



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** X **Month of publication:** October 2026

DOI: <https://doi.org/10.22214/ijraset.2026.85010>

www.ijraset.com

Call: ☎ 08813907089

E-mail ID: ijraset@gmail.com

Fake Job Detection using Hybrid Machine Learning and Explainable Risk Analysis

Ch. Radhika¹, Sarayu Vodnala², Sathwika Kurna³, A. Sanvi Reddy⁴

¹Professor, Department of CSE, G Narayanamma Institute of Technology and Science Hyderabad, India

^{2,3,4}Student, Department of CSE, G Narayanamma Institute of Technology and Science Hyderabad, India

Abstract: *With the rapid growth of online job portals, identifying genuine job opportunities has become increasingly difficult due to the rise of fake job postings. These fraudulent listings often mislead job seekers, waste their time, and may even expose them to scams or data theft. Manually checking every job posting is time-consuming and inefficient, making it necessary to have an automated system that can detect fake listings accurately.*

To overcome these challenges and improve online job safety, a system is developed to automatically classify job postings as real or fake based on key information such as job details, company data, and posting patterns. This system helps users identify trustworthy job opportunities quickly and reduces the chances of falling victim to fraud. By minimizing manual effort and enhancing reliability, it creates a safer and more efficient online job-seeking experience.

Index Terms: Fake Job Detection, Fraud Detection, Machine Learning, Natural Language Processing, TF-IDF, Explainable AI, Risk Analysis, Text Classification, Web Application, Recruitment Fraud.

I. INTRODUCTION

The rapid growth of online recruitment platforms has significantly transformed the way individuals search for employment opportunities. Millions of job postings are published daily across digital platforms, enabling faster and broader access to career opportunities. However, this expansion has also led to a substantial increase in fraudulent job postings, which pose serious risks to job seekers. These deceptive listings often include misleading information, unrealistic salary offers, and requests for sensitive personal or financial details, resulting in financial loss, identity theft, and reduced trust in online recruitment systems [1]–[4].

Identifying fraudulent job postings is a challenging task due to the unstructured nature of job descriptions and the evolving strategies used by scammers. Traditional approaches, including manual verification and basic filtering techniques, are time-consuming and often ineffective in detecting subtle patterns of fraud. As a result, there is a growing need for automated systems that can analyze job postings efficiently and identify suspicious characteristics with high accuracy [5].

Machine learning (ML) has emerged as a powerful approach for addressing such problems [6] by enabling systems to learn patterns from large datasets. By applying Natural Language Processing (NLP) techniques [7] to textual job descriptions, ML models can capture linguistic and contextual features that distinguish legitimate postings from fraudulent ones. These models can analyze complex relationships between different attributes, improving detection accuracy compared to traditional rule-based methods.

However, relying solely on machine learning models introduces certain limitations. Many high-performing models operate as black-box systems, making it difficult to interpret their predictions [8]. In applications such as fraud detection [9], where user trust is critical, the lack of transparency can reduce system reliability and usability. Users often require clear explanations to understand why a job posting is classified as suspicious, especially when making important career decisions.

To address these challenges, there is a need for systems that not only provide accurate predictions but also offer interpretable and meaningful insights. Combining machine learning techniques with additional analysis mechanisms can enhance both detection performance and transparency, enabling users to better evaluate the authenticity of job postings.

This paper presents an investigation into the application of intelligent techniques for detecting fraudulent job postings. The study focuses on analyzing textual and contextual information from job descriptions to identify patterns associated with recruitment fraud.

The rest of the paper is organized as follows: Section II discusses the background and related work in job fraud detection. Section III describes the system framework and methodology used for analysis. Section IV presents the experimental results and evaluation. Section V discusses the implications of the findings, and Section VI concludes the paper with future directions.

Finally, there is an evident problem concerning the transparency of decision-making processes within prediction algorithms. High-performance systems based on machine learning, including ensembles and boosting techniques, are mostly black-boxes with no

explanation of how they make decisions. In the context of the discussed project, which aims to detect fraudulent job listings in order to assist people with career choices, lack of transparency decreases usability and reliability.

