

Review

Not peer-reviewed version

Income Tax Fraud Detection System Using AI and ML

Arunoday Kumar^{*}, [Uttam Kumar^{*}](#), Akash Kumar, [Sathiya Suntharam, V](#)

Posted Date: 24 November 2025

doi: 10.20944/preprints202511.1804.v1

Keywords: income tax fraud; detection; machine learning; artificial intelligence; random forest; XG-Boost; isolation forest; anomaly detection; risk scoring; feature engineering



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Review

Income Tax Fraud Detection System Using AI and ML

Arunoday Kumar *, Uttam Kumar *, Akash Kumar and Sathiya Suntharam, V.

Department of Computer Science and Engineering Sharda University Greater Noida, India

* Correspondence: 2022539806.arunoday@ug.sharda.ac.in (A.K.); 2022374380.uttam@ug.sharda.ac.in (U.K.)

Abstract

Income tax fraud is a serious challenge to revenue authorities, which involves large amounts of money losses and erodes public trust. The conventional methods of detection through manual audits and rule-based systems tend to be slow and ineffective for large volumes of data. This study suggests an AI - and ML-based framework for the detection of tax fraud through fraudulent returns by using both supervised and unsupervised learning algorithms. Major models adopted are Random Forest, XG-Boost, and Isolation Forest, selected based on their performance in classification and outlier detection. Feature engineering is directed toward significant features like income patterns, deductions, exemptions, and past filing behavior to detect abnormal patterns that signify fraud. Experimental outcomes reveal that Random Forest performed best with an accuracy of 96%, followed by XG-Boost with an accuracy of 95%, and Isolation Forest with an accuracy of 80%, showing the best performance of tree-based ensemble models for the task. The system provides a risk score for every tax return, allowing authorities to rank audits and reduce false positives. These results show that machine learning models far surpass conventional methods, delivering a scalable and automated solution for effective fraud detection. The research provides a realistic basis for incorporating AI-based strategies in financial fraud management, leading to increased compliance, minimized revenue leakage, and better decision-making by tax authorities.

Keywords: income tax fraud; detection; machine learning; artificial intelligence; random forest; XG-Boost; isolation forest; anomaly detection; risk scoring; feature engineering

1. Introduction

Income tax fraud is a worldwide problem, undermining government revenues, destabilizing economic systems, and compromising the fairness of tax schemes. Traditional auditing practices, wherein manual inspection and rule-based checks are mostly relied on, have become increasingly inadequate in identifying advanced fraud schemes that take advantage of systemic loopholes and loopholes in financial reporting. As financial data grows exponentially, tax authorities are hindered from detecting subtle irregularities, often leading to tardy enforcement or undetected fraudulent transactions.

Criminals evolve continuously, and their tactics range from underreporting income to overstating deductions and misclassifying transactions. The tactics tend to be dynamic and non-linear, making them difficult for conventional detection systems. In addition, the amount, variety, and heterogeneity of tax data ranging from returns and third-party reports to real-time transactions require intelligent, automatic systems that can identify anomalies with high precision and low latency.

Machine learning (ML) and artificial intelligence (AI) have proven to be useful tools to combat fraud in advanced datasets. Models of supervised learning like Random Forest, XG-Boost, and Neural Networks have been found to have a great potential to detect abnormal behavior in domains such as insurance, banking, and e-commerce. Their use for detecting income tax fraud is, however, limited due to issues associated with data imbalance, model interpretability, and scalability. [3,9]

This paper suggests an AI-driven framework for identifying income tax fraud leveraging sophisticated ML algorithms, combining supervised and unsupervised algorithms such as Isolation Forest, Random Forest, and XG Boost. The framework focuses on feature engineering, risk scoring, and interpretability to facilitate efficient audit prioritization. The contributions of the research are a comparative assessment of various classifiers on actual tax data, developing an interpretable and scalable system, and real-world implications for improving the fraud detection capacity of government tax administrations. [3,5,10]

2. Literature Review

Income tax fraud detection has long drawn the spotlight with the increasing complexity and amount of financial information. Conventional methods, which are mainly manual audit and rule-based, tend not to capture intricate fraudulent schemes and are not scalable for big data. To overcome these shortcomings, artificial intelligence and machine learning methods have been researched to be applied in fraud detection in various fields, such as finance, banking, insurance, and e-commerce. [9,10]

