

## Article

# A Hybrid DNN-Transformer-AE Framework for Corporate Tax Risk Supervision and Risk Level Assessment

Zhenzhen Song <sup>1</sup>, Nanxi Wang <sup>2</sup> and Hongji Li <sup>3,\*</sup><sup>1</sup> Northwest A&F University, Xianyang, Shaanxi, China<sup>2</sup> USC Viterbi School of Engineering, University of Southern California, Los Angeles, CA, USA<sup>3</sup> Columbia University, New York, NY, USA

\* Correspondence: Hongji Li, Columbia University, New York, NY, USA

**Abstract:** Tax risk supervision has become a critical component of modern financial governance, as irregular tax behaviors and hidden compliance risks pose significant challenges to regulatory authorities and enterprises alike. Traditional rule-based methods often struggle to capture complex and dynamic tax-related anomalies in large-scale enterprise data. To address this issue, this paper proposes a hybrid deep learning framework (DNN-Transformer-Autoencoder) for corporate tax risk supervision and risk level assessment. The framework integrates three complementary modules: a Deep Neural Network (DNN) for modeling static enterprise attributes, a Transformer-based architecture for capturing long-term dependencies in historical financial time series, and an Autoencoder (AE) for unsupervised detection of anomalous tax behaviors. The outputs of these modules are fused to generate a comprehensive risk score, which is further mapped into discrete risk levels (high, medium, low). Experimental evaluations on a real-world enterprise tax dataset demonstrate the effectiveness of the proposed framework, achieving an accuracy of 0.91 and a Macro F1-score of 0.88. These results indicate that the hybrid model not only improves classification performance but also enhances interpretability and applicability in practical tax regulation scenarios. This study provides both methodological innovation and regulatory implications for intelligent tax risk management.

**Keywords:** tax risk supervision; risk level assessment; deep learning; DNN-Transformer-AE; anomaly detection

## 1. Introduction

Taxation plays a fundamental role in modern economies, serving both as a primary source of government revenue and as a regulatory mechanism to sustain market order and fairness. However, with the increasing complexity of corporate operations, enterprises often engage in behaviors that introduce compliance risks, such as misreporting, revenue concealment, and aggressive tax avoidance. These practices not only undermine the equity and efficiency of tax collection but also complicate risk monitoring for tax authorities. Traditional supervision approaches, typically based on manual inspection or rule-based models, face inherent limitations: they struggle to scale with massive enterprise data and often fail to capture increasingly concealed and dynamic risk patterns [1].

To address these limitations, data-driven intelligent regulatory approaches have attracted significant attention. By employing machine learning and deep learning methods, researchers can extract latent patterns from multidimensional enterprise data and automate tax risk assessments [2,3]. Such tasks can be formulated either as classification (e.g., high, medium, and low risk) or anomaly detection problems, thereby enabling tax authorities to identify high-risk enterprises, allocate supervisory resources more efficiently, and improve early-warning capabilities. Nevertheless, two challenges remain. First, enterprise data is inherently multimodal, encompassing static attributes such as firm size, industry,

Received: 19 August 2025

Revised: 01 September 2025

Accepted: 28 September 2025

Published: 08 October 2025



**Copyright:** © 2025 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

and registration records, alongside dynamic temporal information including financial statements and tax histories. Second, risky behaviors are often hidden within subtle anomalies, which are difficult to capture using traditional supervised models [4].

In response, this paper proposes a hybrid deep learning framework, the DNN-Transformer-AE model, designed for comprehensive corporate tax risk supervision. The framework integrates classification and anomaly detection to address practical supervisory needs, employs multi-source fusion of static features, temporal dependencies, and anomaly signals to improve predictive accuracy, and demonstrates its effectiveness on real-world corporate tax data. The findings not only contribute methodological innovations but also provide actionable implications for intelligent tax risk management in modern regulatory systems.

## 2. Related Work

Research on corporate tax risk supervision has attracted increasing attention in both academic and practical domains, particularly as governments and regulatory agencies seek more effective ways to detect and prevent tax non-compliance. Existing studies can be broadly categorized into three streams: traditional rule-based approaches, machine learning-based tax risk assessment, and deep learning methods for anomaly detection and risk prediction.

