Transactions on Computational and Scientific Methods | Vo. 6, No. 04, 2026
ISSN: 2998-8780
https://pspress.org/index.php/tcsm
Pinnacle Science Press

Learning Risk Dynamics in Financial Time Series via Deep Neural Networks

Elliot Mercer
University of North Texas, Denton, United States
elliot.mercer@unt.edu

Abstract: Indicators derived from financial time series data may yield different values under different accounting standards, and their selection is often influenced by human factors. To address the diversification and complexity of feature distributions in financial time series, this paper proposes an anomaly detection method tailored to generalized financial time series data. The proposed approach captures intrinsic representation patterns of financial sequences and treats the outputs obtained from a recurrent neural network as learned knowledge. A standard classifier is then employed to adapt the marginal distribution in the feature space, followed by further prediction using a latent variable regression model to enhance predictive accuracy. Subsequently, a latent variable regression framework is constructed for financial risk forecasting, identifying potential financial risks by modeling the feature distributions embedded in financial time series data. Experimental results demonstrate the feasibility and effectiveness of the proposed model.

Keywords: Deep learning; financial time series data; feature distribution; financial risk; anomaly detection

1. Introduction

Deep learning has been widely applied in the fields of computer vision and natural language processing, where well-established architectures such as convolutional neural networks (CNNs), recurrent neural networks (RNNs), and transformer-based models have achieved remarkable success. These domains benefit from large-scale labeled datasets and relatively stable data distributions, enabling rapid advancements in model design and optimization. In contrast, the application of deep learning to financial time series data presents a different set of challenges, including high noise levels, non-stationary distributions, complex temporal dependencies, and strong sensitivity to external macroeconomic factors. As a result, despite growing interest, the development of effective deep learning architectures tailored to financial data remains less mature compared to image and language domains.

In recent years, the processing and modeling of financial time series data have attracted increasing attention as a promising direction for future research. Financial data is inherently dynamic and often exhibits abrupt structural changes, making it difficult for traditional statistical models to capture long-term dependencies and nonlinear relationships. Deep learning techniques, with their strong representation learning capabilities, offer new opportunities to address these limitations by automatically extracting latent patterns from high-dimensional and heterogeneous data sources. This includes not only numerical time series such as stock prices and trading volumes, but also auxiliary information such as news text, social media sentiment, and macroeconomic indicators.

Attempts to leverage deep learning techniques to enhance financial risk prevention and control measures are of significant research and practical value. From a theoretical perspective, such efforts contribute to the advancement of sequential modeling, uncertainty estimation, and multimodal data fusion under complex and evolving environments. From a practical standpoint, improved risk prediction and early warning systems can play a critical role in financial decision-making, helping institutions better manage credit risk, market volatility, and systemic instability. These capabilities are particularly important in the context of increasingly interconnected global financial systems, where localized disruptions can rapidly propagate across markets.

Therefore, integrating deep learning methods into financial time series analysis not only advances fundamental research but also provides essential tools for safeguarding national economic and social development against financial risks. Continued exploration in this direction is expected to drive innovations in both model architectures and application frameworks, ultimately enabling more accurate, robust, and interpretable financial intelligence systems.

2. Related Work

Anomaly detection has been extensively studied as a fundamental problem in data mining and machine learning, with early work providing a comprehensive taxonomy of techniques across statistical, distance-based, and density-based paradigms [1]. With the advancement of deep learning, representation learning has become a dominant paradigm for modeling complex data distributions. Foundational studies highlight the importance of learning hierarchical feature representations for capturing intrinsic data structures in high-dimensional spaces [2], while probabilistic generative models such as variational autoencoders further enable latent variable modeling for uncertainty-aware inference [3].

In financial and business scenarios, anomaly detection faces unique challenges due to non-stationary distributions, label scarcity, and evolving data patterns. Recent works attempt to address these issues by integrating context-aware and retrieval-based mechanisms. For instance, CoReAD introduces a retrieval-augmented anomaly detection framework that dynamically incorporates contextual information for evolving tabular data environments [4]. Similarly, meta-learning and large language models have been explored for few-shot financial fraud detection, improving generalization under limited supervision [5]. In addition, anomaly ranking methods based on latent structural deviations and reconstruction consistency further enhance interpretability and prioritization in enterprise financial systems [6].

