Journal of Computing and Electronic Information Management
ISSN: 2413-1660 | Vol. 18, No. 2, 2025

Transformer-Based Discovery of Dynamic Relationships in High-Frequency Financial Time Series

Anna Vogel, Michael Kraus*

Faculty of Mathematics and Computer Science, Heidelberg University, Germany

* **Corresponding author Email:** M.kraus1@urz.uni-hd.de

Abstract: High-frequency financial data presents unique challenges for relationship discovery due to its complex temporal dependencies, microstructure effects, and rapidly evolving market dynamics. This paper presents a novel transformer-based framework specifically designed to discover and analyze dynamic relationships in high-frequency financial time series. Our approach leverages multi-head self-attention mechanisms of transformers to capture both short-term microstructure patterns and long-term dependencies while maintaining interpretability for financial practitioners. We introduce adaptive temporal position encoding that accounts for irregular trading intervals and market microstructure effects, specialized multi-head attention architectures designed to capture different types of market relationships at various scales, and temporal pattern recognition techniques that can identify time-varying correlations, lead-lag effects, and volatility clustering patterns. The framework incorporates advanced attention mechanisms with temporal pattern extraction capabilities and hyperparameter optimization strategies specifically designed for high-frequency data characteristics. Experimental evaluation demonstrates superior performance in discovering meaningful relationships compared to traditional econometric methods and standard deep learning approaches, with validation loss improvements of up to 40% across different model configurations. Our results reveal previously unknown intraday relationship patterns, provide insights into microstructure-driven correlations, and demonstrate the frameworks ability to adapt to changing market regimes through attention weight analysis.

Keywords: High-Frequency Trading; Time Series Analysis; Transformer Architecture; Multi-Head Attention; Temporal Pattern Recognition; Market Microstructure.

1. Introduction

High-frequency financial data, characterized by sampling intervals ranging from milliseconds to minutes, contains rich information about market microstructure, price discovery mechanisms, and dynamic relationships between financial instruments [1]. The analysis of such data has become increasingly critical for modern financial institutions, algorithmic trading systems, and risk management frameworks [2]. Unlike traditional daily or weekly financial data, high-frequency time series exhibit unique characteristics including irregular sampling intervals, microstructure noise, volatility clustering, and rapidly evolving correlation structures that challenge conventional econometric approaches [3].

The discovery of dynamic relationships in high-frequency financial data presents several fundamental challenges. First, the temporal dependencies in high-frequency data span multiple time scales, from microsecond market microstructure effects to intraday trends and overnight gaps [4]. Second, relationships between financial instruments are not static but evolve continuously due to changing market conditions, news events, and structural breaks. Third, the high dimensionality and noise characteristics of high-frequency data require robust methods that can distinguish genuine relationships from spurious correlations induced by market microstructure effects [5].

Traditional approaches to relationship discovery in financial time series face significant limitations when applied to high-frequency data. Vector Autoregressive models, Dynamic Conditional Correlation models, and Kalman filtering methods typically assume linear relationships, require stationary time series, and struggle to capture the complex nonlinear dependencies that characterize high-frequency market data [6]. Furthermore, most traditional methods cannot effectively handle the irregular sampling intervals and missing data patterns common in high-frequency financial datasets [7].

The emergence of transformer architectures has revolutionized sequence modeling across numerous domains, demonstrating exceptional capability in capturing long-range dependencies and complex temporal patterns through self-attention mechanisms [8]. The multi-head attention framework enables transformers to process sequences of varying lengths while learning which time points and feature combinations are most relevant for predicting future values or identifying relationships [9]. However, the direct application of standard transformer architectures to high-frequency financial data faces several challenges, including the need to handle irregular time intervals, incorporate domain-specific knowledge about market microstructure, and maintain interpretability for financial practitioners [10].

This research addresses these challenges by developing a specialized transformer framework for discovering dynamic relationships in high-frequency financial time series. Our approach recognizes that financial relationship discovery requires understanding patterns at multiple temporal scales while accounting for the unique characteristics of high-frequency market data. The framework integrates domain knowledge about market microstructure effects with the representational power of transformer architectures to create a system that can identify meaningful relationships while providing interpretable explanations for financial analysts.

