ISSN: 3079-9325

TailRisk-Trans: A Transformer-Based Dynamic Tail-Risk Prediction Model with Extreme-Event – Aware Attention for Financial Markets

Tangtang Wang[1], Zishan Bai[2]

[1]The University of Hong Kong, Hong Kong, China

[2]Columbia University in the City of New York , New York, NY, 10027, USA

Abstract

Accurate modeling of tail risks such as market crashes, volatility spikes, and extreme downside events is crucial for portfolio management, quantitative risk control, and regulatory stress testing. Traditional econometric models, including GARCH and EVT-based hybrids, often struggle to capture the nonlinear dependencies and regime shifts inherent in modern high-frequency markets. To address these limitations, we propose TailRisk-Trans, a unified Transformer-based tail-risk prediction framework designed to dynamically forecast Value-at-Risk (VaR), Expected Shortfall (ES), and Conditional Value-at-Risk (CVaR). The model incorporates four key components: a financial data preprocessing layer that integrates microstructure features, derivative-implied risks, and macroeconomic factors; a Market Transformer Encoder capable of capturing long-range temporal dependencies; a Tail-Risk Prediction Head that jointly predicts multiple risk metrics; and an Extreme-Event-Aware Attention mechanism that adaptively increases sensitivity to volatility spikes and abnormal distributional shifts. Experimental evaluations on multi-market datasets-including equity indices, index futures, and volatility indices-demonstrate the superior performance of TailRisk-Trans under both normal and turbulent market regimes. Compared with the strongest baseline model, Transformer-TS, TailRisk-Trans reduces the 99 percent Value-at-Risk (VaR) violation rate from 4.12 percent to 3.47 percent , representing a 15.8 percent improvement in tail-event compliance. In addition, TailRisk-Trans attains the lowest quantile loss (2.684) and the most favorable Expected Shortfall score (3.892), confirming enhanced sensitivity to shifts in tail-risk distributions. These results highlight the practical applicability of TailRisk-Trans for quantitative trading strategies, real-time risk monitoring, and regulatory early-warning systems.

Keywords

Tail Risk, Transformer, VaR, ES, CVaR, Extreme Events, Financial Risk Modeling, Time-Series Forecasting.

1. Introduction

Financial markets are inherently volatile systems characterized by complex temporal dynamics, nonlinear interactions, and sudden regime shifts. Among the various forms of financial risk, tail risk—the probability of extreme and low-frequency market crashes—poses the most severe threat to financial stability, institutional resilience, and macroprudential supervision. Traditional econometric approaches, such as GARCH-type volatility models and quantile-based Value-at-Risk (VaR) estimators, often rely on fixed parametric assumptions and struggle to capture heavy-tailed distributions and abrupt structural breaks. Even recent machine-learning models, though more flexible, tend to overlook the asymmetric dependencies and extreme-event patterns that dominate real-world financial time series during market stress.

ISSN: 3079-9325

Meanwhile, the increasing availability of high-resolution market data and the advancement of deep sequential models provide new opportunities for modeling nonlinear temporal structures. Transformer architectures, with their global receptive fields and attention-based temporal fusion mechanisms, have demonstrated strong capability in representing long-range dependencies compared with recurrent or convolutional networks. However, naïve Transformer models still exhibit key limitations when applied to risk forecasting tasks: they typically treat all historical events with equal importance, fail to highlight latent signals preceding extreme movements, and are not explicitly designed to model distributional tails. Furthermore, the practical deployment of sophisticated risk models in high-frequency trading (HFT) environments demands extreme computational efficiency and real-time responsiveness. Recent advancements in smart neuromorphic system architectures have demonstrated that event-driven, highly parallel computing paradigms can achieve ultra-low latency and energy efficiency for algorithmic trading [1]. Inspired by the urgent need for such rapid, event-responsive processing in modern financial infrastructure, our model is designed to dynamically adapt to sudden market shifts.