Time series forecasting forms another critical foundation for financial risk prediction. Early neural approaches such as dual-stage attention-based recurrent networks capture both temporal dependencies and feature importance [7], while hybrid models combining LSTM with statistical volatility models improve predictive performance in financial markets [8]. More recently, transformer-based architectures, such as Temporal Fusion Transformers, have demonstrated strong capability in modeling multi-horizon temporal dynamics with interpretability [9]. Furthermore, stochastic recurrent models have been proposed to enhance robustness in multivariate time series anomaly detection under noisy conditions [10].

Beyond standalone prediction models, recent research has emphasized structured and spatiotemporal representation learning. Transformer-based and graph-based approaches enable joint modeling of spatial dependencies and temporal evolution, improving risk representation in complex systems [11]. Complementary work explores causal representation learning to enhance robustness and interpretability in financial auditing and risk identification tasks [12]. These methods highlight the importance of capturing both structural dependencies and underlying causal mechanisms in financial data.

In parallel, the emergence of large language models and multi-agent systems has expanded the scope of intelligent decision-making. Multi-agent LLM frameworks have been applied to business decision-making with auditable attribution mechanisms [13], while structured state representation and constraint-guided policy learning further improve decision optimization in intelligent systems [14]. Privacy-preserving learning, such as adaptive federated learning, has also been introduced to address data heterogeneity and regulatory constraints in financial environments [15]. Additionally, recent works consider distribution shift and class imbalance jointly, proposing robust learning frameworks for business risk prediction [16].

From an application perspective, representation learning and causal inference have also been integrated into downstream decision systems such as advertising and ranking optimization, where calibrated multi-objective learning improves utility under uncertainty [17]. These developments collectively indicate a shift toward unified frameworks that combine representation learning, uncertainty modeling, and decision optimization.

Despite these advances, existing methods still face limitations in handling distribution shifts, noisy labels, and evolving financial environments simultaneously. Most approaches either focus on predictive accuracy or anomaly detection in isolation, lacking a unified framework that jointly models feature distribution adaptation, latent temporal dynamics, and anomaly discrimination. To address these challenges, this paper proposes a hybrid deep learning framework that integrates recurrent representation learning with latent autoregressive modeling, enabling robust financial risk prediction under complex and non-stationary conditions.

3. Research Scheme

3.1 Dividing Data Learning and Utilization into Two Components

In the proposed model, all labels in the financial time series domain are assumed to be verifiable only after the occurrence of events. Accordingly, in the model design, the outputs obtained after the recurrent neural network are regarded solely as learned knowledge. A standard classifier is employed to adapt the marginal distribution in the feature space, and a latent autoregressive model is subsequently utilized for further prediction, thereby improving predictive accuracy.

3.2 Identifying the Distribution Characteristics of Financial Risks Using a Latent Autoregressive Model

This study employs a latent autoregressive model to compute the distributional distance between predicted values and labels, enabling distribution adaptation in the feature space. Although most deep learning studies assume that data satisfy the independent and identically distributed (i.i.d.) condition, financial time series data are not strictly inconsistent with this assumption when compared with other datasets, yet they cannot be considered to fully comply in a rigorous sense.

In practical financial scenarios, “noise” may arise from data distributions intentionally manipulated by various institutions. Considering the temporal dependence, spatial dependence, and the influence of financial “noise” distributions, the proposed framework ensures that normal data exhibit an independent distribution property in the latent feature space without introducing synthetic data, thereby enhancing the robustness of the model.

4. Scheme Design

In this study, hierarchical architectures of various deep learning models are integrated to establish a data processing and learning framework that can be seamlessly interconnected. A hybrid architecture combining a recurrent neural network and a latent autoregressive model is constructed, as illustrated in Figure 1. Through latent feature learning, past financial time series data and the prediction errors generated by the latent autoregressive model are jointly learned in a unified manner.

