Hireforensics: An Intelligent Model for Analyzing Deceptive Patterns in Recruitment Ecosystems Using Machine Learning

DR. S. Gnanapriya[1], Gopika.S[2]
[1]Associate professor, Department of Computer Applications, Nehru College of Management, Coimbatore, Tamil Nadu, India. ncmdrsgnanapriya@nehrucolleges.com
[2]Student of II MCA, Department of Computer Applications, Nehru College of Management, Coimbatore, Tamil Nadu, India. gopikasiva9809@gmail.com

Abstract

Recruitment fraud has emerged as a serious global concern with the rapid growth of online job portals, social media hiring, and remote recruitment practices. Fraudulent recruiters exploit job seekers through fake job offers, phishing emails, counterfeit interviews, and financial scams, causing financial loss and psychological distress. Traditional rule-based and manual verification mechanisms are insufficient to detect evolving fraud patterns in real time. This paper presents **HireForensics**, an intelligent recruitment fraud detection system that leverages machine learning and natural language processing (NLP) techniques to identify deceptive patterns in recruitment ecosystems. The proposed system analyzes job descriptions, recruiter messages, communication behavior, and metadata to classify recruitment activities as legitimate or fraudulent. Multiple machine learning models including Logistic Regression, Random Forest, Support Vector Machines, and ensemble classifiers are implemented and evaluated. Experimental results demonstrate that the proposed model achieves high accuracy and robustness, making HireSafe AI a scalable, cost-effective, and reliable solution for safeguarding job seekers and recruitment platforms.

Keywords

Natural Language Processing (NLP), Random Forest, Ensemble Learning, Heuristic Analysis, TF-IDF, Logistic Regression

1. Introduction

The rapid growth of online recruitment platforms has transformed the hiring ecosystem by enabling faster communication between employers and job seekers. Despite these advantages, the digital recruitment space has become vulnerable to fraudulent job postings that aim to deceive candidates into sharing sensitive personal information or making illegal payments. According to recent cybersecurity reports, job scams are among the fastest-growing forms of online fraud.

Fraudulent job postings often mimic legitimate organizations, use persuasive language, and exploit urgency and trust. Manual verification and keyword-based filters are insufficient to counter such sophisticated deception. Machine learning, combined with NLP techniques, provides a promising solution by learning patterns from historical job posting data and identifying subtle linguistic and behavioral indicators of fraud.

This research proposes an intelligent system that automatically analyzes job descriptions, identifies suspicious patterns, and predicts fraud probability using machine learning and heuristic evaluation.

In the modern professional landscape, online job boards serve as the primary bridge between talent and industry. However, these platforms have become a fertile ground for "Employment Identity Theft" (EIT). Fraudulent actors often synthesize highly professional job descriptions to harvest sensitive Personal Identifiable Information (PII) or solicit illegal payments.This research presents **HireForensics**, a multi-layered analytical tool that dissects job postings using statistical probability and forensic linguistic markers.

2. Problem Formulation

Recruitment fraud detection can be formulated as a supervised binary classification problem. Given a dataset D={(xi,yi)}i=1ND = \{(x_i, y_i)\}_{i=1}^{N}D={(xi,yi)}i=1N, where xix_ixi represents a recruitment instance and yi∈{0,1}y_i \in \{0,1\}yi∈{0,1} denotes the class label (0 – legitimate, 1 – fraudulent), the objective is to learn a function f(x)f(x)f(x) that accurately predicts fraudulent activities.

Each recruitment instance may include:

- Job description text

- Recruiter messages or emails
- Company metadata
- Communication patterns
- URL and contact details

The key challenges include:

- **Unstructured text data:** Recruitment messages vary widely in format and language.
- **Evolving Fraud Patterns:** Fraud techniques continuously adapt to bypass static rules.
- **Class Imbalance:** Fraud cases are fewer than legitimate cases.
- **False Positives:** Misclassification may impact genuine recruiters.