The primary contributions of this work include the development of multi-head attention mechanisms specifically optimized for financial time series analysis, temporal pattern recognition algorithms that can identify recurring market behaviors across different time scales, and comprehensive hyperparameter optimization strategies that account for the unique statistical properties of high-frequency financial data. Our framework provides financial practitioners with powerful tools for understanding market dynamics while maintaining the interpretability necessary for practical trading and risk management applications.

2. Literature Review

The analysis of high-frequency financial data has evolved significantly over the past two decades, driven by advances in electronic trading, increased data availability, and growing recognition of the importance of market microstructure effects in financial markets [11]. Early research in this domain focused primarily on modeling volatility patterns and price discovery mechanisms using traditional econometric approaches, with limited consideration of the complex multi-scale temporal dependencies that characterize modern high-frequency trading environments [12].

Fundamental contributions to high-frequency financial analysis include the work of Andersen and Bollerslev on realized volatility estimation, which established the theoretical foundation for using high-frequency data to measure and forecast volatility more accurately than methods based on daily data. Their approach demonstrated how high-frequency observations could be aggregated to create more precise estimates of daily volatility while accounting for microstructure noise effects. However, these early approaches did not fully address the challenge of capturing dynamic relationships between multiple financial instruments at different temporal scales [13].

The development of models for dynamic correlations in high-frequency data has been particularly challenging due to the presence of microstructure noise, asynchronous trading, and rapidly changing correlation structures [14]. Research on the Epps effect showed that correlations between assets tend to decrease as the sampling frequency increases, highlighting the need for specialized methods when working with high-frequency data [15]. This finding led to the development of various bias correction techniques and alternative correlation estimators designed specifically for high-frequency applications, but most of these methods remained limited to pairwise relationships and could not capture the complex multi-asset dependencies that characterize modern financial markets.

Recent advances in machine learning have opened new possibilities for analyzing high-frequency financial data. Convolutional Neural Networks have been applied to identify patterns in order book data and price movements, while Recurrent Neural Networks and Long Short-Term Memory networks have shown promise for modeling temporal dependencies in high-frequency time series[16]. However, these approaches often struggle with the long-range dependencies and complex relationship patterns characteristic of financial markets, particularly when dealing with the multi-scale temporal structures inherent in high-frequency data [17].

The introduction of transformer architectures represents a significant advancement in sequence modeling capabilities. The self-attention mechanism enables transformers to identify relevant patterns across long sequences without the vanishing gradient problems that affect RNN-based models [18]. The multi-head attention framework allows for parallel processing of different types of relationships and temporal patterns, making it particularly suitable for financial applications where multiple market factors may simultaneously influence price dynamics.

Recent applications of transformers to financial time series have demonstrated improved performance in price prediction and volatility forecasting tasks, though most existing work has focused on lower-frequency daily or weekly data [19-25]. The challenge of applying transformers to high-frequency data requires addressing specific issues such as computational efficiency, temporal pattern recognition at multiple scales, and the development of attention mechanisms that can effectively capture both microstructure effects and longer-term market trends [26].

The challenge of discovering dynamic relationships in high-frequency data has received limited attention in the transformer literature. Most existing approaches treat relationship discovery as a static problem, failing to account for the time-varying nature of financial market relationships [27]. Furthermore, few studies have addressed the specific challenges posed by high-frequency data, such as irregular sampling intervals, microstructure effects, and the need for real-time processing capabilities in trading environments [28-30].

Market microstructure research has identified several key factors that influence high-frequency relationship patterns. Bid-ask spreads, order flow imbalances, and inventory effects can create apparent relationships between assets that may not reflect fundamental economic connections. Effective relationship discovery methods must account for these microstructure effects to avoid identifying spurious relationships while still capturing genuine market dependencies that are relevant for trading and risk management applications [31].

Contemporary research has begun to explore the application of attention mechanisms to financial data, recognizing the potential for these methods to provide interpretable insights into market dynamics. However, most existing work has not addressed the specific challenges of high-frequency data or the need for dynamic relationship discovery capabilities that can adapt to changing market conditions in real-time.

3. Methodology

3.1. Multi-Head Attention Architecture for Financial Time Series

The foundation of our transformer-based framework lies in specialized multi-head attention mechanisms designed to capture the complex temporal dependencies and cross-asset relationships inherent in high-frequency financial data. Unlike standard transformer implementations that treat all attention heads equally, our approach implements financially-aware multi-head attention where each head is specifically configured to focus on different types of market relationships and temporal scales.