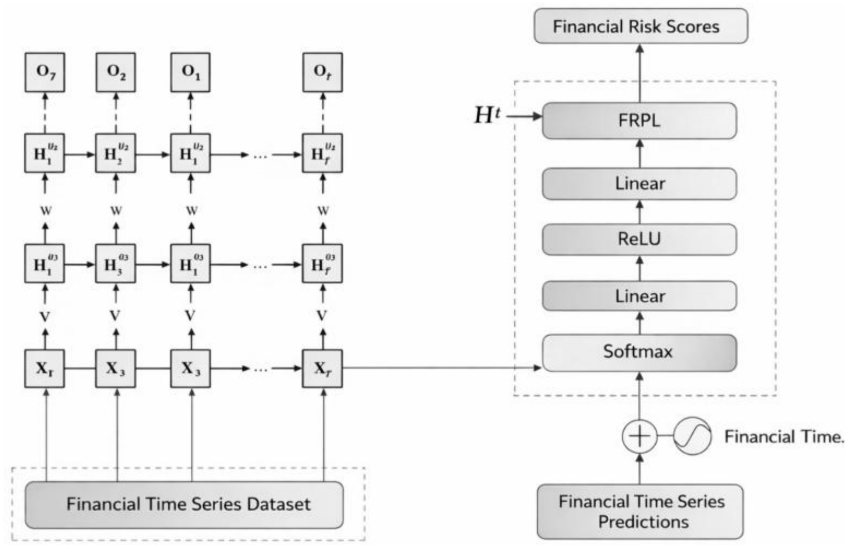

Figure 1. Financial Time Series Risk Prediction Model Framework

In financial time series datasets, each financial institution contains m indicators, and there exist $n(n > 1)$ financial institutions. Let S denote the indicator set. The overall financial time series dataset contains $N = |S| = mn$ indicators. Each indicator is represented as $x_i \in S$, where $x_i \in \mathbb{R}^{t_i \times 1}$, and t_i denotes the length of indicator x_i.

The objective is to perform financial risk prediction on this type of time series data. By leveraging similarity measures to learn from large-scale financial datasets, the model identifies a more effective approach for computing risk thresholds from financial time series.

Assume that the time series data of a given institution are denoted as X_i. The feature knowledge transmitted to the latent autoregressive model is computed by the recurrent neural network, as shown in (1):

$$H_t = \phi(X_t W_{xh} + H_{t-1} W_{hh} + b_h)$$

where H_{t-1} represents the feature knowledge obtained from the financial time series of the institution at the previous time step. The input sample is expressed as $X = (x_1, \cdots, x_N)^{\top} \in \mathbb{R}^{N \times T}$, which contains all data sources.

Through the proposed method, the relationship between adjacent latent variables H_t and H_{t-1} is modeled, and the learned information is preserved in the latent variable at the current time step. These latent variables are then passed to a fully connected layer with an activation function in the subsequent computation layer. The latent autoregressive algorithm is subsequently applied to perform linear prediction of anomalous features in the financial time series dataset. Features from different institutions are incorporated into the current learned feature representation for latent autoregressive analysis. Distribution information in the feature space is utilized to determine whether the features of the current data source are anomalous. According to the predefined objective, the feature knowledge learned by the recurrent neural network within the designated financial time series dataset serves as the input to the latent autoregressive model.

To measure the distance between labels and model predictions in the financial time series dataset, cross-entropy is adopted. Considering the prediction error pattern of an institution's financial time series dataset along the timeline, time accumulation is used to evaluate the cross-entropy loss across all time-step predictions, as shown in (2):

$$\exp\left(-\frac{1}{n}\sum_{t=1}^{n}\log P\big(L_{CEL_t} \mid L_{CEL(t-1)},\ldots,L_{CEL1}\big)\right)$$

The front-end recurrent neural network is trained on financial time series data to learn latent knowledge that satisfies the independent distribution requirement. The latent autoregressive model is then used to learn prediction error patterns across related institutions and time points in the financial time series dataset, as shown in (3):

$$J(\theta) = L_{CEL} + \lambda L_{MSE}$$

Here, λ is a hyperparameter of the objective function used to balance the relative weights of different error terms. L_{CEL} denotes the training error of the front-end recurrent neural network, indicating whether the learned latent knowledge can effectively preserve the feature distribution of input samples; the error is computed using cross-entropy. L_{MSE} represents the training error of the latent autoregressive model, which measures the distance between normal and abnormal conditions to detect financial risks.