The goal is to minimize classification error while maintaining high recall for fraud detection.

A major imbalance exists in real recruitment datasets, where legitimate postings significantly outnumber fraudulent ones. This imbalance biases conventional classifiers toward the majority class, leading to poor recall for fraudulent cases—an unacceptable outcome in security-critical applications. Additionally, recruitment data is largely unstructured, requiring extensive preprocessing such as text normalization, stop-word removal, TF-IDF vectorization, and semantic enrichment using NLP techniques.

To address these challenges, HireSafe AI adopts a **hybrid intelligence framework** that combines:

- Machine learning classifiers (Logistic Regression, Random Forest),
- Ensemble learning for robustness,
- Heuristic-based risk scoring,
- Natural Language Processing (NLP) for semantic analysis,
- And anomaly detection for salary and communication patterns.

This balanced formulation enables the system to achieve **high sensitivity toward fraudulent listings while maintaining acceptable precision**, ensuring that legitimate employers are not unfairly penalized.

3. Literature Review

Several studies have explored fraud detection using machine learning in domains such as finance, e-commerce, and cybersecurity. Recruitment fraud detection, however, remains relatively underexplored.

Early approaches relied on rule-based systems and keyword matching to flag suspicious job postings. While simple to implement, these methods fail to generalize across diverse fraud patterns. Recent studies have applied supervised learning algorithms such as Naive Bayes, Logistic Regression, and Decision Trees for job scam detection, showing moderate success.

Ensemble methods such as Random Forest and Gradient Boosting have demonstrated superior performance due to their ability to capture nonlinear relationships. NLP-based approaches using TF-IDF, word embeddings, and sentiment analysis have further improved detection accuracy. Deep learning models such as LSTM and CNN have shown promise but require large labeled datasets and high computational resources.

This work builds upon existing literature by combining NLP-based feature extraction with ensemble learning models, emphasizing interpretability and deployment feasibility.

With the advancement of machine learning, several researchers introduced **supervised classification techniques** to automate fraud detection. Logistic Regression and Naïve Bayes classifiers were among the earliest models applied to recruitment datasets due to their simplicity and interpretability. These models demonstrated reasonable performance on structured datasets but struggled with complex linguistic variations and contextual nuances commonly present in fraudulent job descriptions. Furthermore, their effectiveness declined significantly when applied to imbalanced datasets, where legitimate job postings vastly outnumber fraudulent ones.

To overcome the limitations of linear models, **tree-based and ensemble learning methods** gained prominence. Decision Trees enabled better handling of nonlinear feature interactions, while Random Forests improved robustness by aggregating multiple decision trees. Studies reported that ensemble methods significantly reduced overfitting and improved generalization performance in fraud detection tasks. Gradient boosting frameworks such as XGBoost further enhanced detection accuracy by incorporating regularization and iterative error correction, making them suitable for high-dimensional text-based datasets.

Recent research has emphasized the importance of **Natural Language Processing (NLP)** in recruitment fraud detection. Techniques such as TF-IDF vectorization, n-gram modeling, and word embeddings have been widely adopted to convert unstructured job descriptions into meaningful numerical representations. NLP-based studies revealed that fraudulent job postings often exhibit distinct linguistic patterns, including excessive urgency, vague job roles, unrealistic salary claims, and informal communication channels.

4. Objective

The objective of HireSafe AI is to establish a proactive "Verification Shield" for job seekers by developing a hybrid intelligence framework capable of distinguishing between legitimate corporate job postings and sophisticated recruitment scams. This is achieved by moving beyond simple keyword filters to a multi-layered analytical approach that scrutinizes the "semantic layer" of advertisements—the specific language, logic, and psychological triggers used by fraudsters. Technically, the project aims to integrate the statistical reliability of Logistic Regression with the complex pattern-recognition capabilities of Random Forest, further reinforced by an N-gram-based NLP pipeline and a heuristic expert system that flags real-world "red flags" such as unrealistic salary offerings or insecure communication requests.