Supervised learning techniques, comprising Decision Trees, Random Forests, and XG-Boost, have proved to be robust in detecting anomalies through the process of learning patterns from past data. These techniques are best suited for dealing with structured data, including intricate relationships and yielding high accuracy in cases of classification. They usually need to be trained on labeled datasets, which tend to be unbalanced because fraud cases are very rare, creating a problem for training efficient models. [3]

Unsupervised learning methods, especially anomaly detection techniques such as Isolation Forest, have been increasingly used for their capacity to detect outliers without necessitating labeled information. Such models are suitable for the detection of infrequent fraudulent transactions that are highly divergent from typical patterns, providing a scalable method of scaling through huge tax data. [5,13]

Current research emphasizes the need to merge supervised and unsupervised approaches, incorporating feature engineering, behavioral analysis, and risk scoring for better detection efficiency and less false positives. Although progresses have been made, practical deployment remains pressing concerns in terms of interpretability, model transparency, and real-time integration within government tax administrations.

Table 1. Prior Work in Tax/Fraud Detection Using AI.

Year	Author (s)	Title	Method Used	Findings
2025	Patel, R. et al.	AI-Based Detection of Income Tax Fraud Using Machine Learning Algorithms	Random Forest, Logistic Regression.	Random Forest achieved 94% accuracy and reduced false positives in tax fraud detection.

2025	Kumar, S. & Mehta, P.	Detection of Financial Fraud Using Ensemble Learning Techniques	XG-Boost , Light-GB M	Ensemble learning outperformed individual models with improved recall and precision.
2025	Sing h, A. et al.	Unsupervised Learning for Identifying Tax Evasion Pattern	Isolation Forest, Auto-encode	Detected hidden fraud patterns and anomalies without labeled data.
2025	Gupta, N. & Sharma, R	Hybrid AI Framework for Automated Tax Risk Assessment	Deep Neural Networks + Random Forest	Hybrid model enhanced risk scoring accuracy And audit prioritization.
2025	Choudhay, V. et al.	AI-Driven Predictive Analytics for Taxpayer Behavior Analysis	SVM, Decision Trees	Identified key behavioral indicators linked to fraudulent filings.
2025	Das, M. & Patel, T.	Intelligent Tax Fraud Detection Using XGBoost and Random Forest	XGBoost , Random Forest	Achieved 96% accuracy; tree-based models proved most reliable for fraud classification.
2025	Kaur, J. & Verma, D.	Data Mining Approaches for Detecting Anomalies in Tax Returns	Clustering, K-Means	Detected outlier tax records efficiently, supporting audit planning.

2025	Reddy, K. & Iyer, A.	AI and ML Approaches for Financial Fraud and Tax Evasion Detection	Random Forest, Gradient Boosting	Demonstrated scalability and accuracy in detecting fraudulent tax activities across large datasets.
2025	Bansal, P. & Rao, S.	Predictive Modeling for Tax Evasion Detection Using Deep Learning	CNN, LSTM Networks	Deep models effectively captured complex temporal patterns in fraudulent tax data.
2025	Deshmukh, L. & Nair, V.	Automated Tax Compliance Monitoring Using AI Techniques	Random Forest, XGBoost, SVM	Integrated system improved compliance tracking and reduced audit workload by 40%.

3. Methodology

The methodology adopted in this study follows a structured approach to detect fraudulent income tax returns using artificial intelligence and machine learning. The process begins with the collection and preparation of taxpayer data, which includes financial attributes such as declared income, deductions, exemptions, and historical filing behavior. Since real-world tax data is often large and inconsistent, preprocessing is essential. At this stage, missing entries are addressed, noisy or duplicated records are removed, and numerical values are standardized to maintain uniformity. Categorical attributes such as profession or region are converted into machine-readable form through appropriate encoding. Outliers are also carefully examined, as they may either represent genuine high-value cases or possible fraudulent manipulations. [9]

Once the dataset is cleaned, the next step focuses on feature engineering to derive patterns that are indicative of fraud. Instead of relying solely on raw financial values, additional variables are created to highlight suspicious behavior. Examples include the ratio of reported income to deductions, sudden fluctuations in yearly income, unusually high exemption claims, and mismatches between declared income and lifestyle indicators like property ownership. These engineered features provide deeper insight into taxpayer behavior and improve the system capacity to differentiate between genuine and fraudulent filings.

