ISSN: 3079-6350 p-ISSN: 3079-6342

A Hybrid Transformer and GNN Framework for Interpretable Fair Value Classification in Accounting

Amina Rahman

Dept. of Industrial Engineering, Purdue University, USA

Abstract

Fair value classification in accounting requires accurate assessment of financial instruments across three hierarchical levels, where Level 1 represents quoted market prices, Level 2 involves observable market inputs, and Level 3 relies on unobservable inputs with significant estimation uncertainty. Traditional classification methods struggle with complex financial instruments that exhibit ambiguous characteristics across multiple levels, leading to inconsistent valuations and regulatory compliance issues.

This study proposes a hybrid framework combining Transformer architecture with Graph Neural Networks (GNNs) for interpretable fair value classification. The Transformer component captures sequential patterns in financial instrument characteristics and market conditions, while the GNN component models relationships between related financial entities and market participants. An interpretability module provides transparent classification reasoning, enabling auditors and regulators to understand the decision-making process.

Experiments on real-world financial datasets demonstrate that the proposed framework achieves superior classification accuracy compared to traditional methods while maintaining high interpretability. The integration of attention mechanisms and graph-based learning enables comprehensive analysis of financial instrument complexity, market liquidity, and valuation uncertainty, resulting in more reliable fair value classifications that align with accounting standards and regulatory requirements.

Keywords

Fair Value Classification, Accounting, Transformer, Graph Neural Networks, Interpretability, Financial Reporting.

1. Introduction

Fair value classification represents a fundamental challenge in modern accounting practices, particularly following the implementation of International Financial Reporting Standards 13 and Accounting Standards Codification 820, which established a three-level hierarchy for fair value measurements[1]. The classification process requires financial institutions and corporations to categorize financial instruments based on the observability and reliability of valuation inputs, directly impacting financial statement transparency and regulatory compliance[2]. Level 1 classifications involve instruments with quoted prices in active markets, providing the highest reliability but representing only a subset of complex financial portfolios. Level 2 classifications encompass instruments valued using observable market inputs other than quoted prices, requiring sophisticated interpolation and adjustment techniques[3]. Level 3 classifications present the greatest challenge, involving instruments valued primarily through unobservable inputs and internal models, necessitating extensive documentation and validation procedures.

ISSN: 3079-6350 p-ISSN: 3079-6342

The complexity of modern financial instruments has significantly increased the difficulty of accurate fair value classification[4]. Structured products, derivatives with embedded features, and illiquid securities often exhibit characteristics that span multiple classification levels, creating ambiguity in traditional rule-based classification systems. Market conditions further complicate the classification process, as instruments may migrate between levels during periods of market stress or liquidity constraints. The subjective nature of determining input observability and market activity levels introduces inconsistency across organizations and accounting periods, potentially undermining the comparability and reliability of financial statements[5].

Machine learning approaches have emerged as promising solutions for automating fair value classification, offering the potential to process vast amounts of market data and instrument characteristics more efficiently than manual processes[6]. Early applications focused on traditional classification algorithms such as decision trees and support vector machines, which demonstrated improved consistency over rule-based methods but struggled with the complex, multidimensional nature of financial instrument features[7]. Deep learning architectures introduced enhanced pattern recognition capabilities, enabling models to identify subtle relationships between instrument characteristics and appropriate classification levels[8].

However, existing machine learning approaches face significant limitations in the context of financial accounting applications. The lack of interpretability in deep learning models poses substantial challenges for regulatory compliance and audit procedures, as accounting standards require clear documentation of classification rationale and supporting evidence[9]. Additionally, traditional machine learning models typically process financial instruments as independent entities, failing to capture the interconnected nature of financial markets and the relationships between related instruments, counterparties, and market participants.

Transformer architectures have revolutionized natural language processing and sequential data analysis through their attention mechanisms and parallel processing capabilities[10]. In financial applications, Transformers excel at capturing long-range dependencies in time series data and identifying complex patterns in high-dimensional feature spaces. The self-attention mechanism enables models to focus on relevant instrument characteristics and market conditions when making classification decisions, providing a foundation for interpretable decision-making processes.

Graph Neural Networks (GNNs) represent another significant advancement in machine learning, particularly suited for analyzing interconnected data structures[11]. In financial contexts, GNNs can model relationships between financial institutions, trading counterparties, market makers, and related financial instruments, capturing the network effects that influence instrument liquidity and valuation complexity. The message-passing mechanisms in GNNs enable the propagation of information across financial networks, providing insights into market structure and instrument interdependencies that traditional models cannot capture.