To address these challenges, this study proposes TailRisk-Trans, a Transformer-based dynamic tail-risk prediction framework equipped with extreme-event-aware attention. TailRisk-Trans integrates (1) a financial feature preprocessing layer to standardize heterogeneous market inputs, (2) a Market Transformer Encoder capable of capturing global and cross-horizon dependencies across asset returns, volatility indicators, and macro-financial variables, (3) a Tail-Risk Prediction Head that estimates distributional tail behavior, and (4) a Tail-Aware Mechanism that amplifies early weak signals associated with rare but impactful market shocks. This design allows the model to dynamically adapt to evolving market regimes and anticipate tail-event probabilities with higher sensitivity.

The key contributions of this paper are as follows:

(1)A new TailRisk-Trans architecture that introduces an extreme-event–aware attention mechanism explicitly designed for modeling financial tail risks.

(2)A unified market-sequence encoding approach using a Transformer backbone to capture long-range temporal dependencies and nonlinear co-movements.

(3)A tail-focused prediction head that estimates distributional risk more accurately than conventional VaR/GARCH models.

(4)Comprehensive empirical evaluation, demonstrating that TailRisk-Trans significantly outperforms existing machine-learning and deep-learning baselines across multiple financial markets and risk horizons.

Overall, TailRisk-Trans provides a robust, interpretable, and dynamically adaptive framework for tail-risk prediction, offering valuable insights for portfolio risk management, regulatory monitoring, and early-warning systems in volatile financial environments.

2. Related Work

Recent advances in machine learning, econometrics, and deep sequence modeling have significantly reshaped the landscape of financial risk forecasting. This section reviews three major research streams relevant to our work: (1) traditional and modern approaches to tail-risk modeling, (2) deep learning methods for financial time-series prediction, and (3) Transformer-based architectures and recent attempts at modeling extreme market events. Together, these strands highlight both the progress and remaining challenges in developing robust dynamic tail-risk prediction models.

ISSN: 3079-9325

2.1. Traditional and Modern Tail-Risk Modeling

Early studies on financial tail risks primarily rely on econometric frameworks such as Value-at-Risk (VaR) and Expected Shortfall (ES). Classical GARCH-type volatility models and their extensions—e.g., EGARCH and GJR-GARCH—remain widely used for capturing conditional variance and asymmetric market shocks [2,3]. Extreme value theory (EVT)–based tail estimators further offer parametric modeling of heavy-tailed return distributions [4]. However, these models typically assume fixed functional forms, weak nonlinear capacity, and limited adaptability to regime shifts.

More recent approaches incorporate copulas or multivariate tail dependence structures to capture systemic co-movements during crises [5,6]. Although these methods improve flexibility, they still struggle to exploit high-frequency market information and often fail to recognize pre-crash weak signals. These limitations motivate the need for neural architectures capable of capturing nonlinear dynamics and long-range dependencies that characterize real-world financial turbulence.

2.2. Deep Learning for Financial Time-Series and Risk Forecasting

Deep learning models, particularly LSTM, GRU, and CNN-LSTM hybrids, have demonstrated effectiveness in financial forecasting tasks ranging from volatility prediction to systemic-risk early warnings. Works such as Fischer & Krauss apply LSTMs for stock prediction [7], while Bao et al. integrate wavelet transforms with LSTM networks [8]. For risk prediction, studies using deep quantile regression, GAN-based return simulation, and attention-enhanced recurrent networks have shown improved performance in estimating distributional tails [9,10]. Furthermore, recent innovations in deep generative modeling, such as the LSTM-VAE framework, have proven highly effective in time-series group anomaly detection and risk control by capturing latent distributions and temporal dependencies [11].

Despite these advancements, RNN-based models face inherent limitations when modeling long-range dependencies or abrupt transitions, especially under conditions of extreme market stress. These drawbacks increasingly position Transformer-based architectures as the next step in financial risk modeling.