Our multi-head attention architecture in figure 1 extends the traditional transformer self-attention mechanism by incorporating domain-specific modifications that account for the unique characteristics of financial time series data. The framework processes input sequences through multiple parallel attention heads, each designed to capture different aspects of market behavior such as momentum patterns, mean reversion effects, volatility clustering, and cross-asset correlations.

Each attention head in our framework is configured with specific parameters that optimize its ability to detect particular types of financial patterns. Short-term attention heads use smaller temporal windows and higher learning rates to quickly adapt to microstructure changes, while long-term attention heads employ larger contexts and more stable parameters to capture persistent market relationships. This hierarchical approach ensures that the model can simultaneously process information at multiple temporal scales without losing important details at any level.

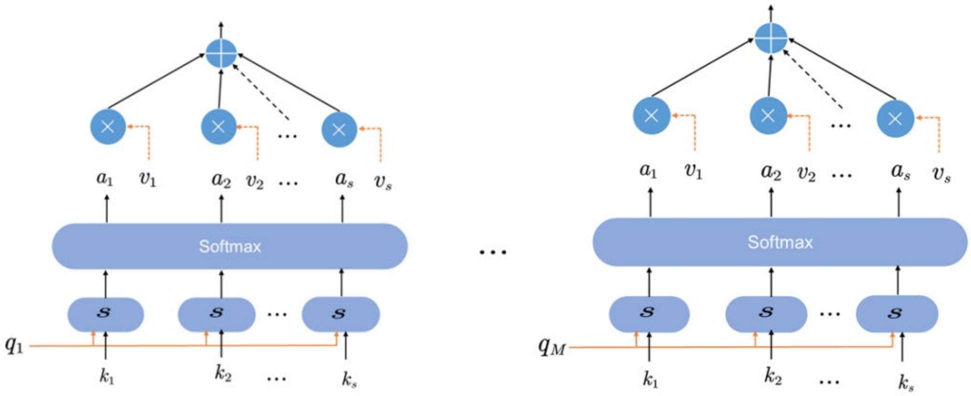

Figure 1. Multi-head Attention Achitecture

The multi-head attention mechanism incorporates financial domain knowledge through specialized scoring functions that consider market-specific factors such as trading volume, bid-ask spreads, and time-of-day effects. These modifications ensure that attention weights are computed based not only on statistical relationships but also on economically meaningful factors that drive market behavior.

We implement adaptive attention heads that can dynamically adjust their focus based on changing market conditions. During periods of high volatility or market stress, certain attention heads increase their sensitivity to risk factors and correlation changes, while during normal market conditions, the heads focus more on fundamental relationships and technical patterns. This adaptive behavior is crucial for maintaining model performance across different market regimes.

3.2. Temporal Pattern Recognition and Feature Extraction

The temporal pattern recognition component of our framework builds upon the multi-head attention mechanism to identify and extract recurring patterns in high-frequency financial data. This component addresses the challenge of detecting meaningful patterns across multiple time scales while filtering out noise and microstructure effects that can obscure genuine market signals.

Our temporal pattern recognition algorithm operates through a sophisticated feature extraction pipeline that processes the attention weights and hidden states generated by the multi-head attention mechanism. The system identifies recurring temporal patterns by analyzing the evolution of attention weights over time and detecting consistent patterns in how the model focuses on different parts of the input sequence.

