International Journal of Scientific Research and Engineering Development-– Volume 9 Issue 2, Mar-Apr 2026
Available at www.ijsred.com

RESEARCH ARTICLE OPEN ACCESS

Fake Job Post Prediction System Using Machine Learning Approach

Pranavi V. Armal[1], Suyash R. Tamkhane[2], Bhumi S. Parsawar[3], Aryan S. Rathod[4], Dr. R. R. Keole[5]
[1234]Undergraduate Student, [5]Head of Department
Department of Information Technology
Shree H.V.P.M. COET, Amravati (MH, 27)
pranaviarmal@gmail.com, suyashtamkhane@gmail.com

----------**********----------

Abstract:

The rapid growth of online recruitment platforms has simplified job searching, but it has also increased the risk of fraudulent job advertisements that mislead applicants and exploit personal information. This study presents a machine learning-based Fake Job Post Prediction System to identify suspicious job postings using both textual and structured job-related attributes. The proposed framework incorporates advanced text preprocessing and feature extraction using TF-IDF, along with classification models such as Passive-Aggressive Classifier and Multi-Layer Perceptron. To address class imbalance, the Synthetic Minority Over-sampling Technique (SMOTE) is applied, improving the model's ability to detect fraudulent cases. Experimental results on a publicly available dataset demonstrate enhanced performance compared to baseline approaches, with improved accuracy and reduced false positives. Furthermore, the system is integrated into a web-based application, enabling real-time prediction and user interaction. The proposed solution provides a scalable and efficient approach to enhancing trust and security in online recruitment systems.

Keywords — **fake job post, prediction system, machine learning, mlp classifier, passive-aggressive classifier.**

----------**********----------

I. INTRODUCTION

The rapid expansion of online recruitment platforms has transformed the hiring process by making job discovery faster, broader, and more convenient for both employers and applicants. Organizations now rely heavily on digital job portals to advertise vacancies, filter candidates, and manage recruitment workflows. At the same time, this shift has also created an attractive opportunity for cybercriminals. Fraudulent job advertisements are increasingly used to deceive job seekers into sharing personal information, paying illegal fees, or engaging with fake employers under the impression that they are applying for legitimate opportunities [1]. As online hiring becomes more common, the ability to automatically distinguish genuine job posts from fake ones has become an important research and practical problem. Fake job postings are difficult to detect manually because many of them imitate the appearance and language of real recruitment messages. Fraudsters often use persuasive titles, vague descriptions, unrealistic benefits, and incomplete company information to attract victims. In some cases, the purpose is identity theft; in others, it may involve financial scams, affiliate marketing traps, or corporate impersonation [1], [5]. The problem is especially serious because job seekers are often in vulnerable situations and may respond quickly to attractive but deceptive opportunities. This makes the cost of false trust extremely high,

involving privacy loss, financial damage, and reputational risk.

This paper presents a Fake Job Post Prediction System using machine learning, built on the publicly known 17,880-record employment scam dataset introduced by Vidros et al. [1]. The implemented system focuses on the textual core of a job advertisement by combining the title, company profile, description, requirements, and benefits fields into a unified text representation. The text is transformed using TF-IDF vectorization, and the class imbalance problem is handled using SMOTE. Two classification models are used: a Passive-Aggressive Classifier and a Multi-Layer Perceptron (MLP). In addition to these learned models, the deployed application includes rule-based support from structured fields such as telecommuting status, company logo presence, screening questions, employment type, education, and experience level. This hybrid design helps the system remain useful even when the text is sparse, noisy, or intentionally manipulated.

II. LITERATURE REVIEW

Research on online recruitment fraud is relatively recent when compared with broader fraud detection or spam classification literature. One of the foundational studies in this field was presented by Vidros et al. [1], who formally described online recruitment fraud as a distinct cybersecurity problem and introduced the Employment Scam Aegean Dataset (EMSCAD), a public dataset of 17,880 annotated job advertisements collected from a real-life recruitment platform. Their work was important not only because it framed the problem clearly, but also because it made systematic experimentation possible for later studies. Vidros et al. argued that textual and metadata-based classification can act as a useful starting point, although they also noted that future solutions may require a more composite approach involving user, organization, and network-level evidence.