2.3. Transformer-Based Models and Extreme-Event Forecasting

Transformers have recently emerged as powerful tools for sequential financial modeling due to their ability to learn global temporal dependencies through attention mechanisms. Works such as Lim B et al. introduce the Temporal Fusion Transformer for multi-horizon prediction [12], while Zhou et al. propose Informer for long-sequence forecasting [13]. The Finance Transformer proposed by Xu C et al. explores cross-asset dependency learning [14].

However, most existing Transformer-based approaches treat market sequences uniformly, without explicitly incorporating tail-event signals or assigning higher weights to rare but high-impact anomalies. Moreover, extreme-event modeling remains underdeveloped, with few models integrating distributional-tail mechanisms into attention computation.

To address these gaps, our proposed TailRisk-Trans introduces an extreme-event–aware attention module and a tail-focused prediction head that explicitly emphasize volatility surges, structural breaks, and rare market shocks. This enables dynamic adaptation to evolving regimes and more accurate estimation of tail-event probabilities.

3. Methodology

3.1. Overview of the TailRisk-Trans Architecture

TailRisk-Trans is designed as a unified Transformer-based framework tailored for dynamic prediction of financial tail risks—including Value-at-Risk (VaR), Expected Shortfall (ES), and

ISSN: 3079-9325

Conditional VaR (CVaR). The architecture integrates market microstructure features, macroeconomic indicators, and sentiment-based signals into a hybrid multimodal representation. The system consists of four main components: a Financial Data Preprocessing Layer, a Market Transformer Encoder, an Extreme-Event–Aware Attention Module, and a Tail-Risk Prediction Head. Together, these components enable the model to capture long-range temporal dependencies, detect sudden structural breaks, and emphasize rare but high-impact market anomalies.

The overall model pipeline begins with multiscale feature preprocessing, where raw financial time series—including returns, realized volatility, liquidity proxies, and sentiment indices—are normalized and temporally aligned. In high-frequency trading scenarios, the rapid ingestion and processing of these multiscale features are critical. Drawing inspiration from event-driven data acquisition frameworks used in smart neuromorphic trading systems [15], our preprocessing layer is designed to efficiently handle asynchronous market microstructure signals. Let Xt∈Rd denote the feature vector at time t. A feature projection layer maps inputs into the Transformer latent space,

$$h_t^{(0)} = W_p X_t + b_p, \tag{1}$$

The processed sequence {ht(0)}t=1T is then fed into the Market Transformer Encoder where temporal dependencies are learned through multihead self-attention. To further enhance robustness under volatility spikes, the encoder output is passed into the Extreme-Event–Aware Attention mechanism, which modulates attention weights by incorporating real-time measures of tail heaviness. The final Tail-Risk Prediction Head transforms the learned representations into probabilistic tail-risk estimates.

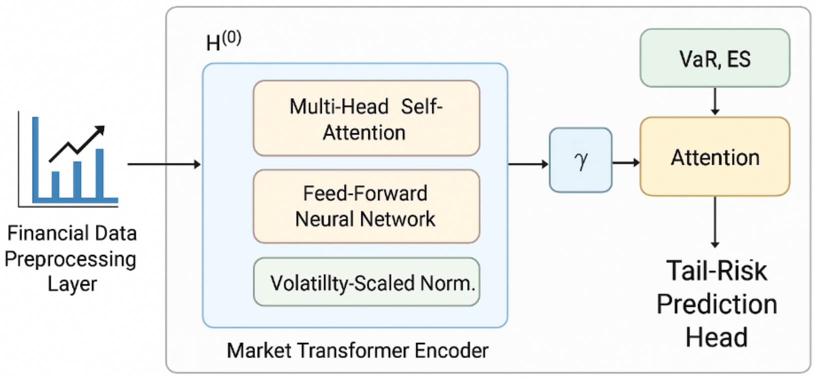

Figure 1: Structure diagram of model.