Moreover, since job scammers constantly adapt their strategies in order to evade existing detection mechanisms, there is a danger that models trained with static datasets might become outdated rather fast [10]. Thus, it becomes critical to have systems capable of incorporating numerous detection methodologies.

A. Towards Explainable and Reliable Detection Systems

Modern research efforts in the area are directed at developing systems which not only achieve high performance in terms of classification accuracy but also generate an explanation of their decision process. Contrary to simple yes-no answers provided by many existing models, nowadays the focus is on generating an interpretation of what features contribute to the prediction of fraud.

One of the factors ensuring that a system is reliable is its capacity to consistently and accurately generate prediction regardless of inputs [11]. It means that an explainable prediction should be accompanied by several analysis tools: text feature extraction, heuristic rules, and additional verification of communication channels or links.

There are grounds to argue that systems using machine learning along with rule-based analysis [12], [13] and verification procedures are more efficient when it comes to identification of fraudulent job advertisements.

Our approach also employs this strategy while building upon existing advancements.

II. EXPERIMENTAL FRAMEWORK

In this section, the approach used in this study for detecting fraudulent job postings is described, along with the use of the dataset, preprocessing of data, feature extraction, and predictive model construction.

A. Dataset and Preprocessing Approach

- 1) *Dataset Description:* The dataset used in this study is called Employment Scam Aegean Dataset (EMSCAD), a commonly used benchmark dataset for fraudulent job detection task [14]. The dataset is available from the Kaggle website and comprises real-world examples of job advertisements. The EMSCAD dataset includes many attributes such as job titles, job description, company profile, location, job requirements, benefits, salary, among others, and also a label indicating whether the job posting is either legitimate or fraudulent. Since the EMSCAD dataset is composed of textual information on the job listings with a variety of structured data attributes, it provides an opportunity to conduct NLP and classification tasks. As a minority class of examples in the dataset, fraudulent job postings create an unbalanced distribution of data.
- 2) *Data Preprocessing:* Preprocessing of data before employing machine learning algorithms for predictive modeling is an essential part of any study. For example, data cleaning involves the removal of non-textual parts of a message or email. Text fields are normalized through case transformations. Other common steps include URL and other special character filtering, tokenization and removal of stop words to obtain preprocessed texts ready for further processing. Missing or irrelevant attribute values in a dataset are addressed by removal of unnecessary columns and keeping only those required for modeling. To compensate for the unbalanced distribution between legitimate and fraudulent job postings in the dataset, SMOTE technique is employed, helping balance the dataset and avoid model bias toward a particular type of posts. Finally, the textual preprocessed data is converted into feature vectors based on TF-IDF technique [7] [15].
- 3) *Dealing With Missing Values:* One of the data quality problems that arise from the EMSCAD dataset includes incomplete values in various fields like company profile, salary information, and job description among others. Several records in the dataset have missing or inconsistent information in some of their attributes which means removal of such cases from the data would lead to significant reduction in the size of the data and might end up deleting valuable information regarding fraudsters' job postings. To avoid such scenarios, the data should be preprocessed such that irrelevant attributes are removed, while those that can aid in predicting a post's legitimacy are left. In case of textual attributes, missing values are replaced with empty strings in order to allow consistent preprocessing by the natural language processing technique.
- 4) *Data Splitting and Data Representation:* The data that has been cleaned and preprocessed will then be split into training and testing datasets using an 80/20 split to test whether models perform well in identifying fraudulent job posts. Preprocessing techniques should also be used to balance the distribution of fraudulent and non-fraudulent job postings in the data. While in numerical data there is no need of transforming it prior to using it for modeling purposes, in the case of texts, transformation is necessary. This entails conversion of textual information to feature vectors using techniques like Term Frequency-Inverse Document Frequency (TF-IDF).

TF-IDF generates high dimensional feature representations which do not require further scaling prior to training any machine learning model.

B. Machine Learning Algorithms

To figure out if a job description is real or not we looked at machine learning algorithms that can predict things ...traditional statistical approaches for fraud detection [16]. We started by making some machine learning classifiers using job description words and their importance. We then compared how well they worked by looking at how accurate they were how precise they were, how well they remembered things and their F1 score.