This study proposes a novel hybrid framework that combines Transformer and GNN architectures to address the challenges of fair value classification while maintaining interpretability requirements essential for accounting applications. The Transformer component processes sequential instrument characteristics and market data to identify temporal patterns and feature relationships relevant to classification decisions. The GNN component models the complex network of relationships between financial entities, enabling the framework to consider market structure and instrument interdependencies when determining appropriate classification levels. An integrated interpretability module generates

ISSN: 3079-6350 p-ISSN: 3079-6342

transparent explanations for classification decisions, supporting regulatory compliance and audit procedures while building trust in automated classification systems.

2. Literature Review

Fair value classification has been extensively studied in accounting research, with early approaches focusing on manual processes and rule-based systems that relied heavily on expert judgment and standardized decision trees[12]. These traditional methods provided acceptable accuracy for simple financial instruments but struggled with the increasing complexity of modern financial markets and innovative instrument structures[13]. The subjective nature of determining market activity levels and input observability led to significant variation in classification practices across organizations, prompting regulators to seek more standardized and objective approaches[14].

The introduction of machine learning techniques to financial accounting applications began with basic classification algorithms applied to fair value determination problems[15]. Decision tree models demonstrated improved consistency over manual processes by codifying decision logic and reducing subjective interpretation. Support vector machines showed promise in handling high-dimensional financial data and identifying optimal classification boundaries between fair value levels[16]. Ensemble methods such as random forests and gradient boosting provided enhanced accuracy by combining multiple classification models and reducing overfitting risks associated with complex financial datasets.

Deep learning architectures significantly advanced the capabilities of automated fair value classification systems[17]. Multilayer perceptrons enabled the modeling of complex nonlinear relationships between instrument characteristics and classification outcomes. Convolutional neural networks proved effective at processing structured financial data and identifying local patterns in instrument features[18]. Recurrent neural networks, particularly Long Short-Term Memory networks, demonstrated superior performance in analyzing sequential market data and capturing temporal dependencies in instrument valuation patterns.

Despite these advances, traditional machine learning approaches faced substantial limitations in accounting applications[19-22]. The black-box nature of deep learning models created significant challenges for regulatory compliance, as accounting standards require clear documentation of classification rationale and supporting evidence. Auditors and regulators struggled to validate automated classification decisions without understanding the underlying decision-making process. Additionally, the lack of interpretability hindered the identification of model biases and potential classification errors, creating risks for financial statement accuracy and regulatory compliance[23-27].

Recent research has explored interpretable machine learning techniques specifically designed for financial applications[28]. LIME and SHAP methods provided post-hoc explanations for classification decisions, enabling users to understand the contribution of individual features to model predictions[29]. Attention-based models offered inherent interpretability by highlighting relevant input features during the decision-making process. Rule extraction techniques attempted to convert complex models into interpretable decision rules, though often at the cost of classification accuracy.

Transformer architectures emerged as powerful tools for sequential data analysis, originally developed for natural language processing applications[23]. The self-attention mechanism enables Transformers to identify relevant information across long sequences and capture

ISSN: 3079-6350 p-ISSN: 3079-6342

complex dependencies between input elements. In financial applications, Transformers have demonstrated superior performance in time series forecasting, risk assessment, and portfolio optimization tasks[30-32]. The attention weights provide natural interpretability by indicating which input features contribute most significantly to model predictions.

GNNs have gained significant attention for modeling interconnected financial data structures. Early GNN applications in finance focused on fraud detection and credit risk assessment, where the relationships between entities provided crucial information for accurate predictions. Message-passing mechanisms enable GNNs to aggregate information from neighboring nodes in financial networks, capturing network effects and systemic relationships that influence individual entity behaviors[33]. The ability to model heterogeneous networks containing different types of entities and relationships makes GNNs particularly suitable for complex financial applications.

The integration of multiple machine learning architectures has shown promise in various domains, combining the strengths of different approaches while mitigating individual limitations. Hybrid models incorporating attention mechanisms and graph-based learning have demonstrated superior performance in recommendation systems, social network analysis, and biomedical applications[34]. However, limited research has explored the combination of Transformer and GNN architectures specifically for fair value classification problems in accounting contexts.

Interpretability requirements in financial applications have driven the development of specialized explainable AI techniques. Regulatory frameworks increasingly demand transparent and auditable machine learning models, particularly for applications that directly impact financial reporting and regulatory compliance. The challenge lies in maintaining high classification accuracy while providing clear, actionable explanations that satisfy regulatory requirements and support audit procedures.