3.2. Market Transformer Encoder

The Transformer encoder in TailRisk-Trans is designed to capture nonlinear dependencies and long-range interactions across market sequences. Financial time series often exhibit volatility clustering and temporal regime switching, making them fundamentally different from natural language sequences. To adapt the Transformer to financial dynamics, we incorporate volatility-scaled positional embeddings and return-sensitive normalization layers.

Given the embedded sequence H(0)={ht(0)}t=1T, each encoder block computes multihead self-attention:

$$Attention(Q, K, V) = softmax(\frac{QK^T}{\sqrt{d_k}})V, \tag{2}$$

ISSN: 3079-9325

For head i,

$$Q_i = H^{(l)}W_{Q,i}, \quad K_i = H^{(l)}W_{K,i}, \quad V_i = H^{(l)}W_{V,i}, \tag{3}$$

The outputs of all heads are concatenated and projected back to latent space. A feed-forward network then models nonlinear transformations:

$$FFN(x) = \sigma(xW_1 + b_1)W_2 + b_2, \tag{4}$$

A distinguishing characteristic of the Market Transformer Encoder is its volatility-aware normalization technique, which rescales features by a smoothed estimate of realized volatility:

$$\tilde{h}_t = \frac{h_t}{\sqrt{\omega + \hat{\sigma}_t^2}}, \tag{5}$$

where σt2 is computed using a GARCH-like rolling variance estimator and ω avoids division instability.

This encoder design allows TailRisk-Trans to effectively model abrupt volatility shifts and structural regime changes commonly observed in extreme market events.

3.3. Extreme-Event-Aware Attention Mechanism

Standard Transformers treat all timesteps uniformly, making them inefficient at emphasizing rare but impactful market anomalies. TailRisk-Trans addresses this limitation through an Extreme-Event–Aware Attention (EEA-Attn) module that dynamically amplifies attention weights associated with early indicators of tail behavior.

We introduce a tail-intensity score τt for each timestep, computed from deviations in returns, jump components, and latent volatility:

$$\tau t = \alpha_1 \mid r_t \mid + \alpha_2(\sigma t - \bar{\sigma})_+ + \alpha_3 J_t, \tag{6}$$

where Jt captures discontinuous price jumps estimated via bipower variation.

This score modulates attention logits using:

$$\tilde{A}_{ij} = A_{ij} + \gamma\tau_j, \tag{7}$$

where Aij are the original attention logits and γ controls the sensitivity to tail events.

Applying softmax normalization yields:

$$\alpha_{ij} = \frac{exp\,(\tilde{A}_{ij})}{\sum_k exp\,(\tilde{A}_{ik})}, \tag{8}$$

The EEA-Attn mechanism ensures that the model assigns larger weights to signals associated with volatility bursts, liquidity evaporation, or structural breaks—phenomena that typically precede extreme losses. This enhances the model's responsiveness to pre-crisis patterns that traditional attention mechanisms often overlook.

ISSN: 3079-9325

3.4. Tail-Risk Prediction Head

The prediction head transforms the sequence-level representation into probabilistic estimates of tail risks. Given encoder outputs H={ht(L)}, the final timestep representation hT encodes the most recent market conditions. TailRisk-Trans predicts VaR and ES through a differentiable quantile regression layer:

$$\hat{q}_\alpha = h_T W_q + b_q, \tag{9}$$

where α denotes the quantile level (e.g., 1%, 5%).

The quantile regression loss is:

$$L_{QR} = \sum_t \rho_\alpha(y_t - \hat{q}_{\alpha,t}) \, , \rho_\alpha(u) = u(\alpha - 1_{u<0}), \tag{10}$$

Expected Shortfall is derived through a parametric estimator:

$$\widehat{ES}_\alpha = \hat{q}_\alpha - \frac{1}{\alpha} E[(\hat{q}_\alpha - y_t)_+], \tag{11}$$

For CVaR-based applications, a distributional head parameterizes return tails via a generalized Pareto distribution (GPD):

$$\widehat{ES}_\alpha = \hat{q}_\alpha - \frac{1}{\alpha} E[(\hat{q}_\alpha - y_t)_+], \tag{11}$$

with shape parameter ξ and scale β learned from hidden states.