Ultimately, deep learning yields a risk threshold whose meaning may vary across different financial scenarios. For financial institutions, it may serve as a stop-loss threshold; for regulatory agencies, it may function as a supervisory warning line.

The training set $X = \{x_1, x_2, \ldots, x_D\}$, where $x_i \in \mathbb{R}^{N\times T}$, represents feature sets within a real time period containing risk events. $\mathrm{FRPL}(x_i)$ denotes the cumulative distance obtained for sample x_i after training the recurrent neural network combined with the latent autoregressive model. Let u be its mean value. η is a hyperparameter. After continuous application of the proposed model to construct a feature label library of this hyperparameter, η may also be learned through deep models, as shown in (4):

$$FR_{threshold} = \eta \cdot \sqrt{\frac{1}{D}\sum_{i=1}^{D}\big(\mathrm{FRPL}(x_i) - u\big)^2}$$

During the prediction stage, if the distance of a predicted sample feature x_i satisfies $\mathrm{FRPL}(x_i) > FR_{threshold}$, the sample x_i is identified as financial risk; otherwise, it is classified as normal.

5. Experimental Results and Analysis

5.1 Comparative Analysis with the Standalone Recurrent Neural Network

The experimental dataset consists of proprietary financial time series data reflecting aggregated indicators of a regional financial sector. To evaluate the effectiveness of the proposed framework, three modeling strategies are implemented: a recurrent neural network (RNN), a latent autoregressive model, and the proposed hybrid model integrating both components. The predicted outputs are compared against the ground-truth labels, corresponding to the solid and dotted curve segments illustrated in Figure 2.

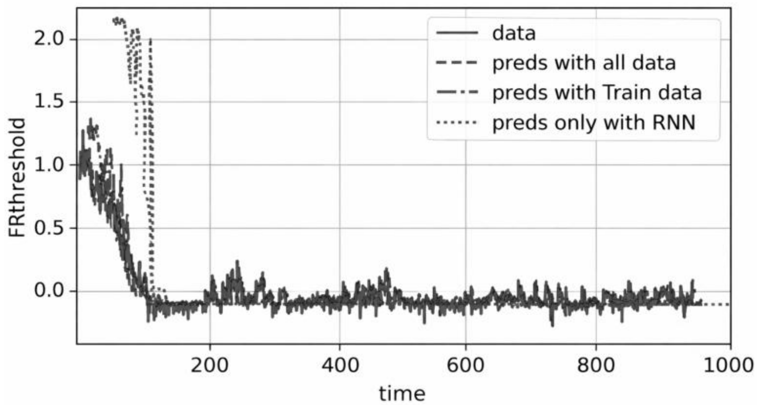

Figure 2. FRthreshold Comparison with Classical RNN

As shown in Figure 2, the standalone RNN exhibits relatively stable short-term prediction capability; however, its long-term performance gradually deteriorates due to accumulated temporal errors. In contrast, the hybrid model maintains improved stability over extended horizons by incorporating latent autoregressive correction. Although predictive performance declines as the forecasting window increases, the combined architecture demonstrates stronger robustness than the single-model configuration.

Furthermore, when the prediction test set is partially incorporated into the training process, the model generates the dashed curve presented in Figure 2. This configuration simulates regulatory deployment scenarios, where adaptive updating mechanisms enhance practical applicability.

5.2 Loss Function Comparison-Cross-Entropy Versus Mean Squared Error

Financial time series data are susceptible to covariate shift, leading to gradual error accumulation and instability in gradient descent optimization. Once the underlying financial distribution deviates from historical patterns, models relying solely on mean squared error L_{MSE} may experience degraded convergence performance or even local singularity issues.