Following this early contribution, later researchers began investigating more specialized forms of fraud detection in recruitment systems. Alghamdi and Alharby [2] proposed an intelligent model for online recruitment fraud detection and emphasized the value of combining data mining with machine learning to reduce privacy and financial risks faced by job seekers. Their work highlighted that recruitment fraud is not merely a text classification problem, but part of a broader security challenge in online platforms. This perspective is relevant because many fake postings are intentionally crafted to appear legitimate and therefore require more than simple keyword spotting.

A different direction was explored by Tabassum et al. [3], who studied machine learning methods for online recruitment fraud detection in a conference setting. Their work demonstrated that conventional classifiers can be effectively used for this task and helped strengthen the case for supervised learning as a practical solution. Such research supports the choice of using established models like Passive Aggressive classifiers and neural networks in applied systems such as the one developed in this project. More recent work has moved beyond binary fraud detection and examined the internal types of fraudulent job advertisements. Naude, Adebayo, and Nanda [5] used the EMSCAD dataset to classify fraudulent job types rather than simply identifying whether a post was fake or real. Their findings showed that different forms of fraud, such as identity theft, corporate identity theft, and multi-level marketing schemes, may have distinguishable linguistic and structural patterns. They also observed that combinations of rule-based features, part-of-speech information, and text-based representations can improve classification effectiveness. This supports the idea that hybrid systems, which combine learned textual features with crafted structured indicators, may be more robust than purely text only solutions.

Mahbub, Pardede, and Kayes [4] contributed another important insight by focusing on contextual features in Australian job industries. Their study suggested that recruitment fraud detection can benefit from domain specific contextual information rather than relying solely on surface textual cues. This is consistent with the design of the current project, where structured fields such as company logo presence, telecommuting setting, and screening questions are used alongside model predictions. In practice, a suspicious posting is often revealed not only by what it says, but also by what information it omits or how its structured profile differs from legitimate hiring norms.

Regional and platform-specific studies have also been conducted. Dake [6], for example, proposed a machine learning-based model for Ghanaian job websites and reported that Random Forest achieved 91.86% accuracy on the studied dataset. This work is useful because it shows that the recruitment fraud problem is not limited to one geography or one platform type. It also suggests that

data distribution and platform context can meaningfully influence which model performs best. In comparison, the current project uses the broader EMSCAD-style dataset and relies heavily on text feature engineering with TF-IDF, which likely contributes to the stronger observed classification performance.

Deep learning approaches have also started to appear in this area. Pillai [7] proposed a Bidirectional LSTM model for fake job posting detection and reported 98.71% accuracy with a ROC-AUC of 0.91. This work demonstrates that sequence aware deep learning architectures can capture richer dependencies in textual job advertisements. However, deep learning models generally require more computational resources and maybe more difficult to deploy in lightweight web applications, especially when interpretability and inference speed matter. In the current project, the MLP provides a simpler neural approach that still performs extremely well on the available data, while the Passive Aggressive model offers even faster inference.

Another related deep learning study described fake news over job posts using Bi-LSTM-based detection [8]. Although the framing overlaps partly with fake-news terminology, the core idea remains aligned with fraudulent job post classification. The study underlines the continuing shift from manual feature engineering toward models that learn sequential language patterns directly from data. Still, for many applied systems, sparse-vector approaches such as TF-IDF remain competitive because they are efficient, interpretable, and relatively easy to integrate into web-based decision tools. A broad pattern emerges across the literature.

First, the EMSCAD dataset introduced by Vidros et al. [1] remains the most influential benchmark in this domain. Second, text processing is consistently central to the problem because the meaning, tone, and completeness of a job advertisement carry important fraud signals. Third, context and structured metadata improve detection when combined with text-based models [4], [5]. Finally, there is no single universally best model across all datasets and settings. Some studies report good results with Random Forest [6], others with feature-rich hybrid models [5], and recent work also shows promise for deep learning [7], [8].

The current project fits naturally within this literature. It follows the core benchmark tradition established by EMSCAD, uses TF-IDF to capture textual evidence, applies SMOTE to address imbalance, compares a linear classifier and a neural classifier, and extends model prediction with structured rule-based scoring. In that sense, the system does not merely repeat prior work; rather, it reflects a practical synthesis of major ideas already established in the literature: benchmark-driven learning, text-centered representation, imbalance handling, and hybrid decision support for deployment-oriented fraud detection.