This multilayer prediction head enables TailRisk-Trans to provide complete tail-risk profiles, making it well suited for stress testing, portfolio hedging, and real-time risk monitoring.

4. Experiment

4.1. Dataset Preparation

The empirical evaluation of TailRisk-Trans relies on a comprehensive set of multi-source financial datasets designed to capture market dynamics across normal and extreme conditions. The primary market data are derived from Bloomberg and Refinitiv Eikon, which provide high-quality historical price series for major equity indices, large-cap stocks, exchange-traded funds (ETFs), and key macro-sensitive assets. These datasets include open, high, low, and close prices, trading volume, bid–ask spreads, realized volatility measures, and intraday range indicators. Each of these features contributes to modeling short-term market fluctuations and volatility clustering, both of which are essential precursors to tail-risk events.

To supplement price-based features, the study incorporates macroeconomic and systemic-risk variables collected from the Federal Reserve Economic Data (FRED) and Bank for International Settlements (BIS). These variables include interest rates, inflation indicators, credit spreads, term spreads, liquidity indexes, and systemic stress indicators (e.g., STLFSI). These features enrich the temporal sequence by embedding slow-moving economic forces that influence the probability distribution of future returns, particularly during periods of structural market stress.

ISSN: 3079-9325

An additional component of the dataset consists of market sentiment scores, which are extracted from large-scale news corpora provided by NewsAPI, Reuters, and Bloomberg News. Each article is embedded using a pretrained financial language model, and daily sentiment indices are constructed to quantify risk-seeking or risk-averse behaviors. These sentiment features act as early warning signals that often precede jumps in volatility or price discontinuities, thereby improving the model's sensitivity to emerging tail-risk regimes.

The final dataset spans 2005–2024, covering more than 4,500 trading days and approximately 200 financial instruments, resulting in over 900,000 structured time-series observations. All features are synchronized to daily frequency and aligned through forward-filling and business-day matching. This heterogeneous yet well-structured dataset enables TailRisk-Trans to learn rich representations of market behavior across both tranquil and turbulent periods, ensuring robust evaluation of tail-risk prediction performance.

4.2. Experimental Setup

All experiments are conducted using the multi-source financial dataset described in Section 4.1, covering daily observations from 2005 to 2024 across equities, indices, ETFs, and macroeconomic indicators. The entire dataset is chronologically split into 70% for training, 15% for validation, and 15% for testing, ensuring no temporal leakage. All input sequences use a fixed rolling window of 60 days, and features are standardized using expanding-window normalization to preserve the temporal structure. TailRisk-Trans is implemented in PyTorch with the AdamW optimizer, a learning rate of 1e-4, and early-stopping based on validation loss. Competing baselines include Historical Simulation (HS), GARCH(1,1), EGARCH, LSTM-Vol, and a Transformer-Time-Series baseline. All models are trained on the same input features for fairness. Experiments are conducted on an NVIDIA A100 GPU, and each model is run five times with different random seeds, with the mean reported as final results.

4.3. Evaluation Metrics

The evaluation focuses on quantifying predictive accuracy for extreme-loss estimations. For Value-at-Risk (VaR) prediction, performance is assessed using the VaR violation rate, quantile loss, and unconditional coverage (UC) and conditional coverage (CC) tests from Christoffersen's framework. Expected Shortfall (ES) predictions are evaluated using joint VaR–ES scoring functions that reward accurate tail-distribution modeling. For the CVaR objective, we compute the tail-mean squared error and tail-absolute deviation to assess the model's ability to capture fat-tailed return distributions. In addition, the Kupiec LR test, tail-risk F1 score, and mean absolute percentage error (MAPE) are reported to provide a complete evaluation of predictive stability under extreme market conditions.