Lan et al. employ a two-layer deep-learning framework to forecast quarterly taxable-input risks and depreciation, continuously updating cumulative exposure with prior claimable fractions; the model enhances enterprise tax-compliance, risk-management, and financial stability while providing an adaptive benchmark for subsequent audits [5].

He et al. integrate a frilled-lizard-optimized gated LSTM into a Python tax-risk platform, achieving 97% accuracy and 2.3 ms per record; their comprehensive pipeline offers subsequent studies a clear deep-learning baseline for pediatric-GI-inspired affinity screening while highlighting scalability gaps that our GNN framework aims to fill [6].

Didimo et al. introduce MALDIVE, a four-step network-mining framework that integrates SNA metrics, ML classifiers, information diffusion, and interactive visualization to flag risky taxpayers for the Italian Revenue Agency, thereby offering a reusable reference for tax-evasion detection and clarifying our own contribution within graph-based risk-assessment research [7].

De Roux et al. propose an unsupervised pipeline that flags under-reporting tax declarations without audited labels, shrinking the audit pool while uncovering previously undetected suspicious cases; their work offers a cost-efficient springboard for subsequent hybrid fraud-detection frameworks and clarifies our own innovation niche amid label-scarce tax analytics [8].

## 3. Methodology

### 3.1. Overall Framework Design

The proposed hybrid framework for corporate tax risk supervision is built upon a DNN-Transformer-AE structure, designed to jointly model multidimensional enterprise information and output both risk scores and risk levels. Specifically, the model consists of three functional modules: a Deep Neural Network (DNN) [9] for static enterprise feature modeling, a Transformer for temporal financial data modeling, and an Autoencoder (AE) for anomaly detection and feature reconstruction. The outputs of these three modules are integrated through a feature fusion layer and subsequently mapped into a final risk score  $y$  via fully connected layers. The score is then categorized into high-risk, medium-risk, or low-risk levels based on pre-defined thresholds.

Formally, the input data can be represented as:

$$X = \{X_s, X_t\} \quad (1)$$

Where  $X_s$  denotes static enterprise attributes (e.g., firm size, industry, registration information), and  $X_t$  represents historical financial and taxation time series data, with  $t$

as the time horizon and  $s$  as the feature dimension. The objective of the model is to learn a mapping function:

$$f: (X_s, X_t) \rightarrow \hat{y}, \quad (2)$$

Where  $y$  is the predicted corporate tax risk score.

### 3.2. DNN Module: Static Enterprise Feature Modeling

Static enterprise features include basic attributes such as firm size, industry type, registered capital, and ownership structure. These features are relatively stable and serve as fundamental indicators of long-term risk. To capture their nonlinear relationships, a deep neural network (DNN) is employed.

Given the static input  $X_s$ , the features are first normalized or embedded, then fed into a multilayer fully connected network:

$$h_s^{(l)} = \sigma(W_s^{(l)} h_s^{(l-1)} + b_s^{(l)}) \quad (3)$$

Where  $h_s^{(0)} = X_s$ ,  $W_s^{(l)}$  and  $b_s^{(l)}$  denote the weight matrix and bias vector of the  $l$ -th layer, and  $\sigma$  is a nonlinear activation function (e.g., ReLU).

### 3.3. Transformer Module: Temporal Data Modeling

Financial and taxation records of enterprises inherently exhibit temporal dependencies. Indicators such as revenue, tax payments, input and output VAT fluctuations may implicitly reflect evolving risk patterns. Conventional recurrent neural networks (RNNs) such as LSTMs often face limitations in capturing long-range dependencies and efficiency. To overcome this, a Transformer-based architecture is applied [10,11].

The time series input  $X_t$  is first embedded and augmented with positional encoding:

$$Z_0 = X_t W_e + PE \quad (4)$$

Where  $W_e$  is the embedding matrix and  $PE$  denotes positional encodings that preserve sequential order.