- 1) *Logistic Regression (LR)*: Logistic regression is a machine learning algorithm that can be used to classify things. It is good at giving us probabilities of something belonging to a group based on what we put in. Since it is simple logistic regression helps us see how each word in a job description affects the result. However it is not good at finding patterns in fake job postings because it is too simple.

$$P(Y = 1|X) = \frac{1}{1 + e^{-z}}$$

Where:

$$z = w_0 + w_1x_1 + w_2x_2 + \dots + w_nx_n$$

- 2) *Random Forest (RF)*: Random Forest is another algorithm that helps us predict what group something belongs to [17]. It is a way of learning that makes decision trees and uses them to predict things. Random Forest is better than regression because it can find more complicated patterns in the data. This means it can detect job postings more accurately by looking at the importance of words.

$$\hat{y} = \text{mode}\{T_1(x), T_2(x), T_3(x), \dots, T_n(x)\}$$

Where:

- $T_i(x)$ represents the prediction of the i^{th} decision tree
- $\hat{y}$ is the final predicted output

- 3) *XGBoost (Extreme Gradient Boosting)*: We also used XGBoost, which is a good machine learning algorithm for classification problems [16]. XGBoost is like a version of making many decision trees. It makes decision trees one after the other to fix the mistakes of the one, which makes it more accurate [18]. So the XGBoost algorithm can look at words in a job description and detect fake job posts more accurately, than other methods.

$$Obj = \sum_{i=1}^n l(y_i, \hat{y}_i) + \sum_{k=1}^K \Omega(f_k)$$

Where:

- $l(y_i, \hat{y}_i)$ represents the loss function
- $\Omega(f_k)$ represents the regularization term
- f_k represents each decision tree

TABLE I
COMPARISON OF MACHINE LEARNING CLASSIFIER PERFORMANCE

Algorithm	Accuracy	Precision	Recall	F1-Score
Logistic Regression	96.53%	96.72%	32.60%	48.76%
Random Forest	97.76%	99.03%	56.35%	71.83%
XGBoost	97.85%	95.61%	60.22%	73.90%

C. How We Analyze Explanations

- 1) *Understanding Risky Job Postings:* Our system has a special feature that helps us understand why a job posting is marked as suspicious. This feature looks at things like keywords strange salary patterns, bad links, grammar mistakes and unusual communication methods. When our system finds something it adds points to a fraud risk score. This score is then shown to the user along with an explanation of why the job posting's flagged. This way users can see why a job posting is suspicious. Our system is more transparent and trustworthy because of this. The explanations we provide come from both computer predictions and rule-based checks. They reflect the characteristics of the job posting.
- 2) *Combining Computer Predictions and Rule-Based Checks:* Our system uses both computer predictions and rule-based checks to analyze job postings. This combination gives us a risk assessment. The computer model looks for patterns in data and extra checks look at things like whether a domain's real whether application links are safe and whether a company is verified. This approach makes our system easier to understand. It links each prediction to features of the job posting, which makes our system more practical for real-world use. Our sys-tem is more practical than models that're hard to understand.

D. How We Measure Performance

We use a strategy to evaluate how well our system works. We compare the performance of computer models using clas-sification metrics. We evaluate our models using metrics like accu-racy, preci-sion, recall and F1-score. Accuracy measures how often our predictions are correct. Precision evaluates the proportion of job postings that we correctly identify as suspicious. Recall measures our systems ability to detect cases. F1-score provides a balance between precision and recall.

- 1) *Dealing with Imbalanced Data:* Since job postings are only a small part of our dataset we need to use evaluation metrics like recall and F1-score. These metrics are important because they ensure that our model can effectively identify cases even if they are rare. We do not just focus on accuracy.
- 2) *Evaluating Our System:* We evaluate how well our system works by looking at its ability to provide risk anal-ysis, including the clarity of explanations the identification of warning signs and the consistency of results across job postings. By combining computer predictions with rule-based checks we improve detection accuracy. Make our system more interpretable. This makes our system more reliable, for real-world applications.
- 3) *Comparing Machine Learning Models:* To figure out which machine learning models are the best we tried them out on our dataset. We looked at things like how accurate they're how precise they are how well they recall things and what their F1-score is to see how well each model can find job postings that are real [18].