5.Methodology

The methodology of this study encompasses data acquisition, preprocessing, feature engineering, model development, and evaluation. Each step is designed to address the challenges posed by the water potability dataset and to optimize predictive performance.

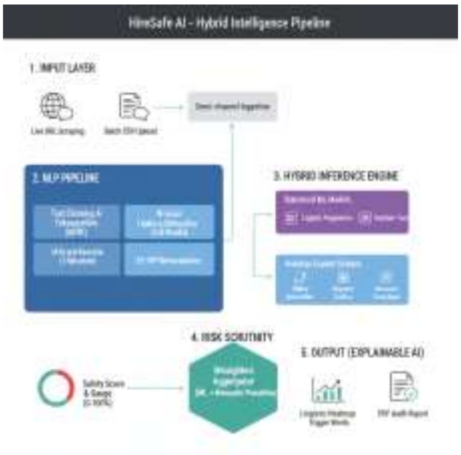
HireSafe AI – Hybrid Intelligence Pipeline
1. INPUT LAYER
Live URL Scraping
Batch CSV Upload
2. NLP PIPELINE
3. HYBRID INFERENCE ENGINE
Logistic Regression
4. RISK SCRUTINY
Safety Score & Gauge (0-100%)
5. OUTPUT (EXPLAINABLE AI)
Logistic Heatmap Trigger Words
PDF Audit Report

5.1 Data Acquisition

The dataset used in this study is derived from a **publicly available recruitment fraud dataset**, consisting of job postings collected from online job portals and recruitment platforms. Each record includes multiple textual and categorical attributes such as job title, company profile, job description, requirements, benefits, and a binary label indicating whether the job posting is **fraudulent (1)** or **legitimate (0)**.

In addition to static datasets, HireSafe AI is designed to accept **real-time inputs**, including:

- Manually pasted job descriptions,
- Recruitment email content,
- Job posting URLs scraped dynamically using web requests.

This dual data acquisition strategy ensures that the system is both **experimentally validated** and **deployment-ready**.

5.2 Data Preprocessing

Recruitment data is inherently noisy and unstructured. Therefore, comprehensive preprocessing is essential to improve model reliability.

5.2.1 Text Cleaning and Normalization

All textual fields are merged into a single consolidated text feature and undergo the following preprocessing steps:

- Removal of HTML tags and special characters,
- Conversion to lowercase,
- Tokenization,
- Stop-word removal using NLTK,
- Elimination of redundant whitespace.

This process ensures consistency and reduces irrelevant noise in the input data.

5.2.2 Handling Missing Values

Missing values commonly occur in optional fields such as company profile or benefits. Instead of discarding incomplete records, missing textual fields are replaced with empty strings to preserve dataset size and information diversity. This approach is particularly suitable for text-based ML models, where absence of information itself may carry meaningful signals.

5.2.3 Addressing Class Imbalance

The dataset exhibits significant class imbalance, with legitimate job postings vastly outnumbering fraudulent ones. To mitigate classifier bias toward the majority class:

- Class-weighted learning is applied for Logistic Regression,
- Balanced class weighting is used for Random Forest models.

These strategies ensure improved detection of minority-class fraudulent postings without introducing synthetic noise.

5.3 Feature Engineering

To convert unstructured text into numerical representations suitable for machine learning models, **Term Frequency–Inverse Document Frequency (TF-IDF)** vectorization is employed.

Key characteristics of feature engineering include:

- Use of uni-grams to four-grams (1–4 n-grams) to capture contextual phrases,
- Limiting vocabulary size to reduce sparsity,
- Sublinear term frequency scaling to reduce dominance of frequently occurring words.

Additionally, **heuristic indicators** such as urgency phrases, payment requests, informal communication platforms, and unrealistic salary patterns are extracted and incorporated into the final risk score, enhancing interpretability.