Within each Transformer encoder block, temporal dependencies are modeled using multi-head self-attention (MHSA):

$$Attention(Q, K, V) = Softmax\left(\frac{QK^T}{\sqrt{d_K}}\right)V \quad (5)$$

Where  $Q$ ,  $K$ ,  $V$  represent query, key, and value projections, respectively.

### 3.4. Autoencoder (AE) Module: Anomaly Detection

In tax risk supervision, high-risk behaviors are often characterized by deviations from normal patterns. To capture such latent anomalies, an Autoencoder (AE) [12] is integrated for feature reconstruction and anomaly detection in the figure 1.


Specifically, the concatenated features  $Z = [hs, ht]$  are compressed into a latent representation  $z$  through the encoder, and then reconstructed via the decoder:

$$z = \sigma(W_{enc}Z + b_{enc}), \quad \hat{Z} = \sigma(W_{dec}z + b_{dec}) \quad (6)$$

Where  $W_{enc}$ ,  $b_{enc}$  and  $W_{dec}$ ,  $b_{dec}$  are encoder and decoder parameters, respectively. The reconstruction objective is to minimize:

$$L_{AE} = ||Z - \hat{Z}||^2 \quad (7)$$

By learning the distribution of normal samples, the AE identifies anomalies when reconstruction error exceeds a certain threshold, providing additional risk signals to the overall framework [13].



**Figure 1.** Structure diagram of Autoencoder Module.

### 3.5. Feature Fusion and Risk Level Output

After obtaining the representations from  $h_s$ , Transformer  $h_t$ , and AE  $z$ , these outputs are concatenated into a joint representation:

$$H = [h_s][h_t][z] \quad (8)$$

Which is then passed through a fully connected layer to predict the risk score:

$$\hat{y} = \text{softmax}(WH + b) \quad (9)$$

Where  $\hat{y}$  denotes the probability distribution across high-risk, medium-risk, and low-risk classes. The final enterprise risk level is determined by the highest probability or threshold-based criteria.

## 4. Experimental Result

### 4.1. Dataset

The dataset used in this study originates from real corporate tax declarations and financial records collected by tax regulatory authorities, covering multidimensional information such as enterprise attributes, financial indicators, and taxation behaviors. Specifically, the dataset includes approximately 12,000 enterprises across multiple industries, including manufacturing, retail, and internet services, ensuring the generalizability of the model across different sectors. The features can be grouped into three categories:

- 1) **Static features:** industry category, region, company size, and historical compliance indicators.
- 2) **Time-series features:** quarterly revenue, profit, tax paid, and invoice issuance.
- 3) **Labels:** tax risk level assigned by domain experts (low, medium, high).

Data preprocessing includes normalization of numerical attributes, one-hot encoding of categorical variables, and interpolation for missing time-series values. The dataset is divided into training (70%), validation (15%), and test (15%) sets with temporal stratification.

### 4.2. Classification Performance

In the initial stage of the experiment, this study first selected several traditional machine learning and deep learning methods as baseline models to verify the performance differences of different models in enterprise tax risk rating tasks. Traditional methods include logistic regression (LR), random forest (RF), and XGBoost [14], while deep learning methods include DNN-LSTM [15] models that combine static and temporal features. In order to comprehensively evaluate the performance of the proposed DNN Transformer Autoencoder hybrid model in enterprise tax risk rating, this study adopted multiple classification and detection indicators. Specifically, the main evaluation criteria for the model include Accuracy, Recall, and Macro F1 score, which measure the model's overall prediction accuracy, ability to identify different risk levels, and balance between categories. In addition, due to the limited sample size and imbalanced categories of high-risk enterprises in corporate tax data, this study places special emphasis on the performance of the Macro F1 score to ensure that the model not only predicts low-risk enterprises well but also identifies high-risk enterprises.


Table 1 summarises the classification performance of different models on the test set for corporate tax-risk supervision. The conventional Logistic Regression achieves an accuracy of 0.79, a recall of 0.68, and an F1-score of 0.70, indicating limited capacity to detect potentially risky firms. Random Forest improves overall performance to 0.84 accuracy, 0.75 recall, and 0.77 F1, benefiting from stronger non-linear modelling, whereas XGBoost further raises these metrics to 0.86, 0.77, and 0.79 by exploiting feature interactions. The deep-learning DNN-LSTM advances accuracy to 0.89 with a recall of 0.82 and F1 of 0.85, demonstrating superior capture of temporal dependencies in tax data.