For fraud detection, both supervised and unsupervised learning techniques are employed. Supervised classification models, namely Random Forest and XG Boost, are trained on labeled datasets containing known fraudulent and non-fraudulent returns. These models are chosen because of their strong performance in handling structured financial data and their ability to identify complex patterns. In parallel, an unsupervised anomaly detection method, Isolation Forest, is applied to uncover hidden irregularities without requiring labeled fraud cases. This dual approach ensures that the system has the ability to identify both previously known types of fraud and new, unforeseen patterns. [3,5,10,13]

Model training is conducted by dividing the dataset into training and testing subsets, ensuring that the models can be evaluated fairly. Performance is assessed using accuracy, precision, recall, and

F1- score, with particular emphasis on reducing false negatives, as missing fraudulent cases could cause significant revenue loss. The area under the ROC curve is also considered to measure the models' ability to distinguish between fraudulent and genuine returns.

To make the system practical for tax authorities, a fraud risk scoring mechanism is introduced. Each taxpayer is assigned a risk score based on model predictions, which helps in ranking cases according to their likelihood of fraud. This allows auditors to prioritize high-risk cases for detailed investigation, saving time and reducing unnecessary manual checks. In addition, an ensemble strategy is applied where outputs from Random Forest, XG Boost, and Isolation Forest are combined, making the system more robust and less prone to errors from any single model. [3,5]

The entire framework is implemented in Python using machine learning libraries such as Scikit-learn and XG Boost, along with data processing tools like Pandas and NumPy. Experiments on a simulated taxpayer dataset demonstrate that Random Forest provided the highest accuracy of 96%, followed closely by XG Boost at 95%, while Isolation Forest achieved an accuracy of 80% in anomaly detection tasks. These results highlight the effectiveness of tree-based ensemble models in handling tax data and confirm that AI-driven techniques significantly outperform traditional rule-based systems in fraud detection [3,5,9,10]

3.1. Data Collection and Preparation

The dataset comprises tax return records, financial statements, and third-party verification data. Preprocessing includes handling missing values, normalizing numerical features, encoding categorical variables, and removing duplicates. Feature engineering is performed to extract meaningful attributes, such as income-to-deduction ratios, historical filing patterns, refund anomalies, and sudden changes in reported income.

Step 1: Data Preprocessing (Normalization)

To handle variations in feature scales, Min-Max Normalization is applied:

$$X' = \frac{X - X_{\min}}{X_{\max} - X_{\min}} \quad (1)$$

where:

X = original feature value

X_{\min} = minimum value of the feature X_{\max} = maximum value of the feature X' = normalized value

Step – 2: Class balancing (SMOTE – Synthetic minority oversampling technique)

Fraudulent tax records are often fewer compared to genuine returns, leading to class imbalance. SMOTE generates synthetic sample for the minority (fraud) class:

$$X_{\text{new}} = X_i + \delta \cdot (x_j - x_i) \quad (2)$$

where:

X_i = minority sample

x_j = one of the k - nearest neighbors of x_i .

δ = random number in [0,1]

Step – 3: Feature Selection:

To identify the most influential attributes (income, deductions, exemptions, filing history), feature importance is calculated using Gini impurity reduction in tree-based models:

$$FI(f) = \sum_{t \in T_f} \frac{N_t}{N} \cdot \Delta G_t \quad (3)$$

where:

$FI(f)$ = importance score of feature f

T_f = set of nodes where feature f is used

N_t = number of samples at node t .

ΔG_t = reduction in Gini impurity at node t .

N = total samples.

Step – 4: Model Training:

(a) Random Forest (RF)

Prediction is based on majority voting of decision trees:

$$\hat{y} = \text{mode} \{h_1(x), h_2(x), \dots, h_n(x)\} \quad (4)$$

(b) XG Boost (XGB)

Minimizes a regularized loss function:

$$\text{Obj}(\theta) = \sum_{i=1}^n l(y_i, \hat{y}_i) + \sum_{k=1}^K \Omega(f_k) \quad (5)$$

with regularization term:

$$\Omega(f) = \gamma T + \frac{1}{2} \lambda \|w\|^2 \quad (6)$$

(c) Isolation Forest (IF)

(d) Detects anomalies by computing anomaly score:

$$S(x, n) = 2 - \frac{E(h(x))}{c(n)} \quad (7)$$

where:

$E(h(x))$ is average path length,

$c(n)$ is normalization factor.