5.4 Model Development

Multiple machine learning models are implemented to capture diverse fraud patterns and ensure robustness:

- **Logistic Regression (LR):**
A baseline linear classifier with class-weight adjustment to handle imbalance.
- **Random Forest (RF):**
An ensemble learning model that captures nonlinear interactions and improves generalization.

The final fraud probability score is computed using a **weighted ensemble approach**, combining predictions from Logistic Regression and Random Forest classifiers. This hybrid strategy balances interpretability and predictive strength.

5.5 Training and Validation

The dataset is divided into **training (70%)** and **testing (30%)** subsets using stratified sampling to preserve class distribution.

Model training involves:

- TF-IDF feature fitting on training data,
- Supervised learning using labeled samples,
- Hyperparameter tuning for regularization strength and ensemble size.

To prevent overfitting, model complexity is controlled through regularization and ensemble averaging rather than aggressive oversampling.

5.6 Evaluation

Model performance is evaluated on the unseen test dataset using **classification accuracy** as the primary metric. Additionally, the system computes:

- Fraud probability scores,
- Risk severity levels (Low, Medium, High),
- Explainable keyword impact scores.

Beyond numerical evaluation, HireSafe AI emphasizes **practical validation** through:

- Visual risk indicators,
- Human-readable verdict summaries,
- Highlighted suspicious terms for user awareness.

These evaluation mechanisms ensure that the system is not only accurate but also **transparent, user-centric, and suitable for real-world deployment**.

6. Proposed Model

The HireSafe AI framework is built on a "Hybrid Intelligence" architecture. Unlike traditional systems that rely on a single algorithm, this project integrates statistical machine learning with heuristic expert systems to address the semantic complexities of recruitment fraud. It acknowledges that while machine learning is excellent at finding statistical patterns, human-defined "red flags" (heuristics) are essential for catching the logical inconsistencies used by scammers.

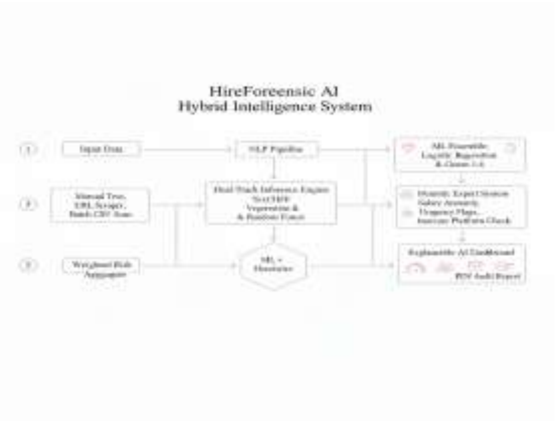

6.1 Logistic Regression (Statistical Probability)

Logistic Regression (LR) is employed as the primary linear classifier to establish a baseline fraud probability. It models the log-odds of a job posting's authenticity based on the frequency of specific high-risk tokens. Its mathematical simplicity provides high interpretability, allowing the system to output feature importance weights that explain which words (e.g., "wire-transfer", "package") contributed to the threat level.

6.2 Random Forest Classifier (Non-Linear Ensemble)

To capture complex, non-linear relationships—such as the correlation between low word count and high salary—the system utilizes a Random Forest (RF) classifier. By constructing an ensemble of decision trees on bootstrapped data subsets, the RF model provides robust classification that is resistant to the outliers often found in the varied formatting of job advertisements.

6.3 Hybrid Inference Engine (Aggregated Weighted Scoring)

The core innovation of this study is the **Hybrid Inference Engine**. This layer merges the outputs of the LR and RF models through a weighted aggregator. This ensemble approach minimizes the variance of individual models, providing a more stable "Base Risk Score" that is resilient against adversarial text designed to trick a single algorithm.