III. PROPOSED METHODOLOGY:

The methodology of the Fake Job Post Prediction System is designed to combine effective text-based learning with practical deployment in a user-facing prediction platform. The system follows a structured pipeline consisting of data collection, preprocessing, feature extraction, imbalance handling, model training, evaluation, and final deployment. The primary objective is to accurately distinguish between real and fraudulent job postings.

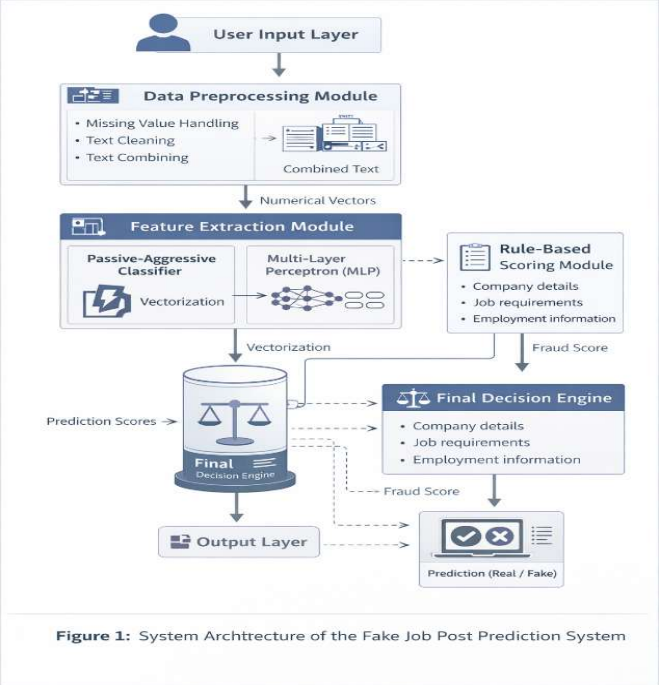

Figure 1: System Architecture of the Fake Job Post Prediction System

Figure 1. System Architecture of fake job post prediction system

A. Dataset Description

The dataset used in this study contains 17,880 job postings with 18 attributes. The target attribute, fraudulent, indicates whether a job post is real or fake.
The dataset is highly imbalanced:

- 17,014 posts are labeled as real
- 866 posts are labeled as fake

This imbalance can bias the model toward the majority class, reducing its ability to detect fraudulent postings. Therefore, handling class imbalance is an essential part of the methodology.

B. Data Preprocessing

In the preprocessing stage, the most relevant text-based columns were selected:

1. Job Title
2. Company Profile
3. Description
4. Requirements
5. Benefits

These fields are highly informative, as fraudulent job postings often contain:

1. Vague or misleading descriptions
2. Unrealistic benefits
3. Missing or incomplete company details

The following preprocessing steps were performed:

- Missing values were replaced with empty strings
- Selected text fields were combined into a single feature called combined_text
- This allowed the model to analyze the complete semantic content of each job post as a unified text input.
- Additionally, the target labels were converted into numeric format:
- Real posts → 0
- Fake posts → 1

This conversion is necessary for machine learning algorithms.

C. Feature Extraction

To convert textual data into numerical format, TF-IDF (Term Frequency–Inverse Document Frequency) vectorization was applied.
TF-IDF transforms each job post into a numerical vector based on word importance relative to the dataset.
The following configurations were used:

- Removal of English stop words
- Maximum feature limit of 10,000

This helped reduce noise and ensured that only meaningful and discriminative words were considered. The resulting TF-IDF matrix was used as the input for model training.

D. Handling Class Imbalance

Due to the imbalanced nature of the dataset, SMOTE (Synthetic Minority Oversampling Technique) was applied to improve the model's ability to detect fraudulent job postings.

- SMOTE generates synthetic samples for the minority (fake job) class
- This helps create a more balanced training dataset
- It reduces bias toward the majority (real job) class

E. Model Training

The dataset was divided into training and testing sets in an 80:20 ratio with stratification.
Two machine learning models were used:

1. Passive-Aggressive Classifier:

- Suitable for high-dimensional sparse data
- Efficient for text classification tasks
- Updates aggressively when misclassification occurs
- Maximum iterations: 1000

2. Multi-Layer Perceptron (MLP):

- A feedforward neural network
- Capable of learning complex non-linear patterns
- Configuration:
- Hidden layer: 100 neurons
- Maximum iterations: 300

E. Model Evaluation

The models were evaluated using standard classification metrics:

1. Accuracy
2. Precision
3. Recall
4. F1-Score

Additional evaluation techniques included:

1. Confusion Matrix
2. ROC Curve
3. Precision-Recall Curve
4. Prediction Distribution Analysis
5. Inference Time Comparison

These metrics provide a comprehensive understanding of model performance beyond accuracy.