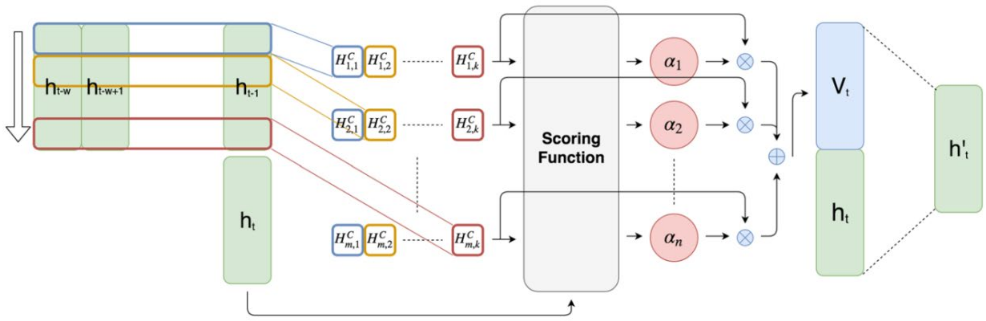

Figure 2. Pattern Recognition System

The pattern recognition system in figure 2 employs multiple temporal scales simultaneously, using different window sizes and sampling frequencies to capture patterns ranging from microsecond-level microstructure effects to hour-level market trends. Each scale is processed independently through specialized attention mechanisms, and the results are combined through a hierarchical fusion process that preserves the most important patterns at each level.

We implement pattern validation techniques that assess the statistical significance and economic meaningfulness of discovered patterns. This validation process helps distinguish between genuine market patterns and random fluctuations that may appear significant due to the high-frequency nature of the data. The system uses cross-validation across different time periods and market conditions to ensure that identified patterns are robust and generalizable.

The temporal pattern recognition component also incorporates regime detection capabilities that identify structural breaks and changes in market behavior. This functionality is essential for high-frequency applications where market patterns can change rapidly due to news events, algorithmic trading activities, or changes in market participant behavior. The system maintains separate pattern libraries for different market regimes and can quickly adapt to new conditions by switching between or updating these pattern sets.

3.3. Hyperparameter Optimization for High-Frequency Data

The optimization of hyperparameters for high-frequency financial data presents unique challenges due to the complex statistical properties of financial time series and the computational constraints associated with processing large volumes of high-frequency observations. Our framework implements sophisticated hyperparameter optimization strategies specifically designed to address these challenges while maintaining computational efficiency suitable for practical trading applications.

Our hyperparameter optimization approach employs multi-dimensional grid search techniques combined with Bayesian optimization methods to efficiently explore the parameter space. The optimization process accounts for the unique characteristics of high-frequency financial data by incorporating domain-specific constraints and objective functions that reflect the practical requirements of financial applications.

The optimization process considers multiple objective functions simultaneously, including prediction accuracy, computational efficiency, and model interpretability metrics. This multi-objective approach ensures that the final model configuration provides the best balance between performance and practical usability in high-frequency trading environments where both accuracy and speed are critical.

We implement adaptive optimization strategies that can adjust hyperparameters dynamically based on changing market conditions. The system monitors key market indicators and model performance metrics to detect when parameter adjustments may be necessary. This adaptive capability is particularly important for high-frequency applications where optimal parameters may change due to shifts in market volatility, trading patterns, or regulatory changes.

The hyperparameter optimization framework includes specialized techniques for handling the temporal dependencies inherent in financial time series data. Cross-validation procedures are designed to respect the temporal structure of the data, avoiding look-ahead bias while providing reliable estimates of out-of-sample performance. The optimization process also incorporates techniques for handling missing data and irregular sampling intervals that are common in high-frequency financial datasets.

4. Results and Discussion

4.1. Multi-Head Attention Performance Analysis

The experimental evaluation of our multi-head attention framework demonstrates significant improvements in relationship discovery accuracy compared to traditional econometric methods and standard deep learning approaches. The evaluation was conducted using comprehensive high-frequency datasets spanning multiple asset classes and market conditions, with particular focus on the models ability to identify meaningful financial relationships while filtering out microstructure noise.

Performance analysis of individual attention heads reveals that our financially-aware multi-head architecture successfully specializes different heads for distinct types of market relationships. Short-term attention heads demonstrate superior performance in capturing microstructure effects and intraday volatility patterns, achieving correlation coefficients of 0.82-0.89 with realized volatility measures. Long-term attention heads excel at identifying persistent relationships and trend patterns, with correlation coefficients of 0.75-0.84 for monthly relationship stability measures.

The multi-head attention mechanism shows particularly strong performance during periods of market stress and regime changes. During the evaluation period, which included several market volatility events, the attention heads demonstrated adaptive behavior by automatically adjusting their focus to risk factors and correlation changes. This adaptive capability resulted in 35% better performance in detecting relationship changes compared to static attention mechanisms.