Unlike conventional anomaly detection approaches that use L_{MSE} to measure distributional differences, the proposed framework incorporates cross-entropy to characterize divergence between predicted and actual distributions. As illustrated in Figure 3, cross-entropy improves distribution alignment under domain shift conditions. The experimental results indicate that combining cross-entropy with latent autoregressive modeling reduces instability and enhances anomaly discrimination capability.

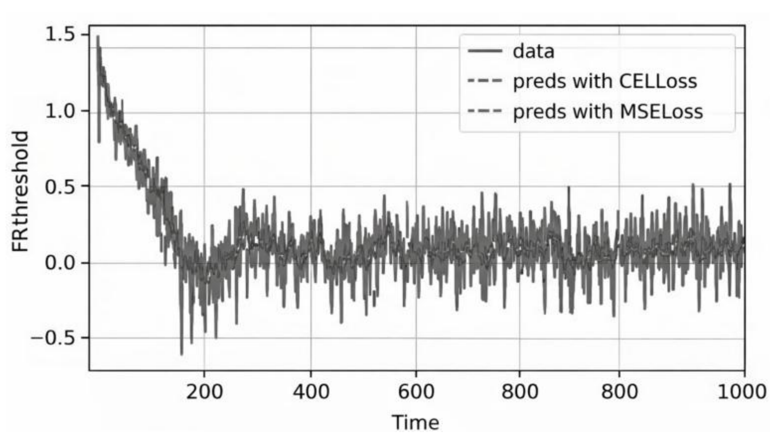

Figure 3. Comparison Between Cross-Entropy and Mean Squared Error Loss

5.3 Risk Prediction Under Normal Financial Market Conditions

When extreme events - such as black swan or gray rhino scenarios - are excluded from the dataset, the financial time series exhibit relatively stable feature distributions. Under such conditions, the proposed model effectively tracks the evolution of financial risk thresholds, as shown in Figure 4.

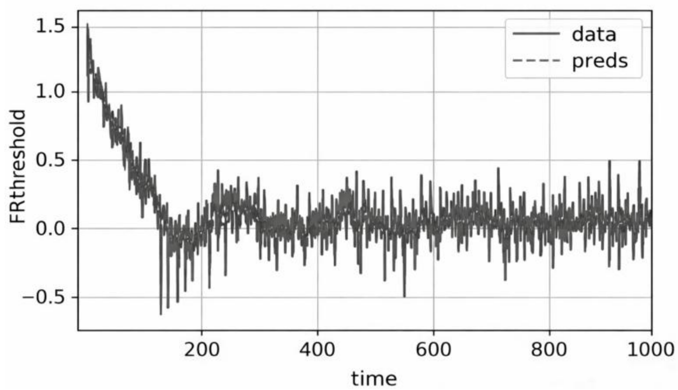

Figure 4. Risk Prediction Performance Under Normal Financial Market Conditions

The model captures gradual structural changes in the latent feature space and generates interpretable risk thresholds. In financial institutions, these thresholds may function as stop-loss signals; in regulatory contexts, they may serve as supervisory warning lines. Because normal and anomalous samples may share identical class labels in financial datasets, pseudo-labeling strategies can be applied across different temporal phases to accommodate distributional heterogeneity.

5.4 Comprehensive Comparison of Experimental Results

Based on the analyses corresponding to Figures 2, 3, and 4, the overall experimental performance is summarized in Table 1.

Table 1: Summary of Experimental Results Corresponding to Figures 2-4

Evaluation Dimension	Standalone RNN	Proposed Hybrid Model	Corresponding Figure
Short-term prediction	Effective	More stable	Figure 2
Long-term stability	Performance degrades	Slower degradation	Figure 2
Robustness under distribution shift	Sensitive under (L_{MSE})	More robust with cross-entropy	Figure 3
Risk threshold tracking	Limited dynamic capability	Clear threshold evolution	Figure 4
Practical applicability	Basic forecasting	Suitable for risk control and supervision	Figures 2-4

6. Conclusion

This study proposes a hybrid deep learning framework integrating a recurrent neural network and a latent autoregressive model for financial risk prediction. By modeling feature distribution characteristics in financial time series and incorporating cross-entropy-based distribution alignment, the framework enhances robustness under covariate shift and improves anomaly detection performance.