F. Deployment Strategy

The system is designed for real-world usability through a user-facing application.
Workflow:

- User inputs job post details
- Input is transformed using saved TF-IDF vectorizer
- Predictions are generated using trained models

The final decision is based on a hybrid approach:

- 70% weight from MLP probability score
- 30% weight from rule-based fraud scoring

Rule-based features include:

- Telecommuting availability
- Company logo presence
- Screening questions
- Employment type
- Required education
- Required experience

Additional Features:

- Safety fallback logic for suspicious inputs
- Handling missing or low-information text
- Detection of unusual attribute combinations

This improves reliability in real-world scenarios where inputs may be noisy or incomplete.

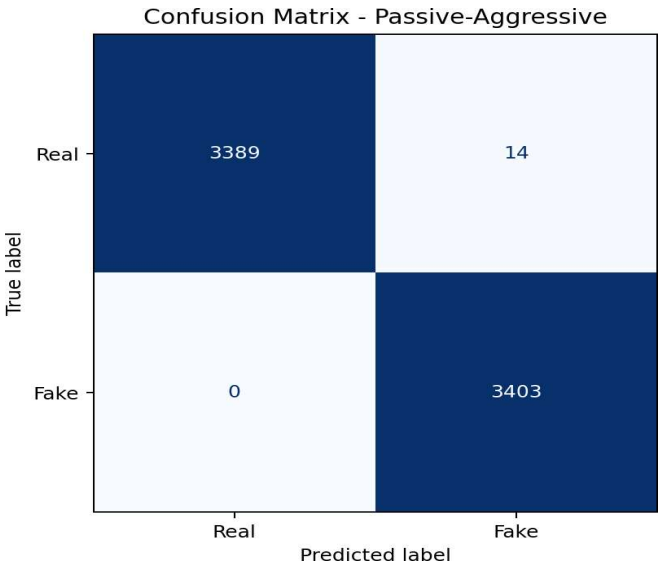

IV. IMPLEMENTATION DETAILS

- Programming Language: Python
- Web Framework: Flask
- Machine Learning Library: Scikit-learn
- Data Handling: Pandas, NumPy
- Model persistence: joblib
- Database: SQLite
- Deployment: Local inference

The system works using both a terminal interface and a

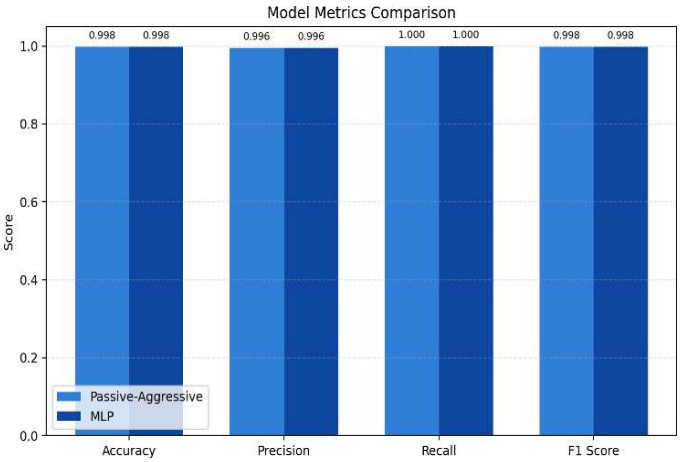

web-based interface.

V. EXPERIMENTAL RESULT

Two machine learning models, Passive Aggressive Classifier and Multi-Layer Perceptron (MLP), were evaluated on the processed dataset. Both models demonstrated strong classification performance across key evaluation metrics, including accuracy, precision, recall, and F1-score. The models were able to effectively distinguish between real and fraudulent job postings.