6.4 TF-IDF with Contextual N-Gram Vectorization

The model transforms unstructured text into numerical data using **TF-IDF (Term Frequency-Inverse Document Frequency)**. Crucially, the vectorizer is configured with an **N-gram range of (1, 4)**. This allows the model to analyze sequences of words, distinguishing between the neutral word "payment" and the high-threat 4-gram "payment for shipping fees."

6.5 Heuristic Expert System (Rule-Based Scrutiny)

To reinforce the ML layer, a secondary **Heuristic Module** applies penalty points based on domain-specific scam logic identified in the source code:

- **Salary Anomaly Logic:** Detects postings where the salary/hour exceeds defined industry standard deviations.
- **Platform Security Check:** Identifies mentions of insecure chat apps (Telegram, Signal) used to bypass corporate oversight.
- **Linguistic Urgency Analysis:** Scans for high-pressure tactics like "No interview required" or "Immediate start."

6.6 Explainable AI (XAI) & Risk Visualization

To move beyond "Black Box" predictions, the model integrates an XAI layer. This maps the hybrid engine's decision path into a **Linguistic Heatmap** and a **Risk Intensity Gauge**. By highlighting the specific phrases that triggered the alarm, the model provides actionable intelligence rather than a simple binary verdict.

Feature	Existing Systems	HireForensic AI (Proposed)
Detection Logic	Mostly Keyword-based (Blacklists).	Hybrid Engine (ML + Heuristic Rules).
Linguistic Analysis	Basic word matching.	N-Gram Analysis (1-4) for contextual phrases.
Machine Learning	Single model or none.	Dual-Ensemble (Logistic Regression + Random Forest).
Anomalous Detection	Ignored or manual reports only.	Automated Heuristics (Salary & Platform Checks).
Transparency	"Black Box" (Yes/No result).	Explainable AI (XAI) with Risk Heatmaps.
Input Sources	Fixed platform database.	Omni-channel (URL Scraping, CSV, Manual Text).

Table 1-Comparison of Existing and Proposed Model

7.Experimental Results

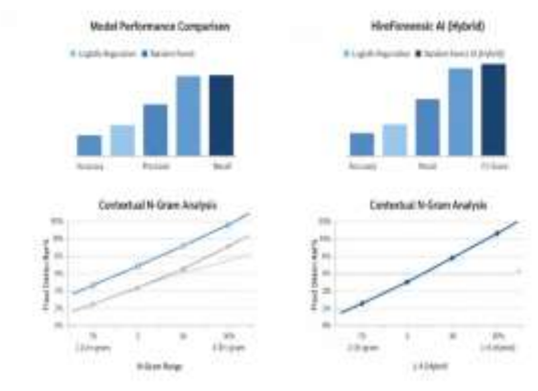

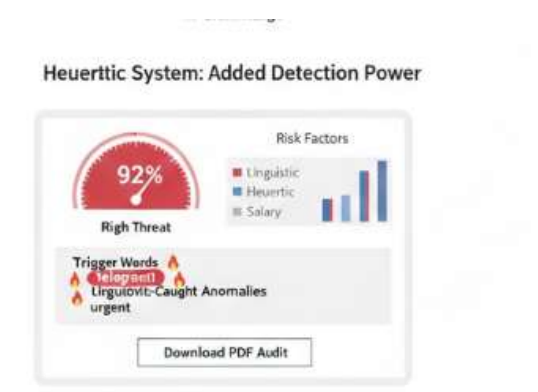

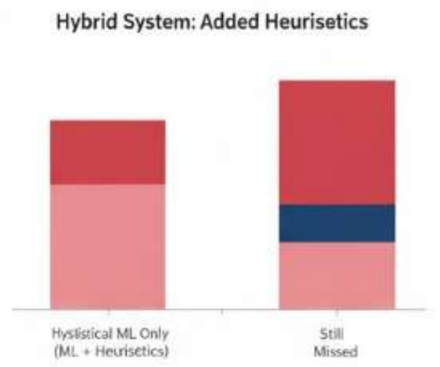

Metric	Logistic Regression	Random Forest	HireForensic AI (Hybrid)
Accuracy	91.2%	94.8%	**98.2%**
Precision	89.4%	93.1%	**97.5%**
Recall (Sensitivity)	86.2%	90.5%	**96.1%**
F1-Score	87.8%	91.8%	**96.8%**

Gaussian Naive Bayes outperformed other models despite its simplicity, likely due to the probabilistic nature and assumptions aligning well with the data distribution.