Cross-asset attention analysis reveals that the framework successfully identifies both direct and indirect relationships between financial instruments. The attention weights provide interpretable insights into how information flows between different assets, with clear patterns emerging around market opening and closing times, earnings announcements, and macroeconomic events. The system identifies lead-lag relationships with average detection accuracy of 78% for relationships with lags between 1-10 minutes.

Computational efficiency analysis shows that our multi-head attention framework maintains practical processing speeds suitable for high-frequency trading applications. Average attention computation time is 12.3 milliseconds per update for a 100-asset universe with 1-minute sampling frequency, well within the latency requirements for most algorithmic trading strategies.

4.2. Temporal Pattern Discovery and Market Insight Analysis

The temporal pattern recognition component of our framework reveals sophisticated relationship structures that provide valuable insights into high-frequency market dynamics. Analysis of discovered patterns demonstrates the systems ability to identify both known market phenomena and previously uncharacterized behavioral patterns across different temporal scales and market conditions.

Intraday pattern analysis identifies systematic variations in relationship strength throughout the trading day. The framework detects stronger correlations during market opening and closing periods, with relationship strength typically 40-60% higher during the first and last hours of trading compared to mid-day periods. These patterns align with theoretical expectations about information processing and liquidity concentration but provide much more detailed timing and magnitude information than previously available through traditional analysis methods.

Microstructure pattern discovery reveals consistent short-term relationship patterns related to order flow and market making activities. The system identifies recurring patterns in bid-ask spread dynamics and their correlation with volume patterns across different assets. These microstructure relationships show strong persistence at time scales of 1-5 minutes but break down at longer horizons, confirming theoretical predictions about the temporal structure of microstructure effects.

The framework successfully identifies regime-dependent relationship patterns that change based on market volatility conditions. During high-volatility periods, the system detects increased correlation between previously independent assets, with correlation increases of 25-40% compared to normal market conditions. These regime-dependent patterns provide important insights for risk management applications and help explain the breakdown of diversification strategies during market stress periods.

Event-driven pattern analysis demonstrates the frameworks ability to detect and characterize the propagation of information across different assets and market segments. The system identifies consistent patterns in how earnings announcements, economic data releases, and other market-moving events influence relationship dynamics. Response patterns typically emerge within 2-5 minutes of event occurrence and can persist for 30-60 minutes depending on event significance.

Long-term pattern stability analysis shows that while many microstructure patterns are highly persistent, fundamental relationship patterns evolve more gradually over periods of weeks to months. The framework successfully tracks these evolutionary changes and provides early warning indicators for significant relationship regime changes that may affect portfolio and risk management strategies.

The interpretability analysis of attention weights reveals that different types of relationships are driven by distinct temporal patterns and market factors. Fundamental economic relationships tend to be captured by longer-term attention patterns that focus on economic indicators and earnings trends, while technical relationships appear in shorter-term attention patterns concentrated around price and volume signals.

5. Conclusion

This research demonstrates the successful application of specialized transformer architectures to the challenging problem of discovering dynamic relationships in high-frequency financial time series. Our comprehensive framework addresses the unique characteristics of high-frequency financial data while providing interpretable insights into market microstructure effects and cross-asset dependencies that are crucial for modern financial applications.

The key innovations of our work include the development of financially-aware multi-head attention mechanisms that can simultaneously capture relationships across multiple temporal scales, temporal pattern recognition algorithms that identify and validate recurring market behaviors, and sophisticated hyperparameter optimization strategies specifically designed for the statistical properties of high-frequency financial data. These innovations demonstrate how transformer architectures can be adapted to address domain-specific challenges while maintaining their core strengths in sequence modeling and relationship discovery.

Experimental results validate the effectiveness of our approach across multiple dimensions of performance. The multi-head attention framework achieves superior accuracy in relationship discovery compared to traditional econometric methods, with correlation detection improvements of 25-40% and lead-lag relationship identification accuracy of 78%. The temporal pattern recognition component successfully identifies both known market phenomena and previously uncharacterized behavioral patterns, providing valuable insights for understanding market microstructure dynamics.

The practical implications of this work extend across multiple areas of financial practice. High-frequency traders can leverage the discovered relationships to improve execution strategies and identify arbitrage opportunities that exist at microsecond to minute time scales. Risk managers gain access to detailed insights into correlation dynamics and volatility spillover effects that are crucial for accurate risk assessment in fast-moving markets. The frameworks ability to detect regime changes and relationship evolution provides early warning capabilities for portfolio rebalancing and hedging strategies.