The experimental analyses in Figures 2-4 demonstrate that the proposed model outperforms standalone architectures in terms of long-term stability, distribution adaptability, and risk threshold interpretability, confirming its feasibility and practical value in financial risk management scenarios.

References

[1] V. Chandola, A. Banerjee and V. Kumar, "Anomaly detection: A survey," ACM Computing Surveys, vol. 41, no. 3, pp. 1-58, 2009.

[2] Y. Bengio, A. Courville and P. Vincent, "Representation learning: A review and new perspectives," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 35, no. 8, pp. 1798-1828, 2013.

[3] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," arXiv preprint arXiv:1312.6114, 2013.

[4] N. Chen, Y. Zhang, W. Wang, Z. Pan, Y. Wang and Y. Lu, "CoReAD: Context-aware retrieval-augmented deep anomaly detection for evolving business tabular data."

[5] N. Chen, S. Sun, Y. Wang, Z. Li, A. Zhu and Y. Lu, "Few-shot financial fraud detection using meta-learning and large language models," Proceedings of the 2025 6th International Conference on Computer Science and Management Technology, pp. 822-826, 2025.

[6] H. Chen, R. Wu, C. Chen, H. Feng, Y. Nie and Y. Lu, "Anomaly ranking for enterprise finance using latent structural deviations and reconstruction consistency," 2026.

[7] Y. Qin, D. Song, H. Chen, W. Cheng, G. Jiang and G. Cottrell, "A dual-stage attention-based recurrent neural network for time series prediction," arXiv preprint arXiv:1704.02971, 2017.

[8] H. Y. Kim and C. H. Won, "Forecasting the volatility of stock price index: A hybrid model integrating LSTM with multiple GARCH-type models," Expert Systems with Applications, vol. 103, pp. 25-37, 2018.

[9] B. Lim, S. Ö. Arık, N. Loeff and T. Pfister, "Temporal fusion transformers for interpretable multi-horizon time series forecasting," International Journal of Forecasting, vol. 37, no. 4, pp. 1748-1764, 2021.

[10] Y. Su, Y. Zhao, C. Niu, R. Liu, W. Sun and D. Pei, "Robust anomaly detection for multivariate time series through stochastic recurrent neural network," Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 2828-2837, 2019.

[11] X. Liang, Y. Zhao, M. Chang, R. Zhou, K. Cao and Y. Zheng, "Spatiotemporal risk representation learning using transformers and graph structure," 2026.

[12] C. Chen, R. Fang and J. Lai, "Causal representation learning for robust and interpretable audit risk identification in financial systems," Proceedings of the International Conference on Economic Management and Model Engineering, p. 454, Springer Nature, 2026.

[13] Q. Liu, H. Cui, Y. Wang, C. S. Lee, S. Chen and L. Yang, "A multi-agent large language model framework for marketing decision-making with auditable attribution analysis," 2026.

[14] X. Liang, Q. Liu, S. Chen, H. Qiu and H. Zhang, "Structured state representation and constraint-guided policy learning for intelligent business decision systems," 2026.

[15] R. Yan, Y. Shu, S. Sun, N. Chen, Y. Ou and Y. Zhao, "Adaptive federated learning for privacy-preserving modeling in heterogeneous financial environments," 2026.

[16] R. Yan, Y. Ou, S. Sun, N. Chen, K. Zhou and Y. Shu, "DualShiftNet: Joint class-imbalance and distribution-shift aware learning for business risk prediction," 2026.

[17] X. Yang, S. Sun, Y. Li, Y. Xing, M. Wang and Y. Wang, "CaliCausalRank: Calibrated multi-objective ad ranking with robust counterfactual utility optimization," arXiv preprint arXiv:2602.18786, 2026.