III. RESULTS AND ANALYSIS

This section presents the results obtained from the implemented machine learning models and analyzes the effective-ness of the proposed system.




Fig. 1. User interface of the JobGuard system showing job input and analysis features.

A. Comparing Model Performance

We evaluated three models: Logistic Regression, Random Forest, and XGBoost on the dataset. The performance comparison of these models is presented in Table I.

The results indicate that XGBoost achieves the best overall performance, particularly in terms of recall and F1-score, making it highly effective for detecting fraudulent job postings. Random Forest also demonstrates strong performance, especially in precision, while Logistic Regression serves as a baseline but fails to detect a significant number of fraudulent cases.

Fig. 2 shows the comparative performance of the three models based on precision, recall, and F1-score. It is observed that XGBoost achieves the highest recall and F1-score, indicating its superior ability to correctly identify fraudulent job postings. Random Forest also performs well, particularly in terms of precision, while Logistic Regression shows lower recall, suggesting that it fails to detect a significant number of fraudulent cases. This highlights the effectiveness of ensemble models in handling complex textual patterns. Overall, ensemble-based models outperform the baseline model [19] as they can capture complex patterns in textual data more effectively.

Fig. 3 presents the confusion matrix of the XGBoost model. The model correctly classifies a large number of legitimate job postings (true negatives) and fraudulent postings (true positives). However, a small number of misclassifications are




Fig. 2. Comparison of precision, recall, and F1-score across Logistic Regression, Random Forest, and XGBoost models.

observed in the form of false positives and false negatives. The low number of false negatives indicates that the model is effective in identifying fraudulent job postings, which is critical in fraud detection systems.




Fig. 3. Confusion matrix of the XGBoost model showing classification performance.

B. System-Level Analysis

The proposed system not only classifies job postings as real or fake but also provides a detailed risk analysis explaining the reasons behind each prediction. The system identifies multiple risk indicators such as suspicious keywords, unrealistic salary values, unverified links, and non-corporate communication methods [20].

The system outputs are analyzed based on observable risk factors rather than internal model behavior. This allows users to clearly understand why a particular job posting is classified as fraudulent.




Fig. 4. Analysis result for a job posting showing high risk score and identified fraud indicators.

Fraudulent job postings typically exhibit multiple suspicious characteristics such as urgency-related keywords, abnormal salary offers, and unverified communication channels [20]. These factors collectively increase the risk score.




Fig. 5. Analysis result for a job posting with moderate risk score.

In contrast, legitimate job postings generally contain structured descriptions, verified company information, and formal communication methods, resulting in lower risk scores.

The system also provides explanations highlighting specific risk factors, improving transparency and user trust.




Fig. 6. Explanation of detected risk factors contributing to the fraud risk score.

C. Case Study

To further evaluate the system, a case study was conducted on a job posting identified as highly suspicious.

- 1) Job Details: The job posting included phrases such as “urgent hiring” and offered a high salary despite requiring minimal experience. It also requested communication through messaging platforms and used a non-corporate email address.
- 2) System Prediction: The system classified the job posting as highly fraudulent with a risk score close to 100%. A second case study was conducted on a legitimate job posting.
- 3) Job Details: The job posting contained structured information, verified company details, and formal communication methods.
- 4) System Prediction: The system assigned a low risk score, indicating that the job posting is likely genuine.

These case studies demonstrate that the system effectively classifies job postings and provides meaningful explanations for its predictions.

D. Evaluation of Explanation Quality

The system’s explanations are evaluated based on consistency and relevance. The identified risk factors accurately reflect the characteristics of fraudulent job postings.

The system also demonstrates stability by producing similar results for job postings with minor variations. Additionally, modifying key risk factors results in noticeable changes in risk scores, confirming that the system correctly identifies influential features.




Fig. 7. Analysis result for a legitimate job posting.




Fig. 8. Detailed analysis of a legitimate job posting.