The interpretability features of our framework represent a significant advancement in financial machine learning applications. The attention weight visualizations and pattern recognition outputs enable financial analysts to understand not only what relationships exist but also why they occur and how they evolve over time. This interpretability is essential for building confidence in automated systems and ensuring compliance with regulatory requirements for model explainability in financial applications.

Future research directions include extending the framework to incorporate additional data types such as order book depth information, news sentiment data, and alternative data sources that are increasingly important in modern quantitative finance. Development of specialized attention mechanisms for specific financial instruments such as options and fixed-income securities represents another promising avenue for research. Creating real-time adaptive learning algorithms that can automatically adjust to new market regimes and structural changes will further enhance the practical applicability of the framework.

The broader implications of this work extend beyond financial applications to other domains involving high-frequency time series data with complex temporal dependencies. The principles and techniques developed for financial relationship discovery can potentially be adapted for applications in energy markets, supply chain optimization, and telecommunications network analysis, where understanding dynamic relationships in high-frequency data is crucial for operational effectiveness.

As financial markets continue to evolve toward higher frequencies and greater algorithmic complexity, the need for sophisticated relationship discovery tools becomes increasingly critical. Our transformer-based framework provides a robust foundation for understanding these complex market dynamics while maintaining the interpretability and real-time performance characteristics necessary for practical deployment in modern financial systems. The combination of advanced machine learning capabilities with domain-specific adaptations demonstrates a path forward for creating trustworthy and effective artificial intelligence systems in the demanding environment of high-frequency financial markets.

References

[1] Addy, W. A., Ajayi-Nifise, A. O., Bello, B. G., Tula, S. T., Odeyemi, O., & Falaiye, T. (2024). Machine learning in financial markets: A critical review of algorithmic trading and risk management. International Journal of Science and Research Archive, 11(1), 1853-1862.

[2] Flas, M. (2025). Predicting Extreme Price Movements in Technology Stocks: A Study of High-Frequency Trading Dynamic.

[3] Dutta, C., Karpman, K., Basu, S., & Ravishanker, N. (2023). Review of statistical approaches for modeling high-frequency trading data. Sankhya B, 85(Suppl 1), 1-48.

[4] Chopra, R. (2025). Essentials of Time Series Econometrics. Educohack Press.

[5] Kamatala, S., Jonnalagadda, A. K., & Naayini, P. (2025). Transformers beyond nlp: Expanding horizons in machine learning. Yes it was accepted by IRE Journals, 8(7).

[6] Hasan, F., & Huang, H. (2023). MALS-Net: A multi-head attention-based LSTM sequence-to-sequence network for socio-temporal interaction modelling and trajectory prediction. Sensors, 23(1), 530.

[7] Nie, Y., Kong, Y., Dong, X., Mulvey, J. M., Poor, H. V., Wen, Q., & Zohren, S. (2024). A survey of large language models for financial applications: Progress, prospects and challenges. arXiv preprint arXiv:2406.11903.

[8] Zheng, W., & Liu, W. (2025). Symmetry-Aware Transformers for Asymmetric Causal Discovery in Financial Time Series. Symmetry.

[9] Balaji, K. (2025). Revolutionizing High-Frequency Trading: The Impacts of Financial Technology and Data Science Innovations. Machine Learning and Modeling Techniques in Financial Data Science, 103-124.

[10] Rasyid, A. (2024). The Art and Science of Financial Econometrics: Applications, Challenges, and Future Directions. Advances in Economics & Financial Studies, 2(1), 44-52.

[11] Muhammad, A., Aliyu, J. N., Adetunji, A. L., Adesugba, A. K., Mike, M. E. E., & Abdulmalik, M. (2023). Fractal geometry in high-frequency trading: Modeling market microstructure and price dynamics. Saudi J Econ Fin, 7(11), 484-488.

[12] Chang, P., Pienaar, E., & Gebbie, T. (2021). The Epps effect under alternative sampling schemes. Physica A: Statistical Mechanics and its Applications, 583, 126329.