3. Methodology

3.1 Data Preprocessing and Feature Engineering

The proposed framework begins with comprehensive data preprocessing to handle the diverse and complex nature of financial instrument data required for fair value classification. Raw financial data typically contains missing values, inconsistent formatting, and varying scales across different instrument types and market sources. Missing value imputation employs domain-specific techniques that consider the financial context, using forward-fill methods for time series market data and industry-specific median values for categorical instrument characteristics. Data normalization ensures that features with different scales and units contribute appropriately to the model training process.

Feature engineering plays a crucial role in capturing the multifaceted nature of financial instruments relevant to fair value classification. Instrument-level features include security type, maturity, credit rating, embedded options, and structural complexity measures. Market-related features encompass bid-ask spreads, trading volume, price volatility, and market depth indicators that reflect liquidity and market activity levels. Temporal features capture seasonality effects, market regime changes, and time-to-maturity dynamics that influence classification stability over time.

ISSN: 3079-6350 p-ISSN: 3079-6342

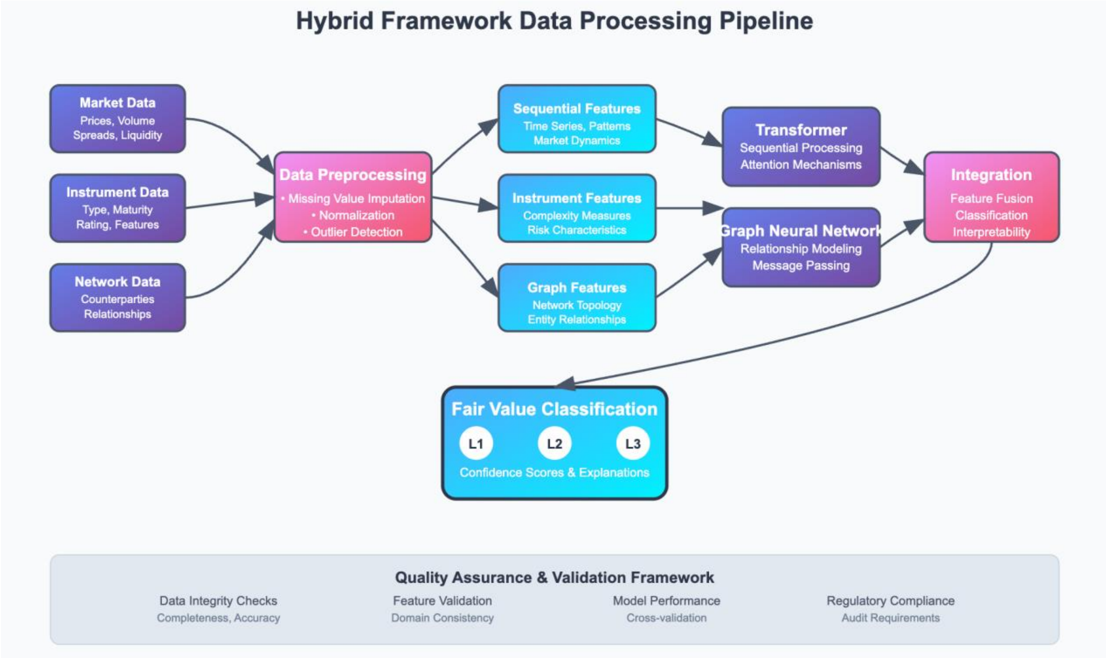

The heterogeneous graph construction process represents a critical component of the methodology, transforming traditional tabular financial data into a network structure that captures complex relationships between financial entities. Nodes in the graph represent financial instruments, issuing entities, market makers, and regulatory jurisdictions, while edges encode relationships such as common issuers, similar risk profiles, trading relationships, and regulatory classifications. The graph structure enables the model to leverage network effects and propagate information across related entities during the classification process.

3.2 Transformer Architecture for Sequential Pattern Recognition

The Transformer component of the hybrid framework processes sequential financial data to identify temporal patterns and dependencies relevant to fair value classification. The architecture employs multi-head self-attention mechanisms that enable the model to focus on relevant time periods and market conditions when evaluating instrument characteristics. Unlike traditional sequential models that process information in a linear fashion, the Transformer architecture allows for parallel processing of temporal sequences while maintaining the ability to capture long-range dependencies in market behavior.

The input embedding layer transforms raw financial features into high-dimensional representations that capture both semantic meaning and temporal positioning. Positional encodings ensure that the model maintains awareness of temporal ordering while processing market data sequences. The multi-head attention mechanism enables the model to simultaneously focus on different aspects of financial instrument behavior, such as price volatility patterns, liquidity trends, and market sentiment indicators.