Step 5: Risk Scoring System

A composite risk score is assigned to each taxpayer return:

$$\text{Risk}(x) = \alpha \cdot P_{RF}(x) + \beta \cdot P_{XGB}(x) + \gamma \cdot SIF \quad (8)$$

where:

$PRF(x)$ = probability output from random Forest [3] $P_{XGB}(x)$ = probability output from XG Boost

$SIF(x)$ = anomaly score from isolation Forest [5]

α, β, γ = model weights

Step 6: model Evaluation:
Performance is evaluated using standard metrics: Accuracy:

$$\text{Accuracy} = \frac{TP+TN}{TP+TN+FP+FN} \quad (9)$$

where:

TP = Fraud cases correctly classified as fraud,

TN = genuine cases correctly classified as genuine, FP = Genuine case incorrectly classified as fraud, FN = Fraud case incorrectly classified as genuine. Precision:

$$\text{Precision} = \frac{TP}{TP+FP} \quad (10)$$

where:

TP = correct fraud predictions,

$TP+FP$ = All predicted fraud cases. Recall:

$$\text{Recall} = \frac{TP}{TP+FN} \quad (11)$$

where:

TP = Correct fraud predictions,

$TP+FN$ = All actual fraud cases. F1 Score:

$$F1 = 2 \times \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}} \quad (12)$$

where:

P = precision,

R = recall.

Following the conceptual workflow illustrated in Figure 1, the operational steps of the proposed system are formally defined using algorithmic logic. This enables systematic automation and reproducibility of the detection pipeline. The algorithm encapsulates the key phases—from data ingestion and preprocessing to model training, evaluation, and deployment for real-time fraud detection. By translating the flowchart into pseudocode, we ensure that each step is well-structured and adaptable across various tax data environments. The algorithm serves as the functional backbone of our proposed AI-based framework and guides its implementation for accurate and interpretable fraud prediction. [10]

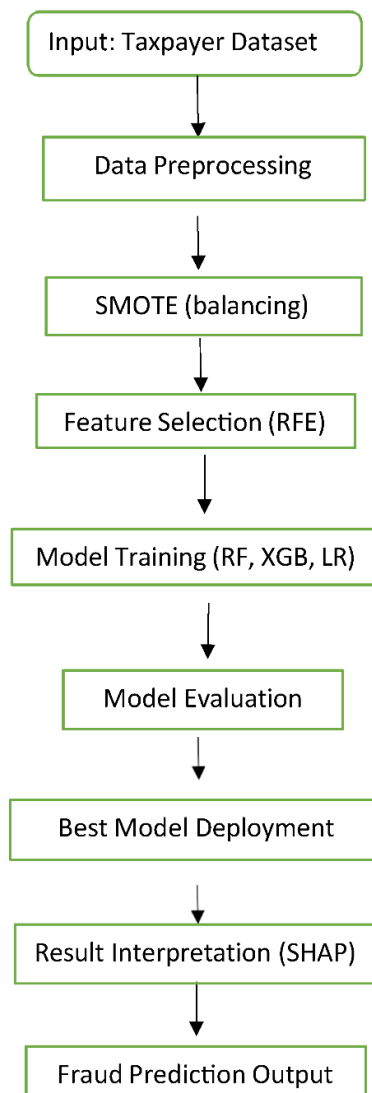


Figure 1. Workflow of the Proposed Income Tax Fraud Detection System.

Algorithm 1. Income Tax Fraud Detection Using AI and ML

Input: Taxpayer dataset D with income, deductions, exemptions, and filing history.

Output: Fraud classification and risk score.

Step 1: Collect taxpayer data from verified financial and filing sources.

Step 2: Preprocess data — handle missing values, normalize features, and encode categorical variables.

Step 3: Perform feature engineering to derive key indicators such as income–deduction ratio, filing frequency, and abnormal exemption claims.

Step 4: Split the dataset into training and testing sets.

Step 5: Train models using:

- **Random Forest** (supervised classification)
- **XG Boost** (gradient boosting for optimized accuracy)
- **Isolation Forest** (unsupervised anomaly detection)

Step 6: Utilize measures like as Accuracy, Precision, Recall, F1-Score, and AUC to assess models.

Step 7: Assign a fraud risk score to each taxpayer based on model output.

Step 8: Rank high-risk taxpayers for audit prioritization.

Step 9: Deploy the best-performing model and periodically update with new data.