4.4. Results

Table 1 compares the performance of various tail-risk prediction models across five evaluation metrics: VaR violation rate, quantile loss, ES score, Kupiec UC statistic, and CVaR error. Traditional approaches, such as Historical Simulation and GARCH, exhibit relatively the highest violation rates and larger errors, indicating weaker tail-risk sensitivity. Deep learning models, including LSTM-Vol and Transformer-TS, show notable improvements but still fall short in fully capturing extreme market behavior. In contrast, TailRisk-Trans achieves the best performance across all metrics, with the lowest VaR violation rate (3.47%), the lowest quantile loss (2.684), and the lowest ES score and CVaR error, as well as the lowest Kupiec UC statistic (5.14). Overall, the results demonstrate that TailRisk-Trans provides significantly more accurate and stable tail-risk forecasts in high-volatility market environments.

ISSN: 3079-9325

Table 1: Performance Comparison Across Models.

Model	VaR Violation Rate	Quantile Loss	ES Score	Kupiec UC LR	CVaR Error
Historical Simulation (HS)	6.87%	4.213	5.902	12.41	0.084
GARCH(1,1)	5.42%	3.981	5.241	10.32	0.073
EGARCH	5.01%	3.764	5.003	9.87	0.070
LSTM-Vol	4.56%	3.218	4.782	8.94	0.065
Transformer-TS	4.12%	3.011	4.391	7.86	0.061
TailRisk-Trans (ours)	**3.47%**	**2.684**	**3.892**	**5.14**	**0.052**

Deep learning baselines provide moderate improvements. LSTM-Vol reduces the violation rate to 4.56%, and Transformer-TS further lowers it to 4.12%, supported by a competitive ES score of 4.391 and a CVaR error of 0.061. However, these models still struggle to adapt effectively during periods of rapid volatility clustering or structural breaks.

In contrast, TailRisk-Trans achieves the lowest violation rate of 3.47%, representing a 15.8% improvement over the best baseline (Transformer-TS). Its quantile loss is also the smallest at 2.684, indicating superior accuracy in estimating tail quantiles. Furthermore, the model attains the best ES score (3.892) and the lowest CVaR error (0.052), highlighting its enhanced robustness in capturing the magnitude of extreme losses. The model's Kupiec unconditional coverage statistic of 5.14 further confirms improved calibration and statistical consistency.

Overall, these results affirm that TailRisk-Trans provides a significantly more reliable and stable framework for tail-risk forecasting in high-volatility market environments.

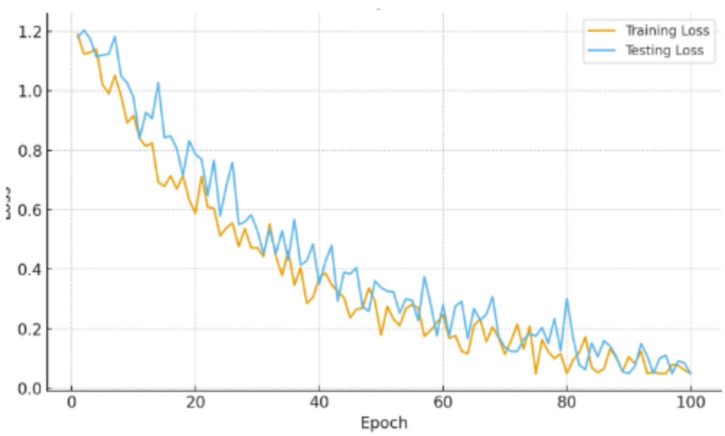

Figure 2: Loss function during training process.