**Table 1.** Performance Comparison of Different Models in Enterprise Tax Risk Supervision Tasks.

Model	Accuracy	Recall	F1-score
Logistic Regression (LR)	0.79	0.68	0.70
Random Forest (RF)	0.84	0.75	0.77
XGBoost	0.86	0.77	0.79
DNN - LSTM	0.89	0.82	0.85
<b>DNN-Transformer-AE</b>	<b>0.91</b>	<b>0.86</b>	<b>0.88</b>

Our proposed hybrid DNN-Transformer-Autoencoder outperforms all baselines, attaining 0.91 accuracy, 0.86 recall, and 0.88 F1. The higher recall enhances sensitivity to high-risk companies, while the elevated F1 reflects a better precision-coverage balance. Overall, the DNN-Transformer-Autoencoder more comprehensively extracts deep representations from corporate tax records, delivering more accurate and reliable risk identification in tax-supervision tasks.

Figure 2 presents the loss curves for both the training and validation datasets. The x-axis represents epochs, which refers to the number of times the model has gone through the entire training data. The y-axis represents loss, a measure of how much error the model makes in its predictions. The training and validation loss are represented by the orange solid line and the blue dashed line, respectively. As the number of epochs increases, both curves exhibit a steady decline and converge at approximately 40 epochs, stabilizing around a loss of 0.05. An important observation is that the validation loss consistently remains lower than the training loss throughout the process. This indicates that the DNN-Transformer-AE model achieves stable convergence and demonstrates strong generalization ability, as the lower validation loss reflects its robustness in capturing meaningful patterns beyond the training data rather than overfitting to the training set.



**Figure 2.** Loss function during the training process.

This performance is particularly significant in the context of corporate tax risk supervision and risk level assessment. Enterprises exhibit complex patterns across static attributes (e.g., registered capital, industry classification) and dynamic financial behaviors (e.g.,



historical tax declarations, profit fluctuations), while abnormal taxation behaviors are often concealed within irregular patterns. The rapid convergence of the hybrid model suggests that the DNN component effectively extracts static enterprise characteristics, the Transformer captures long-term temporal dependencies in financial records, and the Autoencoder identifies subtle anomalies through reconstruction error. The stable validation loss further confirms that the model can generalize to unseen corporate data, enabling accurate detection of high-risk firms.

## 5. Conclusion

This study proposes a hybrid deep-learning framework DNN-Transformer-Autoencoder, for corporate tax-risk supervision and risk-level assessment. By integrating static enterprise attributes with dynamic financial time-series and introducing an autoencoder module for unsupervised detection of anomalous tax behavior, the framework significantly improves the capture of complex risk patterns. Evaluated on a real-world tax dataset, the hybrid model outperforms both traditional methods and single-architecture deep learners, achieving 0.91 accuracy and 0.88 macro-averaged F1 on the test set. These results demonstrate that the proposed approach balances precision and coverage, offering regulators a more reliable tool for grading enterprise risk.

From a regulatory perspective, the work has strong practical relevance. As corporate financial behavior grows more intricate and data volumes expand, rule-based systems struggle to identify latent tax risks. Our hybrid framework not only uncovers hidden compliance threats but also stratifies firms into high, medium, and low risk tiers, providing tax authorities with a new technological lever for precision supervision and resource allocation while giving firms actionable feedback for self-remediation.

Despite its strong performance, this study has limitations. First, the input features are primarily confined to financial and operational indicators, while macroeconomic conditions and inter-firm relationships remain unincorporated. Second, the model's computational complexity may pose challenges for deployment in resource-constrained tax offices. Future research could further explore the integration of macroeconomic and policy data into the feature space and develop interpretable, lightweight models (e.g., attention-based visualizations or model compression) to enhance usability in practice.