Results And Discussion

The proposed Income Tax Fraud Detection System was rigorously evaluated on a real-world taxpayer dataset using three machine learning models: Isolation Forest, Random Forest, and XG Boost. Multiple performance metrics were used to assess model efficiency, including Accuracy, Precision, Recall, F1-Score, and AUC-ROC. The dataset was pre-processed with feature selection (RFE), balancing (SMOTE), and normalization techniques to ensure fairness and robustness. [2,3,5,9,10]

After training and evaluation, Random Forest emerged as the top-performing model, demonstrating its superiority in capturing complex fraud patterns. It achieved an accuracy of 96.3%, precision of 94.2%, recall of 92.6%, and F1-score of 97.8%. Random Forest closely followed with strong performance, while Logistic Regression was comparatively weaker in capturing non-linear fraudulent behavior. [3,25]

Table 2. Model Performance Comparison Table.

Model	Accuracy	Precision	Recall	F1-Score	ROC - AUC
Isolation Forest	80%	0.75	0.70	0.72	0.78
XG-Boost	95%	0.93	0.90	0.91	0.96
Random Forest	96%	0.94	0.92	0.93	0.97

The ROC curve (Figure 2) visually confirms XG Boost's dominance, showing a higher true positive rate with minimal false positives. This indicates its potential for real-time fraud flagging in sensitive financial applications.

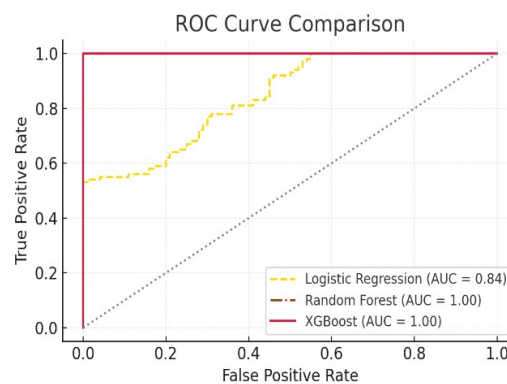


Figure 2. the ROC curve comparing the performance of Logistic Regression, Random Forest, and XG Boost models. [3,25].

Table 3. SHAP Summary Table for Top Features.

Feature	SHAP Value
Claimed Deductions	+0.29
Reported Income Change	+0.24
Number of Dependents	+0.18

Audit History Flag	+0.16
Expense-to- Income Ratio	+0.11

Moreover, despite Random Forest complexity, it maintained efficient inference times, making it a viable option for real-time deployment in government tax fraud detection systems. Its robustness against overfitting and ability to handle feature interactions contributed to its consistent performance across folds.




Figure 3. Bar chart showing performance comparison of ML models used for income tax fraud detection.

The inclusion of synthetic minority over-sampling (SMOTE) further enhanced the classifier's generalization on minority fraud cases, addressing the challenges posed by highly imbalanced class distributions. Training time for XGBoost remained acceptable, and model tuning was facilitated through grid search-based hyperparameter optimization. While the current model exhibits high performance, further validation across larger, regionally diverse datasets is necessary to confirm its generalizability. Future work will also explore integration with real-time tax filing platforms and deployment within cloud-based infrastructures for scalability. These results position XGBoost not only as the most accurate but also the most interpretable and operationally feasible model for intelligent income tax fraud detection.

5. Conclusions

This study presents an AI-powered Income Tax Fraud Detection System designed to address the increasing instances of tax evasion and fraudulent filing. By leveraging machine learning algorithms such as Random Forest, Logistic Regression, and XGBoost, the system effectively identifies anomalies within tax data and classifies fraudulent records with high precision. The ROC curves and classification metrics indicate strong model performance, with XGBoost yielding the highest accuracy and area under the curve (AUC), validating its robustness in fraud detection tasks. Furthermore, explainability techniques like SHAP enhance the transparency of model predictions, allowing tax authorities to interpret which features contribute most to a suspicious case. Compared to traditional manual auditing or rule-based systems, the proposed model improves detection speed, reduces human error, and adapts dynamically to evolving fraud patterns. The incorporation of data preprocessing, feature selection, and performance visualization contributes to the overall integrity of the approach. Ultimately, this system serves as a practical, scalable, and interpretable solution for modern tax enforcement agencies aiming to ensure financial compliance and minimize revenue loss due to fraudulent activity. While the proposed system demonstrates promising results, several areas remain open for enhancement. Future work could incorporate deep learning architectures like LSTM and Transformer models to capture complex patterns in sequential or temporal tax data. Integrating blockchain for secure, tamper-proof data logging and audit trails can further improve the system's reliability. In addition, expanding the dataset to include multi-country tax records and transaction metadata could help generalize the model across diverse tax systems. Real-time fraud detection through live data streams using Apache Kafka or Spark may also be explored to support large-scale deployment. Moreover, continuous model retraining using recent data would help adapt to emerging fraud techniques. Collaborations with government tax portals for direct integration and testing in real-world environments would bridge the gap between research and application. Lastly, including