Overall, the system provides reliable and interpretable ex-planations, improving transparency and user trust.

IV. CONCLUSION

This study shows a way to find job postings by using machine learning and risk analysis. The results show that some models, like XGBoost are really good at finding these job postings. The system also helps users by giving a score of how risky a job posting's explaining why it thinks that about the job posting.

Unlike methods that only say if a job is fake or not this system uses a mix of text analysis, risk evaluation and verification to look at job postings. This helps it find job postings and explain its answers about the job postings. By looking at things like words salaries that're too high and contact information that is not verified the system gives a better idea of what job postings are like.

The results show that using modeling and risk analysis together makes it better at finding job postings and makes users trust the system more when they are looking for a job. The systems ability to analyze things in time and its easy-to-use interface make it useful for people who are really looking for a job.

Future work can focus on using better language models real-time job data and feedback from users to make the system better at finding fake job postings and fighting recruitment fraud that involves fake job postings.

V. DECLARATION ON GENERATIVE AI


We used AI tools to improve grammar and language while writing this article about job postings and the system to find them. We are responsible for the content and integrity of the work, about job postings.

REFERENCES

- [1] A. Vidros, C. Kolias, and G. Kambourakis, "Detecting fraudulent job advertisements using machine learning," in *Proc. 11th Int. Conf. Availability, Reliability and Security (ARES)*, 2016.
- [2] K. Jain and R. Gupta, "Online recruitment fraud detection using machine learning," *IJSREM*, 2022.
- [3] P. Resnick et al., "Reputation systems for online communities," *Communications of the ACM*, vol. 43, no. 12, 2000.
- [4] J. Fette, N. Sadeh, and A. Tomasic, "Learning to detect phishing emails," in *WWW Conference*, 2007.
- [5] H. Shahriar and M. Zulkernine, "Trustworthiness evaluation of job postings using data mining techniques," in *IEEE Conf. on Trust, Security and Privacy in Computing*, 2012.
- [6] M. Alghamdi and H. Alfalqi, "A survey of machine learning techniques for cyber security," *Journal of Information Security*, 2015.
- [7] S. Bird, E. Klein, and E. Loper, *Natural Language Processing with Python*. O'Reilly Media, 2009.



- [8] S. Lundberg and S. Lee, "A unified approach to interpreting model predictions," in NeurIPS, 2017.
- [9] V. Upadhyay and K. Khandelwal, "Applying artificial intelligence in recruitment: A review," International Journal of Advanced Research in Computer Science, 2018.
- [10] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," ACM Computing Surveys, vol. 41, no. 3, 2009.
- [11] J. Sharples and J. Domingue, "The blockchain and kudos: A distributed system for educational record, reputation and reward," in European Conference on Technology Enhanced Learning, 2016.
- [12] S. Roy and D. Banerjee, "Combining textual and behavioral analysis for fraud detection in online recruitment systems," International Journal of Computer Applications, vol. 182, no. 45, 2019.
- [13] R. Bansal, A. Kumar, and S. Singh, "Hybrid machine learning approach for fake job detection using naïve bayes and logistic regression," International Journal of Innovative Science and Research Technology, vol. 5, no. 6, 2020.
- [14] Y. Huang, Q. Qiu, and S. Huang, "Detecting fraudulent job postings using deep learning techniques," IEEE Access, vol. 8, 2020.
- [15] T. Mikolov et al., "Efficient estimation of word representations in vector space," arXiv preprint arXiv:1301.3781, 2013.
- [16] R. Bolton and D. Hand, "Statistical fraud detection: A review," Statistical Science, vol. 17, no. 3, 2002.
- [17] L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.
- [18] S. Bhattacharyya et al., "Data mining for credit card fraud detection: A comparative study," Decision Support Systems, vol. 50, 2011.
- [19] M. Sharma and P. Gupta, "Multi-layered fraud detection system using machine learning techniques," International Journal of Engineering and Technology, vol. 7, no. 4, 2018.
- [20] J. Eberle and L. Holder, "Behavior-based fraud detection in online systems," IEEE Intelligent Systems, 2009.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)