In conclusion, this study, through the integration of DNN, Transformer, and Autoencoder methods, reveals that hybrid deep learning frameworks can effectively capture complex tax risk patterns and stratify firms into risk levels. These insights provide regulators with a robust, data-driven approach to corporate tax supervision, offering both methodological innovation and practical guidance for advancing financial governance.

## References

1. E. Ezeife, E. Kokogho, P. E. Odio, and M. O. Adeyanju, "The future of tax technology in the United States: A conceptual framework for AI-driven tax transformation," *Future*, vol. 2, no. 1, p. 101203, 2021.
2. N. Braun Binder, "Artificial intelligence and taxation: risk management in fully automated taxation procedures," In *Regulating Artificial Intelligence*, 2019, pp. 295-306. doi: 10.1007/978-3-030-32361-5\_13.
3. H. Strauss, T. Fawcett, and D. Schutte, "Tax risk assessment and assurance reform in response to the digitalised economy," *Journal of Telecommunications and the Digital Economy*, vol. 8, no. 4, pp. 96-126, 2020. doi: 10.18080/jtde.v8n4.306.
4. N. Braun Binder, "Ai and taxation: Risk management in fully automated taxation procedures," *Available at SSRN* 3293577, 2018. doi: 10.2139/ssrn.3293577.
5. Y. Lan, "Enterprise Tax Assessment and Risk Avoidance Based on Deep Learning," *International Journal of Computational Intelligence Systems*, vol. 18, no. 1, pp. 1-22, 2025. doi: 10.1007/s44196-025-00837-0.
6. X. He, "Research on Tax Risk Identification and Assessment System Assisted by Software and Deep Learning," *International Journal of High Speed Electronics and Systems*, 2025. doi: 10.1142/s0129156425405613.
7. W. Didimo, L. Grilli, G. Liotta, L. Menconi, F. Montecchiani, and D. Pagliuca, "Combining network visualization and data mining for tax risk assessment," *Ieee Access*, vol. 8, pp. 16073-16086, 2020. doi: 10.1109/access.2020.2967974.

8. D. De Roux, B. Perez, A. Moreno, M. D. P. Villamil, and C. Figueroa, "Tax fraud detection for under-reporting declarations using an unsupervised machine learning approach," In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, July, 2018, pp. 215-222.
9. A. L. E. X. A. N. D. E. R. Raikov, "Decreasing tax evasion by artificial intelligence," *IFAC-PapersOnLine*, vol. 54, no. 13, pp. 172-177, 2021.
10. Y. Li, X. Jiang, and Y. Wang, "TRAM-FIN: A transformer-based real-time assessment model for financial risk detection in multinational corporate statements," *Journal of Advanced Computing Systems*, vol. 3, no. 9, pp. 54-67, 2023. doi: 10.69987/jacs.2023.30905.
11. X. W. Tan, and S. Kok, "Explainable AI for Comprehensive Risk Assessment for Financial Reports: A Lightweight Hierarchical Transformer Network Approach," *arXiv preprint arXiv:2506.23767*, 2025.
12. Z. Chen, "Tax Risk Early Warning System for SMEs Using Auto EncoderBackpropagation Neural Network with Genetic Algorithm Optimization," *Informatica*, vol. 49, no. 6, 2025. doi: 10.31449/inf.v49i6.6830.
13. M. R. Islam, "AI-AUGMENTED TAX RISK SCORING FOR SMALL AND MEDIUM ENTERPRISES: A PANEL DATA STUDY," *ASRC Procedia: Global Perspectives in Science and Scholarship*, vol. 1, no. 01, pp. 501-531, 2025.
14. M. S. Beasley, N. C. Goldman, C. M. Lewellen, and M. McAllister, "Board risk oversight and corporate tax-planning practices," *Journal of Management Accounting Research*, vol. 33, no. 1, pp. 7-32, 2021. doi: 10.2308/jmar-19-056.
15. L. Xin, Y. Zhao, and Q. Zhao, "BeiDou and SAR fusion technology with AI in reservoir dam monitoring for climate-based disaster mitigation," *Discover Artificial Intelligence*, vol. 5, no. 1, p. 100, 2025. doi: 10.1007/s44163-025-00320-x.

**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 GBP and/or the editor(s). GBP 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.