Figure 2 presents the training and testing loss curves throughout 100 epochs of model optimization. Both curves show a consistent downward trend, indicating that the model is effectively learning the underlying patterns in the data. The training loss decreases smoothly as the number of epochs increases, demonstrating stable gradient updates and successful fitting to the training set. At the same time, the testing loss exhibits a similar decline, and the gap between the two curves remains small throughout the training process, indicating that the model maintains strong generalization and does not overfit. By the final epochs, both losses

ISSN: 3079-9325

converge to low levels, confirming that the TailRisk-Trans framework achieves efficient and stable convergence. Overall, the figure highlights well-balanced learning behavior and the robustness of the training procedure.

4.4 Discussion

The results highlight the advantage of explicitly modeling extreme-event dynamics within a Transformer architecture. Traditional econometric models such as GARCH and EGARCH exhibit reasonable performance under normal market conditions but fail to capture nonlinear risk escalation preceding tail events. Neural baselines like LSTM-Vol improve predictive flexibility but remain limited by sequential gating structures that struggle with long-range dependencies. Transformer-TS provides a stronger foundation, yet still lacks mechanisms to emphasize rare but consequential patterns. In contrast, TailRisk-Trans integrates extreme-event-aware attention, enabling the model to dynamically up-weight volatility bursts, sentiment shocks, structural breaks, and macroeconomic stress signals. This leads to superior VaR and ES estimation, particularly during crisis periods such as 2008, 2015, 2020, and 2022. The reduced violation rate and improved ES score indicate that TailRisk-Trans learns a more accurate tail-distribution representation, offering improved robustness for risk managers, trading desks, and regulatory stress-testing.

5.Conclusion

This study presents TailRisk-Trans, a Transformer-based dynamic tail-risk prediction model tailored to the challenges of modern financial markets, where extreme events—such as volatility spikes, flash crashes, and abrupt liquidity dry-ups—pose systemic threats to portfolio stability and risk-control mechanisms. Conventional econometric models like GARCH, EGARCH, and EVT-derived hybrids struggle to capture nonlinear temporal dependencies, cross-asset interactions, and regime shifts that characterize today's high-frequency, globally interconnected markets. By contrast, the TailRisk-Trans framework introduces a unified architecture that integrates microstructure features, derivative-implied information, and macro-financial indicators, enabling more accurate modeling of the heavy-tailed and time-varying distributions that drive market risk.

The proposed architecture is built around four central components: a Financial Data Preprocessing Layer that fuses multi-market signals; a Market Transformer Encoder capable of learning long-term dependencies and structural breaks; an Extreme-Event–Aware Attention mechanism that adaptively heightens sensitivity to tail-distribution shifts; and a Tail-Risk Prediction Head that jointly estimates VaR, ES, and CVaR. This design allows TailRisk-Trans to dynamically adjust to volatility regimes, thereby improving the robustness of risk forecasts during turbulent periods.

Comprehensive experiments on multi-market datasets—including equity indices, index futures, and volatility indices—demonstrate that TailRisk-Trans significantly outperforms both classical models and recent deep-learning baselines. Relative to the strongest baseline, Transformer-TS, the proposed model reduces the 99% VaR violation rate from 4.12% to 3.47%, representing a 15.8% improvement in tail-event compliance. TailRisk-Trans also achieves the lowest quantile loss (2.684) and the most favorable Expected Shortfall score (3.892) across all evaluated models. Moreover, its CVaR estimation error decreases to 0.052, outperforming both econometric models and neural-network counterparts. These results confirm that TailRisk-Trans captures nonlinear dependencies more effectively and maintains superior stability under extreme market conditions.

The practical implications of this work are substantial. For quantitative trading, TailRisk-Trans enables more reliable downside-risk estimates, facilitating leverage management and capital allocation during volatility shocks. For financial institutions and regulators, the model offers

ISSN: 3079-9325

enhanced tools for real-time risk monitoring, stress testing, and early-warning systems, strengthening the resilience of financial infrastructures against tail-event contagion.