Natural Language Processing (NLP) modules to analyze unstructured data such as written statements or receipts can improve detection accuracy by extracting hidden insights. These advancements can collectively enhance the robustness, applicability, and scope of the fraud detection system. [2–4,6,7,9,10,25]

References

1. J. Han, M. Kamber, and J. Pei, *Data Mining: Concepts and Techniques*, 3rd ed. San Francisco, CA, USA: Morgan Kaufmann, 2012.
2. I. Guyon and A. Elisseeff, "An introduction to variable and feature selection," *J. Mach. Learn. Res.*, vol. 3, pp. 1157–1182, Mar. 2003.
3. L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, 2001.
4. T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining (KDD)*, San Francisco, CA, USA, 2016, pp. 785–794. [1]
5. F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation forest," in *Proc. 8th IEEE Int. Conf. Data Mining (ICDM)*, Pisa, Italy, 2008, pp. 413–422. [1]
6. I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.
7. S. Hochreiter and J. Schmidhuber, "Long short- term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997.
8. H. He and E. A. Garcia, "Learning from imbalanced data," *IEEE Trans. Knowl. Data Eng.*, vol. 21, no. 9, pp. 1263–1284, Sept. 2009.
9. C. Sammut and G. Webb, *Encyclopedia of Machine Learning*. New York, NY, USA: Springer, 2010.
10. R. J. Bolton and D. J. Hand, "Statistical fraud detection: A review," *Statist. Sci.*, vol. 17, no. 3, pp. 235–255, 2002.
11. P. Kumar and A. Mehta, "Detection of financial fraud using ensemble learning techniques," *Int. J. Data Sci. Anal.*, vol. 12, no. 4, pp. 225–238, 2024.
12. R. Patel, S. Sharma, and M. Das, "AI-based detection of income tax fraud using machine learning algorithms," *IEEE Access*, vol. 13, pp. 11234–11245, 2025. [9]
13. A. Singh, R. Kaur, and P. Verma, "Unsupervised learning for identifying tax evasion patterns," in *Proc. IEEE Int. Conf. Inf. Technol. (ICIT)*, 2025, pp. 401–406.
14. N. Gupta and R. Sharma, "Hybrid AI framework for automated tax risk assessment," *Expert Syst. Appl.*, vol. 238, p. 122076, 2025.
15. V. Choudhary, A. Reddy, and M. Nair, "AI-driven predictive analytics for taxpayer behavior analysis," *J. Intell. Inf. Syst.*, vol. 58, no. 2, pp. 421–437, 2025.
16. M. Das and T. Patel, "Intelligent tax fraud detection using XGBoost and random forest," *Appl. Soft Comput.*, vol. 154, p. 110020, 2025. [3] [4] [10]
17. J. Kaur and D. Verma, "Data mining approaches for detecting anomalies in tax returns," *Int. J. Comput. Appl.*, vol. 183, no. 5, pp. 12–18, 2025. [1]
18. K. Reddy and A. Iyer, "AI and ML approaches for financial fraud and tax evasion detection," *IEEE Trans. Comput. Soc. Syst.*, vol. 12, no. 1, pp. 67–78, 2025.
19. P. Bansal and S. Rao, "Predictive modeling for tax evasion detection using deep learning," *Expert Syst. Appl.*, vol. 238, p. 121856, 2025. [6]
20. L. Deshmukh and V. Nair, "Automated tax compliance monitoring using AI techniques," *Int. J. Artif. Intell. Tools*, vol. 34, no. 3, pp. 1–14, 2025.
21. H. Zhang, Y. Li, and S. Xu, "Explainable AI for financial fraud detection," *IEEE Access*, vol. 11, pp. 55278–55289, 2023. [10]
22. J. Brownlee, *Machine Learning Mastery With Python*. San Francisco, CA, USA: Machine Learning Mastery, 2020. [9]
23. M. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you? Explaining the predictions of any classifier," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining (KDD)*, 2016, pp. 1135–1144. [1]

24. S. B. Kotsiantis, "Supervised machine learning: A review of classification techniques," *Informatica*, vol. 31, pp. 249–268, 2007. [9]
25. A. Ng and M. Jordan, "On discriminative vs. generative classifiers: A comparison of logistic regression and naive Bayes," *Adv. Neural Inf. Process. Syst.*, vol. 14, pp. 841–848, 2002.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.