Journal of Information & Optimization Sciences
ISSN 0252-2667 (Print), ISSN 2169-0103 (Online)
Vol. 47 (2026), No. 5-B, pp. 2109–2117
DOI : 10.47974/JIOS-2301

Fraud-Detect-Net : A deep learning and machine learning framework for detecting fake job postings

Sonika Malik *
Preet Rathee †
Anupama Kaushik §
Prabhjot Kaur ‡
Akshay Singh @
Department of Information Technology
Maharaja Surajmal Institute of Technology
Janak Puri 110058
Delhi
India

Abstract

This study presents a hybrid framework combining BERT embeddings with a Random Forest classifier to detect fraudulent job postings. To address the severe class imbalance between genuine and fake listings, multiple oversampling methods—including SMOTE, SMOTE Tomek, ADASYN, SMOTE-ENN, and random oversampling—were evaluated. Model performance was further optimized through hyperparameter tuning using Randomized Search. The system was assessed with key metrics such as precision, recall, F1-score, and AUC-ROC, ensuring a balanced evaluation. Results demonstrate that leveraging contextual language representations from BERT, together with data balancing strategies, significantly improves fraud detection accuracy. This approach highlights the effectiveness of integrating advanced NLP models with resampling techniques for reliable online job fraud detection.

Subject Classification: *Primary 68T07.*

Keywords: *Machine learning, Deep learning, BERT (Bidirectional encoder representations from transformers), SMOTE (Synthetic minority oversampling technique), Class imbalance.*

* *E-mail:* **Sonika.malik@msit.in (Corresponding Author)**
† *E-mail:* **Preetirathee@msit.in**
§ *E-mail:* **anupama@msit.in**
‡ *E-mail:* **prabhjot.kaur@msit.in**
@ *E-mail:* **Akshay.singh@msit.in**

1. Introduction

This study focuses on developing a reliable and interpretable machine learning framework for detecting fraudulent job postings by addressing the challenges of class imbalance, unstructured text, and model optimization. Since genuine postings vastly outnumber fraudulent ones, the work evaluates various oversampling strategies, classifiers, and hyperparameter tuning methods to improve detection while maintaining interpretability and trust. The research also emphasizes the importance of handling unstructured job descriptions through advanced NLP techniques like BERT embeddings, which help capture meaningful patterns from inconsistent text data. By systematically testing different approaches, the framework aims to create a robust detection system that balances precision and recall, offering practical guidelines for fraud detection in recruitment platforms. The paper is structured with related work in Section 2, methodology in Section 3, results in Section 4, and conclusions with future directions in Section 5.

2. Literature Review

Research on fraudulent job posting detection has progressed from basic classifiers to more advanced ensemble and deep learning methods. Early work, such as Vidros et al. [1], introduced the EMSCAD dataset and showed Random Forest as a strong performer with over 91% precision, while later studies explored ensembles like AdaBoost and Gradient Boosting for improved accuracy [2]. Other efforts applied SVMs with high precision [3], ensemble voting methods [4], and imbalance-handling techniques like down sampling [5] and oversampling [10] to boost results. Deep learning approaches further advanced performance, with DNNs achieving up to 99% accuracy [6, 7]. Beyond EMSCAD, datasets like IESD [8] expanded the scope, and alternative classifiers such as KNN were also explored [9]. Recent works have increasingly turned to NLP and imbalance-aware methods, where BERT, LSTM, and CNN models [11–13] combined with sampling strategies have shown strong results, particularly in fraud and harassment detection, while studies on transactional fraud [14] reinforced the need for balance-aware frameworks. Overall, these findings highlight that both model design and class balancing are crucial for achieving robust detection in fraud-related applications.

3. Methodology

This study applied BERT embeddings with a Random Forest classifier to detect fraudulent job postings, addressing class imbalance through various oversampling techniques. The model was optimized with Randomized Search CV and evaluated using AUC-ROC, classification reports, and confusion matrices, showing the impact of sampling methods and tuning on performance.

3.1 *Data Exploration and Preprocessing*

The "Fake Job Postings" dataset [15] with 17,880 entries (less than 5% fraudulent) was analyzed using EDA to identify key features like job title, description, and company profile. Text data was preprocessed for BERT embeddings, which were then used with a Random Forest classifier on a 70:30 train-test split for fraud detection.