Figure 2. Comparison of Accuracy, Precision, Recall,

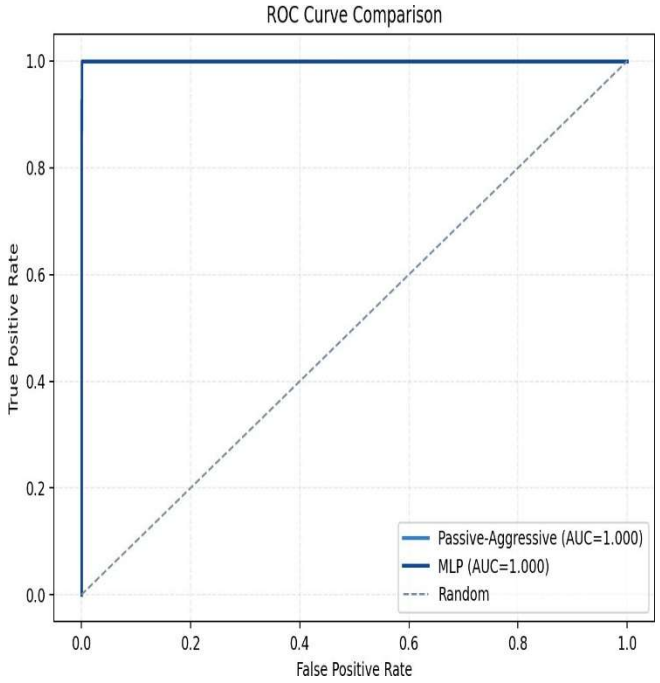

and F1-score for the Passive-Aggressive and MLP models.

Figure 3. Confusion matrix of the Passive-Aggressive Classifier.

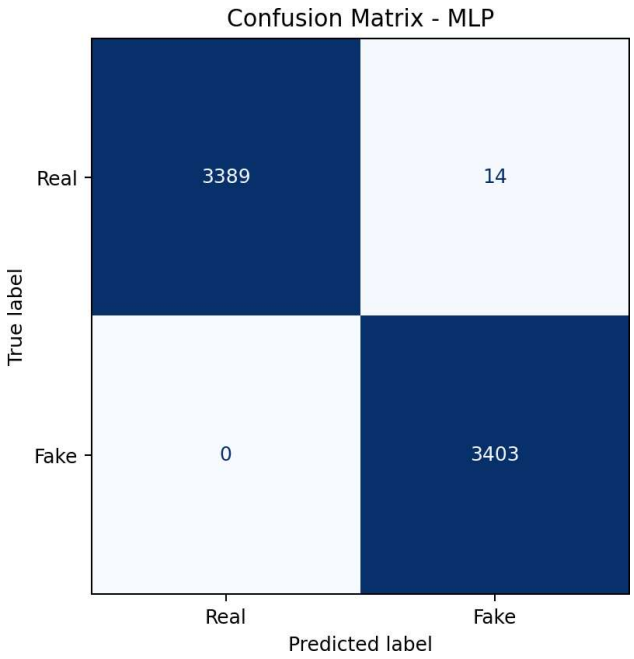

Figure 4. Confusion matrix of Multi-Layer Perceptron

Figure 5. ROC curve comparison of Passive-Aggressive and MLP models.

Both models showed strong performance across evaluation metrics such as accuracy, precision, recall, and F1-score. The results indicate that the system is capable of effectively distinguishing between real and fraudulent job postings. Confusion matrix analysis further demonstrated that most instances were correctly classified, with very few false predictions.

The models achieved high recall, which is particularly important in this domain as it ensures that fraudulent job posts are correctly identified. At the same time, precision was also maintained, reducing the chances of misclassifying genuine job posts.

In terms of performance, the Passive Aggressive Classifier was observed to be faster and more lightweight, making it suitable for large-scale or real-time applications, while the MLP model provided stable and reliable classification results.

Another important observation concerns inference efficiency. For a sample of 1,000 test instances, the Passive-Aggressive model completed prediction in approximately 0.381 ms, while the MLP required about 7.071 ms. Although both are fast enough for practical use, the Passive-Aggressive classifier is clearly more lightweight. This makes it attractive for high-throughput screening or deployment on resource-constrained systems.

On the other hand, the MLP provides probabilistic output and serves as the principal learned component in the final decision logic of the web application. In the deployed prediction flow, the final score is computed using a weighted combination of 70% MLP fraud probability and 30% rule-based structured fraud score. This is a meaningful engineering decision because it balances machine learning confidence with interpretable indicators derived from job metadata.

Overall, the experimental results show that the proposed system is effective, efficient, and suitable for fake job detection. However, the performance may vary in real-world scenarios due to changing fraud patterns and differences in data distribution, which can be explored in future work.

