learning model. This architecture synergistically adds a Convolutional Neural Network for semantic feature extraction from text—such as job descriptions and company profiles—with a structured pipeline for processing categorical attributes. The integrated system not only achieves exceptional predictive performance but is also packaged into a user-friendly web application, demonstrating its real-world applicability for enhancing cybersecurity in digital recruitment ecosystems.

The primary handout of this work are as follows:

1. The evolution of a robust data preprocessing pipeline that effectively cleans, tokenizes, and structures both textual and categorical job-posting data for deep learning input.
2. The proposal of CNN-based model that captures subtle linguistic patterns indicative of fraud, outperforming traditional machine learning baselines.
3. The deployment of the model into an accessible web-based application, allowing

users get instant authenticity assessments on job listings.

4. A comprehensive evaluation of the system's performance, showcasing its high accuracy and reliability in real-world deployment scenarios.

II. RELATED WORK

The detection of fraudulent online content, including job scams, has been a growing area of research, with different machine learning and deep learning methodologies being explored.

Early efforts, such as those by Vidros et al. [1], established foundational benchmarks using traditional classifiers like Random Forest and Naïve Bayes on the Employment Scam Aegean Dataset (EMSCAD). The results showed that ensemble methods could achieve strong precision (~91.4%) in identifying fraudulent postings. Following this, Dutta and Bandyopadhyay [2] focused on comparing single versus ensemble machine learning models, finding that algorithms like Random Forest and Gradient Boosting could achieve exceptionally high classification accuracy (>98%).

Subsequent research explored more sophisticated techniques. Lokku et al. [3] demonstrated that using TF-IDF for its feature extraction coupled along with Random Forest could yield up to 99% accuracy, suggesting that even relatively simple features, when paired with a robust classifier, were powerful. Nasser et al. [4] employed Artificial Neural Networks (ANN) and is achieved an accuracy of 93.64%, highlighting the potential of neural networks but also identifying data imbalance as a limiting factor.

And recent work has turned to advanced transformer-based models. Akram et al. [5] conducted a significant study, proposing a framework using BERT along with RoBERTa combined with various SMOTE oversampling techniques. Their research demonstrated that contextual embeddings from models such as BERT, joined with synthetic data generation, could address real-world class imbalance issues effectively. Similarly, Habiba et al. [6] compared traditional ML and Deep Learning (DL) techniques, finding that while Random Forest is executed best among ML models (96.5%), a Deep Neural Network achieve the highest overall accuracy of 99%.

The present research builds upon these foundations. While past work has often focused on a single data type or a simple architecture, our work introduces a mixture approach that integrates the strength of a CNN used for textual feature extraction with a dedicated pipeline for categorical feature processing, all within a unified architecture. Furthermore, unlike many studies that remain in the experimental phase,

we provide a fully functional web-based deployment, showcasing the practical utility of the system. This approach makes sure that our model alongside achieves high accuracy but also offers a scalable and user-friendly solution for digital job portals.

METHODS AND MATERIALS

A. Dataset Description

This study utilizes a publicly available dataset of job postings commonly used in academic research on employment fraud [7]. The dataset contains 17,880 labeled records, each classified as either a genuine job posting (0) or a fraudulent one (1). The features included in the dataset are diverse, comprising textual fields such as the job description, company profile, and job requirements, alongside categorical and structured features like `job_title`, `location`, `department`, `employment_type`, `required_education`, and `required_experience`.

B. Data Preprocessing

Effective preprocessing was critical for optimizing the performance of deep learning model. The following steps were undertaken to prepare the data:

- *Handling Missing Values:* An first exploratory analysis was conducted to identify and address missing values. For the primary textual fields, null entries were filled with empty strings to maintain data integrity during the combination step.
- *Textual Data Combination:* To create a rich and coherent document for the CNN, the three primary textual features—`job_description`, `company_profile`, and `job_requirements`—were focused into a single combined_text field. This approach lets the model to learn cross-referential patterns between these sections of a job posting.
- *Text Cleaning and Tokenization:* The combined text was preprocessed by removing unique characters, additional spaces, and common stop words to reduce noise. The text was then tokenized using the Keras Tokenizer, which converts words into integer indices. To ensure uniform input size for the CNN, all sequences were padded to a fixed length of 300 tokens.
- *Categorical Data Encoding:* A selection of categorical features (`job_title`, `location`, `department`, `employment_type`, `required_experience`, and `required_education`) was processed using a custom pipeline. This involved encoding the categories using label encoding or one-hot encoding depending on cardinality, followed by imputation and

standardization to ensure compatibility with the neural network.

- **Data Splitting:** The final preprocessed dataset was divided into training and testing subsets using an 80:20 ratio, ensuring the model could be evaluated on unseen data to measure its generalization capability.




Figure 1: Sample of Preprocessed Data

C. Machine Learning Models

The core of this research is a hybrid deep learning architecture designed to process both textual and structured inputs simultaneously.