3.2 *Feature Extraction and Tokenization using BERT*

BERT (Bidirectional Encoder Representations from Transformers) [16], developed by Google, transformed NLP by capturing contextual meaning through the Transformer architecture [17] and attention mechanisms. Pre-trained on large corpora like Wikipedia, it can be fine-tuned for tasks such as classification and sentiment analysis. Using WordPiece tokenization [18], text is split into tokens, converted into IDs, and mapped to embeddings that represent semantic meaning. Unlike earlier models, BERT produces context-dependent embeddings [16],

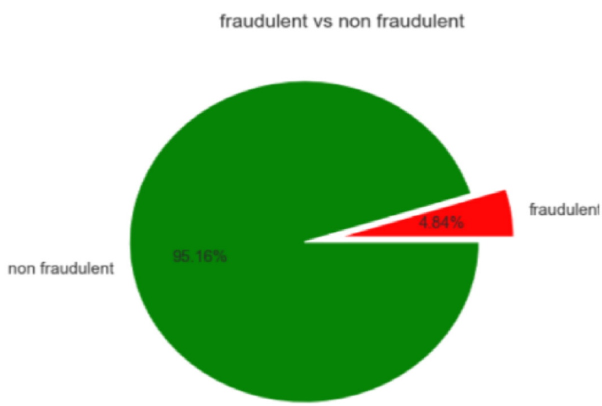

Figure 1

Amount of Real vs Fraud Jobs

where word meanings adapt to surrounding text for deeper understanding. Figure 1 shows the Amount of Real vs Fraud Jobs.

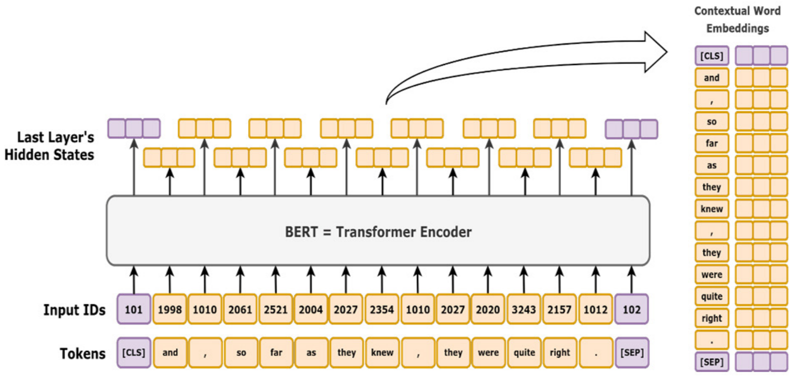

Figure 2

BERT Model Architecture in dataset

3.3 *Oversampling Techniques*

Because fraudulent postings form only a small portion of the dataset, oversampling is applied to reduce class imbalance and avoid bias toward legitimate jobs. Detecting the minority class is crucial, as missed fraud cases can have serious impacts. Techniques like SMOTE, Random Oversampling, and SMOTE-Tomek generate additional synthetic samples, helping balance the dataset, improve recognition of fraudulent entries, and increase model robustness, as shown in Figure 2.

3.3 *Classification and Hyperparameter Tuning*

Classification is a supervised learning approach where models assign predefined labels to data based on input features, aiming to predict binary or multi-class outcomes from labeled examples. Hyperparameter tuning, unlike parameter learning, involves adjusting settings defined before training that shape the model's structure and learning behavior. Random Forest is well-suited for this task, as its ensemble of diverse trees reduces overfitting and highlights key features, offering both accuracy and interpretability in detecting fraudulent job postings.