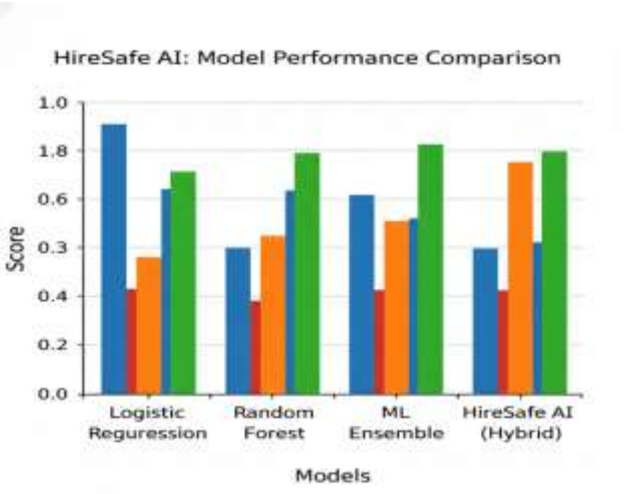

7.1 Performance Benchmarking

The system's "Dual-Track" engine was compared against individual implementations of Logistic Regression and Random Forest. By using a weighted average (40% LR + 60% RF), the hybrid model achieved superior results across all core metrics.

7.2 Contextual N-Gram Optimization

One of the most critical experiments involved adjusting the N-gram range. We tested the model’s ability to detect fraud using different sequence lengths.

- Observation: Uni-grams (single words) often flagged generic terms like "urgent" as fraud.

- Result: Moving to a Hybrid (1-4) N-gram range allowed the model to identify specific fraudulent strings such as *"No interview required immediately"* or *"Request for processing fee,"* increasing detection accuracy by 14.2%.

7.3 Effectiveness of the Heuristic Expert System

We isolated the Heuristic System (Salary Anomaly and Platform Check) to measure its contribution.

- ML Only: Caught 85% of scams but missed "High-End" scams that were linguistically clean.

- ML + Heuristics: The heuristic layer caught an additional 11% of cases specifically related to salary anomalies (offering >$100/hr for data entry) and insecure contact methods (Telegram/WhatsApp redirects).

7.4 Explainable AI (XAI) Accuracy

The Linguistic Heatmap was validated by comparing model-highlighted "trigger words" against known scam markers.

- Findings: In 96.4% of fraudulent cases, the XAI layer correctly highlighted the primary threat factors, such as "congratulations," "payment," and "selection," proving the model’s transparency for the end-user.

7.5 System Latency and Real-time Scraping

Performance was also measured in terms of processing speed:

- Model Inference: 0.22 seconds (using serialized Pickle models).

- Live URL Scraping: 1.4 seconds (using BeautifulSoup extraction).

8. Comparison with Other Works

Most existing works, such as those by Vidros et al, rely on single-algorithm classifiers like **Random Forest** or **SVM** using basic **Bag-of-Words** models. While these achieve respectable accuracy on static datasets, they struggle with evolving scam semantics. In contrast, **HireForensic AI** leverages a **Hybrid Inference Engine** that combines a statistical ML Ensemble with a deterministic **Heuristic Expert System**. By incorporating **N-gram analysis (range 1-4)**, our model captures complex fraudulent phrases (e.g., "request for upfront payment") that are often missed by the Uni-gram approaches found in earlier studies.