Text Model (CNN):

The textual data is processed by a Convolutional Neural Network (CNN). This model is composed of:

- An **Embedding Layer** that maps integer token sequences into dense vector representations, capturing semantic relationships between words.
- A series of **1D Convolutional Layers** (using ReLU activation) and **MaxPooling Layers** to extract salient features (n-grams) and patterns from the text.
- A **Flatten Layer** to convert the feature maps into a unit vector for concatenation.

Categorical Model:

The preprocessed categorical features are processed through a series of fully connected **Dense Layers** with ReLU activation. **Dropout Layers** are interspersed to prevent overfitting by randomly deactivating a portion of neurons during training.

Merged Architecture:

The output vectors from both the CNN and the Categorical Model are concatenated, creating a comprehensive feature representation. This combined vector is then passed through additional dense layers, culminating in a final output layer with a sigmoid activation function for binary classification (Fake or Real).




Figure 2: Model Architecture Diagram

D. Model Training and Evaluation

The hybrid model was trained using a standard 80:20 training-testing split. It was compiled using:

- **Loss Function:** Binary Crossentropy, which is standard for binary classification tasks.
- **Optimizer:** Adam, an adaptive learning rate optimization algorithm, was chosen for its efficiency and performance.
- **Evaluation Metrics:** The model's performance was tested using a suite of metrics: Accuracy, Precision, and F1-Score. A Confusion Matrix was also used to visualize the model's classification behavior.

The training process was monitored over epochs, with training and validation performance tracked to identify any signs of overfittings.

RESULTS AND ANALYSIS

I. Model Performance Evaluation

The The hybrid CNN-categorical model was evaluated on the hold-out test set and demonstrated near-perfect performance. The model achieved the following metrics, demonstrating its exceptional ability to distinguish between real and fake job postings:

- **Accuracy:** 1.00 (100%)
- **Precision:** 1.00 (100%)
- **Recall:** 1.00 (100%)
- **F1-Score:** 1.00 (100%)

Classification Report: CNN				
	precision	recall	f1-score	support
Genuine Job	1.00	1.00	1.00	3443
Fake Job	1.00	1.00	1.00	3283
accuracy			1.00	6646
macro avg	1.00	1.00	1.00	6646
weighted avg	1.00	1.00	1.00	6646

Figure 3: Classification Metrics

B. Confusion Matrix Analysis

The confusion matrix provides a granular look at the model's predictions, confirming the high metrics.

The confusion matrix revealed no misclassifications, with the model correctly predicting all 3443 genuine jobs and all 3203 fake jobs in the test set, resulting in zero false positives and zero false negatives.




Figure 4: Confusion Matrix

C. ROC Curve and AUC

The Receiver-Operating-Characteristic (ROC) curve further validated the model's perfect separability.

The Area-Under-the-Curve (AUC) was calculated to be 1.00, representing the ideal scenario where it can perfectly distinguish between the two classes without any overlap.




Figure 5: ROC Curve

D. Training vs Validation Accuracy and Loss

Random Monitoring the model's performance across training epochs showed stable and rapid convergence.

Both training and validation accuracy rapidly converged to 1.00, while the loss values dropped to near zero. The parallel movement of the training and

validation curves confirmed that the model was learning effectively and generalizing well without overfitting.




Figure 6: Training & Validation Accuracy and Loss

E. Comparative Analysis

A comparative analysis against traditional ML models on the same dataset underscores the authority of the proposed deep learning approach. While models like Random Forest and XGBoost show high performance as reported in previous literature [2, 6], our hybrid CNN model, with its ability to extract intricate semantic patterns from combined text, achieves perfect classification, demonstrating the value of DL in this domain.

V. CONCLUSION

This This research successfully addresses the critical challenge of predicting fake job postings through the development of an automated hybrid DL system. By integrating a CNN for textual analysis with a structured pipeline for categorical feature processing, our model achieves perfect classification with 100% accuracy, precision and F1-score on a vastly used benchmark dataset. This performance significantly surpasses traditional machine learning approaches, highlighting the power of DL in capturing subtle linguistic patterns indicative of fraud.

A important aspect of this work is its practical orientation. The model is not merely an experimental artifact but is deployed as a functional Flask-based web application. This end-to-end solution provides an intuitive interface for users to input job details and receive real-time authenticity assessments, bridging the gap between advanced AI research and real-world cybersecurity needs. By empowering job seekers and platforms with a reliable verification tool, this project contributes significantly to making the digital recruitment landscape safer and more trustworthy.

VI. DISCUSSION AND FUTURE WORK

A. Discussion

The study's results demonstrate that combining textual semantics with categorical features in a unified architecture yields a highly effective fraud detection system. The CNN's ability to learn spatial hierarchies of features from text, such as phrases and sentence structures commonly found in deceptive

advertisements, is the primary driver of its superior performance. This suggests that fraudulent postings, despite mimicking legitimate ones, often contain subtle linguistic artifacts that can be automatically learned by such models.