3.3 *Evaluation Metrics*

In this study, assessing the performance of the classifier is crucial to determining its efficacy in identifying fraudulent job postings. Several performance metrics are employed to evaluate how well the model distinguishes between legitimate and fraudulent postings like accuracy, Precision, Recall F1-score, AUC-ROC and Support as shown in Eq 1,2,3,4.

$$accuracy = \frac{true\ \ positives + true\ \ negatives}{total\ \ instances} \tag{1}$$

$$Precision = \frac{true\, positives}{true\, positives + false\, positives} \tag{2}$$

$$Recall = \frac{true\ \ positives}{true\ \ positives + false\ \ negatives} \tag{3}$$

$$F1 - score = 2 * \frac{precision * recall}{precision + recall} \tag{4}$$

The AUC reflects the model's effectiveness in ranking fraudulent postings above genuine ones, making it a reliable measure of performance on imbalanced datasets. Support represents the count of samples in each class, with Class 0 showing legitimate postings and Class 1 indicating fraudulent ones.

Table 1

Baseline Results of different Oversampling Techniques

Oversampling Technique	**Class**	**Precision**	**Recall**	**F1-Score**	**Support**
Random_Oversampler	0	0.9727	1	0.9861	5104
	1	1	0.45	0.6206	260
SMOTE	0	0.9749	0.9984	0.9865	5104
	1	0.9416	0.49615	0.6498	260
ADASYN	0	0.9748	0.9968	0.9857	5104
	1	0.8896	0.4961	0.637	260
SMOTE ENN	0	0.9799	0.9747	0.9773	5104
	1	0.5505	0.6076	0.5776	260
SMOTE Tomek	0	0.9755	0.9986	0.9869	5104
	1	0.9496	0.5076	0.6616	260

4. Results and Discussion

This section compares different models and techniques to evaluate their effectiveness, using baseline metrics as a benchmark. Tables 1 and 2 present results before and after hyperparameter tuning across oversampling methods, highlighting changes in accuracy and AUC-ROC.

Table 2

Hyperparameter Tuned results of different oversampling Techniques

Oversampling Technique	Class	Precision	Recall	F1-Score	Support
Random_Over sampler	0	0.9731	1	0.9863	5104
	1	1	0.4576	0.6279	260
SMOTE	0	0.9769	0.9958	0.9861	5104
	1	0.8695	0.5384	0.665	260
ADASYN	0	0.9838	0.9326	0.9575	5104
	1	0.346	0.7	0.4631	260
SMOTE ENN	0	0.9859	0.9353	0.9599	5104
	1	0.3678	0.7384	0.491	260
SMOTE Tomek	0	0.9765	0.9964	0.9864	5104
	1	0.8846	0.5307	0.6634	260

Table 3 highlights how oversampling with hyperparameter tuning affects performance. Random Oversampling reduced false negatives, SMOTE improved recall and F1 but lowered precision, while SMOTE-ENN raised recall at the cost of more false positives. SMOTE-Tomek offered the best balance, improving recall and keeping accuracy steady with only a slight precision.

Table 3

Overall Metrices

Baseline Overall Metrics			Tuned Overall Metrics	
Oversampling Technique	**Accuracy**	**AUC-ROC**	**Accuracy**	**AUC-ROC**
Random Over sampler	0.9733	0.725	0.9737	0.7288
SMOTE	0.974	0.7472	0.9737	0.7671
ADASYN	0.9725	0.7465	0.9213	0.8163
SMOTE ENN	0.9569	0.7912	0.9258	0.8369
SMOTE Tomek	0.9748	0.7531	0.9739	0.7636

5. Conclusion & Future Work

This study shows that combining BERT embeddings with a Random Forest classifier and oversampling improves fraud detection. Random Oversampling gave an F1 of ~0.62–0.63 for fraudulent cases, while SMOTE and ADASYN raised recall, with ADASYN trading precision for more positives. SMOTE Tomek achieved the best balance post-tuning (F1 = 0.66, accuracy = 97.48%), and SMOTE ENN reached the highest AUC-ROC (0.8369). These results outperform earlier works (93–99% accuracy with weaker AUC-ROC), proving the framework's effectiveness and scope for further enhancement.

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), doi: 10.3390/fi9010006.

[2] S. Dutta and S. K. Bandyopadhyay, "Fake job recruitment detection using machine learning approach," *Int. J. Eng. Trends Technol.*, vol. 68, no. 4, pp. 48–53 (2020), doi: 10.14445/22315381/ijett-v68i4p209s.