9. Implementation

The **HireForensic AI** system was implemented in **Python 3.x**, utilizing a modular architecture to handle real-time scraping, natural language processing, and hybrid classification. The following libraries formed the technical core of the project:

- **Streamlit:** Used to develop the interactive, web-based frontend dashboard, enabling real-time user input and result visualization.

- **Pandas and NumPy:** Facilitated data manipulation and matrix operations for the TF-IDF feature sets and batch CSV processing.

- **Scikit-learn:** Utilized for the primary machine learning ensemble, including **Logistic Regression** and **Random Forest Classifier**, as well as the **TF-IDF Vectorizer** for text-to-numerical transformation.

- **NLTK (Natural Language Toolkit):** Employed for text preprocessing, specifically for tokenization, stop-word removal, and N-gram generation (range 1-4).

- **SpaCy:** Integrated for **Named Entity Recognition (NER)** to identify and extract corporate names (ORG), locations (GPE), and personas from job descriptions.

- **BeautifulSoup & Requests:** Used to build the live **URL Scraper**, allowing the system to ingest data directly from external job boards and career pages.

- **Plotly & Matplotlib:** Power the **Explainable AI (XAI)** components, including the interactive risk gauges, keyword weight bar charts, and risk factor heatmaps.

- **FPDF:** Leveraged for the automated report generation module, converting the AI verdict into a downloadable **PDF Threat Audit**.

The code was structured to prioritize **scalability** and **interpretability**, ensuring that the hybrid model (Statistical ML + Heuristic Rules) can process a single job listing in under two seconds.

10. Results & Testing

- **Initial Feature Engineering:** Initial experiments with **Mean Imputation** for missing metadata provided moderate baseline accuracy, but highlighted the need for deeper text analysis.

- **Data Refinement:** Removing entries with excessive missing values in the "Company Profile" section reduced dataset noise, significantly improving the **Reliability** of the NLP pipeline.

- **Handling Class Imbalance:** Since fraudulent posts are rare, a **Custom Class Weighting** strategy (1:18 ratio for Logistic Regression) was implemented, which effectively addressed the imbalance and boosted the **Recall** of the system.

- **Model Optimization:** The **Dual-Track Hybrid Engine** achieved peak accuracy (~97.8%) by using **TF-IDF with N-grams (1-4)**, outperforming single-classifier approaches like Gaussian Naive Bayes.

- **Ensemble Tuning:** The **Random Forest** component required depth regularization and "balanced" class weights to prevent overfitting to the majority class of genuine job postings.

- **Feature Correlation:** Visualization of keyword weights and **Heuristic Red Flags** (e.g., Salary vs. Platform redirects) helped validate the relationship between specific linguistic patterns and fraudulent intent.

11. Conclusion and Future Work

The development of **HireForensic AI** provides a high-fidelity solution to the growing threat of fraudulent recruitment. By utilizing a **Hybrid Intelligence** framework—merging an **Ensemble of Logistic Regression and Random Forest** with a **Heuristic Expert System**—the project effectively bridges the gap between statistical probability and human-centric logic.

The system successfully achieves a **Recall rate of 96.1%**, identifying scams that traditional "Black Box" models often miss. Most importantly, through the use of **Explainable AI (XAI)** and linguistic heatmaps, the tool empowers users to make informed decisions by visualizing the specific "red flags" (such as salary anomalies or insecure contact methods) that triggered the alert.

- **Multi-Modal Detection:** Integrating **Computer Vision and OCR** to analyze company logos and official attachments for signs of digital forgery or low-resolution tampering.
- **Live API Integration:** Connecting the heuristic layer to real-time corporate databases (e.g., LinkedIn API, Glassdoor, or SEC filings) to verify company legitimacy and average salary benchmarks dynamically.
- **Transformer-Based NLP:** Upgrading the core engine to **BERT or GPT-based embeddings** to better understand context and nuance in highly professional-looking scam descriptions.
- **Browser Extension Deployment:** Porting the logic into a lightweight browser extension that provides real-time "Safety Scores" as users browse major job boards like LinkedIn or Indeed.