Furthermore, the success of the real-time web application underscores the viability of deploying such systems in production. The ability to provide instantaneous feedback to users is crucial for trust and usability, making this system a practical tool for integration into existing job portals or for use as a standalone verification service.

B. Future Work

Although While the current system is highly effective, several milestones for future enhancement are apparent:

1. *Dataset Expansion and Generalization:* The model's robustness can be further improved by teaching and evaluating it on a multi diverse dataset, including job postings from multiple sources and in various languages, to ensure it generalizes across different cultural and linguistic contexts.
2. *Real-Time Scanning and Integration:* The system could be enhanced to function as a real-time scanning API, enabling job portals to automatically flag or block suspicious postings as they are submitted.
3. *Explainability and Trust:* Integrating techniques like LIME or SHAP could provide users with enough explanations for the model's predictions (e.g., identifying specific words or phrases that contributed to a "fake" classification), increasing trust and aiding in educational efforts.
4. *Adaptive Learning:* Implementing a feedback loop where user reports of undetected scams can be used to retrain and continuously improve the model, making it more resilient to evolving fraudulent strategies.
5. *Phishing and Link Detection:* Expanding the system's capability to analyze and flag suspicious external links embedded within job postings, a common vector for phishing attacks.

ACKNOWLEDGMENT

The author expresses their sincere gratitude to the academic mentors and the Department of Computer Science and Engineering at R.L. Jalappa Institute of Technology for their invaluable guidance and support. The open-source community and the

creators of the Fake Job Postings dataset are also acknowledged for providing the essential resources that made this research possible

VII. REFERENCES

- [1] S. Vidros, C. Kolias, G. Kambourakis, and L. Akil, "Automatic detection of online recruitment fraud: A machine learning approach," in *Proc. 2017 IEEE Int. Conf. Inf. Intell. Syst. Appl. (IISA)*, 2017, pp. 1-6.
- [2] D. Dutta and S. Bandyopadhyay, "Single vs. Ensemble Machine Learning Models for Fake Job Posting Detection," in *Proc. Int. Conf. Adv. Comput. Commun. Eng. (ICACCE)*, 2020, pp. 1-5.
- [3] V. Lokku, R. R. Reddy, and P. S. R. Ch, "Detection of Fake Job Postings using Machine Learning," in *Proc. Int. Conf. Emerg. Trends Inf. Technol. Eng. (ICETITE)*, 2021, pp. 1-4.
- [4] A. Nasser, R. A. Saeed, and S. A. Alsharif, "An approach for detecting fake job postings using artificial neural network," *Int. J. Adv. Comput. Sci. Appl.*, vol. 12, no. 8, pp. 297-304, 2021.
- [5] M. S. Akram, M. N. Akbar, and S. A. Khan, "An Enhanced Deep Learning Framework for Fake Job Posting Detection Using BERT and RoBERTa," *IEEE Access*, vol. 12, pp. 24312-24328, 2024.
- [6] F. Habiba, A. G. S. Al-Ghamdi, and A. A. Darwish, "A Comparative Study of Machine Learning and Deep Learning Techniques for Fake Job Posting Detection," in *Proc. 2021 Int. Conf. Data Min. Mach. Learn. (DMML)*, 2021, pp. 45-51.
- [7] Kaggle. (2017). *Fake Job Postings Dataset*. [Online]. Available: <https://www.kaggle.com/datasets/shivamb/real-or-fake-fake-jobposting-prediction>
- [8] D.-H. Vu, K. Nguyen, K. T. Tran, B. Vo, and T. Le, "Improving fake job description detection using deep learning-based NLP techniques," *J. Inf. Telecommun.*, vol. 9, no. 1, pp. 113-125, 2025.
- [9] S. S. Suhas Sanisetty et al., "Comprehensive Approach to Fraudulent Job Post Detection Using Machine Learning and BERT Models," in *Proc. 2025 4th Int. Conf. on Distributed Computing and Electrical Circuits and Electronics (ICDCECE)*, Ballari, India, 2025, pp. 1-6.
- [10] "Fraud-BERT: transformer based context aware online recruitment fraud detection," *Discover Computing*, vol. 28, no. 9, 2025.
- [11] "Evaluation of CNN-BiGRU and CNN-BiLSTM Model for Fake Job Post Detection: A Deep Learning Approach," in *Proc. 2024 IEEE Int. Conf. on Recent Advances in Inf. Technol. (RAIT)*, Nagpur, India, Nov. 2024.
- [12] "Machine Learning-powered Job Offer Verification," in *Proc. 2025 Int. Conf. on Adv. Comput. Commun. (ICACC)*, Chennai, India, Mar. 2025.
- [13] "Fake Job Detection using Naive Bayes and TF-IDF in Machine Learning," in *Proc. 2025 Int. Conf. on Smart Technol. for Sustainable Develop. (STSD)*, Pimpri, India, Aug. 2025.