VI. LIMITATIONS

- The model is trained and evaluated on a SMOTE-balanced dataset, which may not reflect real-world class distribution.
- Performance may vary in real-world scenarios due to:
    - Changing language patterns
    - Different job platforms
    - Evolving fraud techniques
- The system relies on TF-IDF features, which:
- Cannot capture deep semantic meaning
- May fail on unseen or rare words
- Lack of external validation on real-world or cross-platform datasets.
- The model may require frequent updates to adapt to new fraud strategies.

VII. FUTURE WORK

- Evaluate the model on real-world and large-scale datasets.
- Apply deep learning models (e.g., LSTM, BERT) for better text understanding.

- Improve feature extraction beyond TF-IDF using contextual embeddings.
- Perform cross-domain testing across different job platforms.
- Continuously update the system to handle new fraud patterns.
- Enhance the user interface with:
    - Better visual explanations
    - Real-time feedback

VIII. CONCLUSION

This study presented a Fake Job Post Prediction System that combines machine learning and supporting rule-based analysis to detect fraudulent job advertisements. Using TF IDF feature extraction, SMOTE-based balancing, and two supervised learning models, the system achieved very high performance on the available dataset. Both the Passive Aggressive Classifier and the Multi-Layer Perceptron produced strong and consistent results, showing that textual content in job advertisements contains clear signals that can be used for fraud detection. The system was also integrated into a web-based application, making the model useful in a practical setting rather than only as an offline experiment. Although the results are promising, further improvement is still possible. Future work may include using larger and more recent datasets, incorporating additional structured features, trying deep learning or transformer-based language models, improving explainability, and testing the system on real-time recruitment platforms. These extensions can make the system more robust, scalable, and closer to real-world deployment.

ACKNOWLEDGMENT

We thank our guide and Head of Department for guidance throughout the project.

REFERENCES

[1] S. Vidros, C. Kolias, G. Kambourakis, and L. Akoglu, "Automatic Detection of Online Recruitment Frauds: Characteristics, Methods, and a Public Dataset," Future Internet, vol. 9, no. 1, p. 6, 2017. https://doi.org/10.3390/fi9010006

[2] B. Alghamdi and F. Alharby, "An Intelligent Model for Online Recruitment Fraud Detection," Journal of Information Security, vol. 10, no. 3, pp. 155-176, 2019. https://doi.org/10.4236/jis.2019.103009

[3] H. Tabassum, G. Ghosh, A. Atika, and A. Chakrabarty, "Detecting Online Recruitment Fraud Using Machine Learning," in 2021 9th International Conference on Information and Communication Technology (ICoICT), 2021, pp. 472-477. https://doi.org/10.1109/ICoICT52021.2021.9527477

[4] S. Mahbub, E. Pardede, and A. S. M. Kayes, "Online Recruitment Fraud Detection: A Study on Contextual Features in Australian Job Industries," IEEE Access, vol. 10, pp. 82776-82787, 2022. https://doi.org/10.1109/ACCESS.2022.3197225

[5] M. Naude, K. J. Adebayo, and R. Nanda, "A Machine Learning Approach to Detecting Fraudulent Job Types," AI & Society, vol. 38, pp. 1013-1024, 2023. https://doi.org/10.1007/s00146-022-01469-0

[6] D. K. Dake, "Online Recruitment Fraud Detection: A Machine Learning-based Model for Ghanaian Job Websites," International Journal of Computer Applications, vol. 184, no. 51, pp. 20-28, 2023. https://doi.org/10.5120/ijca2023922639

[7] A. S. Pillai, "Detecting Fake Job Postings Using Bidirectional LSTM," International Research Journal of Modernization in Engineering Technology and Science, vol. 5, no. 3, 2023. arXiv:2304.02019. https://doi.org/10.48550/arXiv.2304.02019

[8] B. G. Banik, "Detecting Fake News Over Job Posts via Bi-Directional Long Short-Term Memory (BIDLSTM)," International Journal of Cognitive Informatics and Natural Intelligence, 2021

[9] A. Amaar, W. Aljedaani, F. Rustam, S. Ullah, V. Rupapara, and S. Ludi, "Detection of Fake Job Postings by Utilizing Machine Learning and Natural Language Processing Approaches," Neural Processing Letters, vol. 54, no. 3, pp. 2219–2247, Jun. 2022, doi: 10.1007/s11063-021-10727-z.