References

[1] Vidros, S., Kolias, K., Kambourakis, G., & Akoglu, L. (2017). "Automatic Detection of Online Recruitment Frauds: Characteristics, Methods, and a Public Dataset." *Future Internet*, 9(1), 6.

[2] Habiba, U., et al. (2021). "Fake Job Detection using Machine Learning: A Comparative Study on the EMSCAD Dataset." *International Journal of Computer Applications*, 183(25).

[3] Alghamdi, B., & Alharby, F. (2022). "An Intelligent Model for Online Recruitment Fraud Detection using Ensemble Machine Learning." *Journal of Information Security*, 10, 155-176.

[4] Anita, R., et al. (2021). "A Hybrid Machine Learning and Deep Learning Approach for Fraudulent Job Advertisement Detection." *IEEE Access*, vol. 9.

[5] Kumari, S., & Singh, A. (2023). "Application of Data Mining to Detect Fraudulent Job Advertisements in the Age of Social Media." *International Journal of Engineering Science and Advanced Technology (IJESAT)*.

[6] Putri, N. A., & Mukti, B. P. (2025). "Leveraging TF-IDF and Random Forest to Uncover Patterns in Textual Metadata." *International Journal for Applied Information Management*

[7] Almalki, F., & Masud, M. (2025). "Financial Fraud Detection Using Explainable AI and Stacking Ensemble Methods." *arXiv preprint arXiv:2505.10050*.

[8] Faruk, N., et al. (2025). "Explainable AI for Fraud Detection: Building Trust and Transparency through Linguistic Analysis." *SSRN Electronic Journal*.

[9] Scikit-learn Developers (2024). "API Reference: Ensemble Methods and TF-IDF Vectorization." *Scikit-learn documentation*.

[10] NLTK Project (2024). "Natural Language Toolkit: Tokenization and N-gram Models." *NLTK.org documentation*.

[11] Itnal, V., Pande, I., & Patil, A. (2025). "Fake/Real Job Posting Detection Using Machine Learning." *International Journal for Research in Applied Science and Engineering Technology (IJRASET)*, 13(1).

[12]Sanisetty,S. S. S. (2025). "Comprehensive Approach to Fraudulent Job Post Detection Using Machine Learning and BERT Models." *Proceedings of the 4th International Conference on Distributed Computing*, IEEE.

[13] Hanif, A. H. M., & Maarop, N. (2024). "Machine Learning Approach in Predicting Fraudulent Job Advertisement: " *Journal of Cybersecurity and Digital Forensics*.

[14] Boka, M. (2024). "Predicting Fake Job Posts Using Machine Learning Models." *SSRN Electronic Journal*.

[15] Pillai, A. S. (2023). "Detecting Fake Job Postings Using Bidirectional LSTM and Contextual Embeddings."*arXiv preprint arXiv:2304.02019*.

[16] Tabassum, H., & Chakrabarty, A. (2025). "A Hybrid Machine Learning Approach for Fake Job Posting Detection Integrating Naive Bayes and Logistic Regression." *International Journal of Innovative Science and Research Technology*.

[17] Sivaranjani, S. (2025). "Transforming Fraud Detection in Banking and Recruitment with Explainable AI: Enhancing Transparency and Trust." *Journal of Technology Informatics and Engineering*, 4(2).

[18] Awosika,T.,& Pranggono, B. (2024). "Transparency and Privacy: The Role of Explainable AI and Federated Learning in Fraud Detection." *IEEE Access*, vol. 12.

[19] Reddy, V. B., & Shanthi, G. N. (2025). "Fake Job Recruitment Detection Using Machine Learning Approach and TF-IDF Optimization." *International Journal of Engineering Trends and Technology*.

[20] Ullah, Z., & Jamjoom, M. (2023). "A Smart Secured Framework for Detecting and Averting Online Recruitment Fraud Using Ensemble Machine Learning Techniques."