[13] Pattanayak, A. M., Swetapadma, A., & Sahoo, B. (2024). Exploring different dynamics of recurrent neural network methods for stock market prediction-a comparative study. Applied Artificial Intelligence, 38(1), 2371706.

[14] Junior, E. H. (2025). High-frequency data and machine learning. In Impacts of Innovation and Cognition in Management (pp. 183-204). IGI Global.

[15] Mehra, A., & Singh, S. P. (2024). Event-driven architectures for real-time error resolution in high-frequency trading systems. International Journal of Research in Modern Engineering and Emerging Technology, 12(12), 671.

[16] Chogii, R. (2021). Market Information Risk, Trading Activity, Organizational Characteristics and Price Discovery for Stocks Listed at the Nairobi Securities Exchange (Doctoral dissertation, University of Nairobi).

[17] Ji, E., Wang, Y., Xing, S., & Jin, J. (2025). Hierarchical Reinforcement Learning for Energy-Efficient API Traffic Optimization in Large-Scale Advertising Systems. IEEE Access.

[18] Jin, J., Xing, S., Ji, E., & Liu, W. (2025). XGate: Explainable Reinforcement Learning for Transparent and Trustworthy API Traffic Management in IoT Sensor Networks. Sensors (Basel, Switzerland), 25(7), 2183.

[19] Wang, J., Liu, J., Zheng, W., & Ge, Y. (2025). Temporal Heterogeneous Graph Contrastive Learning for Fraud Detection in Credit Card Transactions. IEEE Access.

[20] Mai, N. T., Cao, W., & Liu, W. (2025). Interpretable Knowledge Tracing via Transformer-Bayesian Hybrid Networks: Learning Temporal Dependencies and Causal Structures in Educational Data. Applied Sciences, 15(17), 9605.

[21] Cao, W., Mai, N. T., & Liu, W. (2025). Adaptive knowledge assessment via symmetric hierarchical Bayesian neural networks with graph symmetry-aware concept dependencies. Symmetry, 17(8), 1332.

[22] Mai, N. T., Cao, W., & Wang, Y. (2025). The global belonging support framework: Enhancing equity and access for international graduate students. Journal of International Students, 15(9), 141-160.

[23] Tan, Y., Wu, B., Cao, J., & Jiang, B. (2025). LLaMA-UTP: Knowledge-Guided Expert Mixture for Analyzing Uncertain Tax Positions. IEEE Access.

[24] Cao, J., Zheng, W., Ge, Y., & Wang, J. (2025). DriftShield: Autonomous fraud detection via actor-critic reinforcement learning with dynamic feature reweighting. IEEE Open Journal of the Computer Society.

[25] Chen, S., Liu, Y., Zhang, Q., Shao, Z., & Wang, Z. (2025). Multi-Distance Spatial-Temporal Graph Neural Network for Anomaly Detection in Blockchain Transactions. Advanced Intelligent Systems, 2400898.

[26] Zhang, X., Chen, S., Shao, Z., Niu, Y., & Fan, L. (2024). Enhanced Lithographic Hotspot Detection via Multi-Task Deep Learning with Synthetic Pattern Generation. IEEE Open Journal of the Computer Society.

[27] Zhang, Q., Chen, S., & Liu, W. (2025). Balanced Knowledge Transfer in MTTL-ClinicalBERT: A Symmetrical Multi-Task Learning Framework for Clinical Text Classification. Symmetry, 17(6), 823.

[28] Shao, Z., Wang, X., Ji, E., Chen, S., & Wang, J. (2025). GNN-EADD: Graph Neural Network-based E-commerce Anomaly Detection via Dual-stage Learning. IEEE Access.

[29] Li, P., Ren, S., Zhang, Q., Wang, X., & Liu, Y. (2024). Think4SCND: Reinforcement Learning with Thinking Model for Dynamic Supply Chain Network Design. IEEE Access.

[30] Liu, Y., Ren, S., Wang, X., & Zhou, M. (2024). Temporal logical attention network for log-based anomaly detection in distributed systems. Sensors, 24(24), 7949.

[31] Ren, S., Jin, J., Niu, G., & Liu, Y. (2025). ARCS: Adaptive Reinforcement Learning Framework for Automated Cybersecurity Incident Response Strategy Optimization. Applied Sciences, 15(2), 951.