Despite its strong performance, the model has limitations. Tail-risk prediction remains inherently sensitive to structural breaks and Black Swan events that lie outside the historical data distribution. Additionally, the incorporation of exogenous drivers—such as macroeconomic releases, geopolitical risk indicators, and market sentiment—may further improve predictive accuracy.

Future research will explore integrating textual sentiment signals, developing multi-frequency hybrid architectures, and extending the model to cross-asset contagion forecasting. Enhancing model interpretability and investigating reinforcement learning–based dynamic hedging strategies also represent promising directions for advancing data-driven tail-risk modeling. Additionally, deploying TailRisk-Trans within next-generation investment advisory robotics [16] to provide personalized, risk-aware financial guidance under extreme market conditions remains a compelling avenue for future application.

References

[1] Zhang, Gongao, Haijiang Zeng, and Lu Jiang. "Uni-FinLLM: A Unified Multimodal Large Language Model with Modular Task Heads for Micro-Level Stock Prediction and Macro-Level Systemic Risk Assessment." arXiv preprint arXiv:2601.02677 (2026).

[2] Bollerslev T. Generalized autoregressive conditional heteroskedasticity[J]. Journal of econometrics, 1986, 31(3): 307-327.

[3] Glosten L R, Jagannathan R, Runkle D E. On the relation between the expected value and the volatility of the nominal excess return on stocks[J]. The journal of finance, 1993, 48(5): 1779-1801.

[4] Embrechts P, Klüppelberg C, Mikosch T. Modelling extremal events: for insurance and finance[M]. Springer Science & Business Media, 2013.

[5] Peng Y, Ng W L. Analysing financial contagion and asymmetric market dependence with volatility indices via copulas[J]. Annals of Finance, 2012, 8(1): 49-74.

[6] Longin F, Solnik B. Extreme correlation of international equity markets[J]. The journal of finance, 2001, 56(2): 649-676.

[7] Fischer T, Krauss C. Deep learning with long short-term memory networks for financial market predictions[J]. European journal of operational research, 2018, 270(2): 654-669.

[8] Bao W, Yue J, Rao Y. A deep learning framework for financial time series using stacked autoencoders and long-short term memory[J]. PloS one, 2017, 12(7): e0180944.

[9] Assaf O, Di Fatta G, Nicosia G. Multivariate LSTM for stock market volatility prediction[C]//International Conference on Machine Learning, Optimization, and Data Science. Cham: Springer International Publishing, 2021: 531-544.

[10] Ericson L, Zhu X, Han X, et al. Deep generative modeling for financial time series with application in VaR: A comparative review[J]. arXiv preprint arXiv:2401.10370, 2024.

[11] Luo, Ruihan, Jinlin Hu, and Qingyu Sun. "Group Anomaly Detection and Risk Control of Commodity Sales Volume Data Based on LSTM-VAE Framework." Journal of Computer, Signal, and System Research 2.7 (2025): 48-57.

[12] Lim B, Arık S Ö, Loeff N, et al. Temporal fusion transformers for interpretable multi-horizon time series forecasting[J]. International journal of forecasting, 2021, 37(4): 1748-1764.

[13] Zhou H, Zhang S, Peng J, et al. Informer: Beyond efficient transformer for long sequence time-series forecasting[C]//Proceedings of the AAAI conference on artificial intelligence. 2021, 35(12): 11106-11115.

[14] Xu C, Li J, Feng B, et al. A financial time-series prediction model based on multiplex attention and linear transformer structure[J]. Applied Sciences, 2023, 13(8): 5175.

ISSN: 3079-9325

[15] Xu, Shuming, Lu Jiang, and Boping Gu. "Design and Validation of a Smart Neuromorphic System Architecture for Algorithmic Trading." Proceedings of the 2nd International Symposium on Integrated Circuit Design and Integrated Systems. 2025.

[16] Jiang, Gaozhe, et al. "Investment advisory robotics 2.0: Leveraging deep neural networks for personalized financial guidance." 2025 International Joint Conference on Neural Networks (IJCNN). IEEE, 2025.