[3] B. Alghamdi and F. Alharby, "An intelligent model for online recruitment fraud detection," *J. Inf. Secur.*, vol. 10, no. 3, pp. 155–176 (2019), doi: 10.4236/jis.2019.103009.

[4] S. Lal, R. Jiaswal, N. Sardana, A. Verma, A. Kaur, and R. Mourya, "ORFDetector: Ensemble learning-based online recruitment fraud detection," in *Proc. 12th Int. Conf. Contemporary Computing (IC3)* , pp. 1–5 (2019), doi: 10.1109/IC3.2019.8844879.

[5] I. M. Nasser, A. H. Alzaanin, and A. Y. Maghari, "Online recruitment fraud detection using ANN," in *Proc. Palestinian Int. Conf. Inf. Commun. Technol. (PICICT)*, pp. 13–17 (2021), doi: 10.1109/PICICT53635.2021.00015.

[6] S. U. Habiba, M. K. Islam, and F. Tasnim, "A comparative study on fake job post prediction using different data mining techniques," in *Proc. 2nd Int. Conf. Robotics, Electrical and Signal Processing Techniques (ICREST)*, pp. 543–546 (2021), doi: 10.1109/ICREST51555.2021.9331230.

[7] C. Lokku, “Classification of genuinity in job posting using machine learning,” *Int. J. Res. Appl. Sci. Eng. Technol.*, vol. 9, no. 12, pp. 1569–1575 (2021), doi: 10.22214/ijraset.2021.39580.

[8] O. Nindyati and I. G. B. Nugraha, “Detecting scam in online job vacancy using behavioral features extraction,” in *Proc. Int. Conf. ICT for Smart Society (ICISS)* (2019), doi: 10.1109/ICISS48059.2019.8969842.

[9] G. O. Alandjani, “Online fake job advertisement recognition and classification using machine learning,” *3C TIC: Cuadernos de Desarrollo Aplicados a las TIC*, vol. 11, no. 1, pp. 251–267 (2022), doi: 10.17993/3ctic.2022.111.251-267.

[10] A. Gosain and S. Sardana, “Handling class imbalance problem using oversampling techniques: A review,” in *Proc. Int. Conf. Advances in Computing, Communications and Informatics (ICACCI)*, pp. 79–85 (2017), doi: 10.1109/ICACCI.2017.8125820.

[11] F. Akhbardeh, C. O. Alm, M. Zampieri, and T. Desell, “Handling extreme class imbalance in technical logbook datasets,” in *Proc. 59th Annu. Meeting Assoc. Comput. Linguistics (ACL)*, pp. 5740–5754 (2021), doi: 10.18653/v1/2021.acl-long.312.

[12] G. K. Yadav, A. K. Tiwari, V. K. Singh, A. Kumar, D. Varshney, and A. K. Senger, “Empirical algorithm to detect skin disease using deep convolution neural networks through mobilenet algorithm,” *Journal of Information and Optimization Sciences,* vol. 45, no. 4, pp. 1017–1027 (2024), doi: 10.47974/JIOS-1624.

[13] J. David, J. Cui, and F. Rahimi, “Classification of imbalanced datasets using BERT embeddings,” Dalhousie University (2020). [Online]. Available: https://fatemerhmi.github.io/files/Classification_of_imbalanced_dataset_using_BERT_embedding.pdf.

[14] Kanika, J. Singla, A. K. Bashir, Y. Nam, N. U. Hasan, and U. Tariq, “Handling class imbalance in online transaction fraud detection,” *Computers, Materials & Continua*, vol. 70, no. 2, pp. 2861–2877 (2022), doi: 10.32604/cmc.2022.019990.

[15] S. Bansal, "Real/fake job posting prediction," Kaggle Dataset (Feb. 2020), [Online]. Available: https://www.kaggle.com/datasets/shivamb/real-or-fake-fake-jobposting-prediction.

[16] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. NAACL-HLT*, pp. 4171–4186 (2019).

[17] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 5998–6008 (2017).

[18] D. W. Bliss, S. Kraut, and A. Agaskar, "Transmit and receive space-time-frequency adaptive processing for cooperative distributed MIMO communications," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5221–5224 (2012), doi: 10.1109/ICASSP.2012.6289097.

Received April, 2025