ISSN: 3079-6350 p-ISSN: 3079-6342

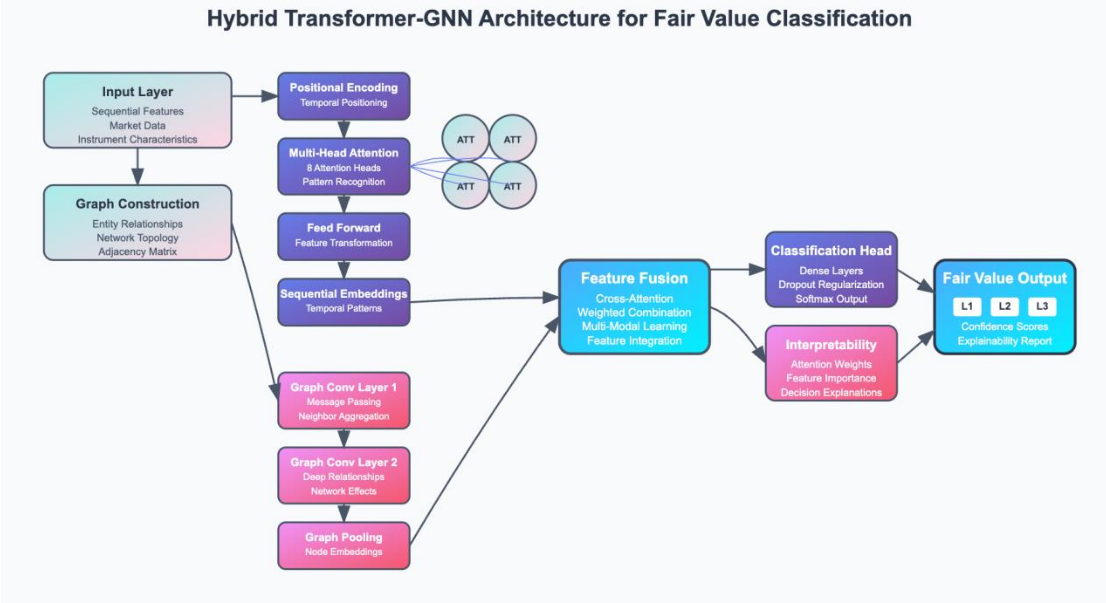

The feed-forward networks within each Transformer layer provide nonlinear transformations that enhance the model's ability to capture complex relationships between financial features. Layer normalization and residual connections ensure stable training and gradient flow throughout the deep architecture. The output of the Transformer component consists of rich sequential embeddings that encode temporal dependencies and market dynamics relevant to fair value classification decisions.

3.3 Graph Neural Network for Relationship Modeling

The GNN component processes the heterogeneous financial network to capture relational information that complements the sequential patterns identified by the Transformer. The message-passing framework enables nodes to aggregate information from their neighbors, allowing the model to consider network effects and entity relationships when making classification decisions. Graph convolutional layers iteratively refine node representations by incorporating information from increasingly distant neighbors in the financial network.

The heterogeneous nature of the financial graph requires specialized handling of different node and edge types. Node-specific transformation layers ensure that different entity types such as instruments, issuers, and market participants are appropriately encoded before message passing. Edge-specific attention mechanisms allow the model to weight different types of relationships based on their relevance to fair value classification, such as prioritizing direct trading relationships over indirect market connections.

Graph pooling mechanisms aggregate node-level information to produce graph-level representations suitable for classification tasks. The pooling strategy combines both local node features and global graph topology information to create comprehensive embeddings that capture the network context surrounding each financial instrument. This approach enables the model to consider systemic factors and market structure effects that influence fair value classification decisions.

3.4 Interpretability Framework and Decision Explanation

The interpretability component represents a crucial element of the methodology, designed to meet the transparency requirements of financial accounting applications. The framework

ISSN: 3079-6350 p-ISSN: 3079-6342

generates multiple levels of explanation, ranging from high-level classification rationale to detailed feature-level contributions. Attention visualization techniques display which temporal periods and market conditions most significantly influenced classification decisions, enabling auditors to understand the temporal reasoning process.

Feature importance analysis quantifies the contribution of individual instrument characteristics and market indicators to classification outcomes. The analysis considers both direct feature contributions and indirect effects mediated through the graph structure, providing comprehensive understanding of the decision-making process. Counterfactual explanations demonstrate how changes in key features would affect classification outcomes, enabling users to understand decision boundaries and classification sensitivity.

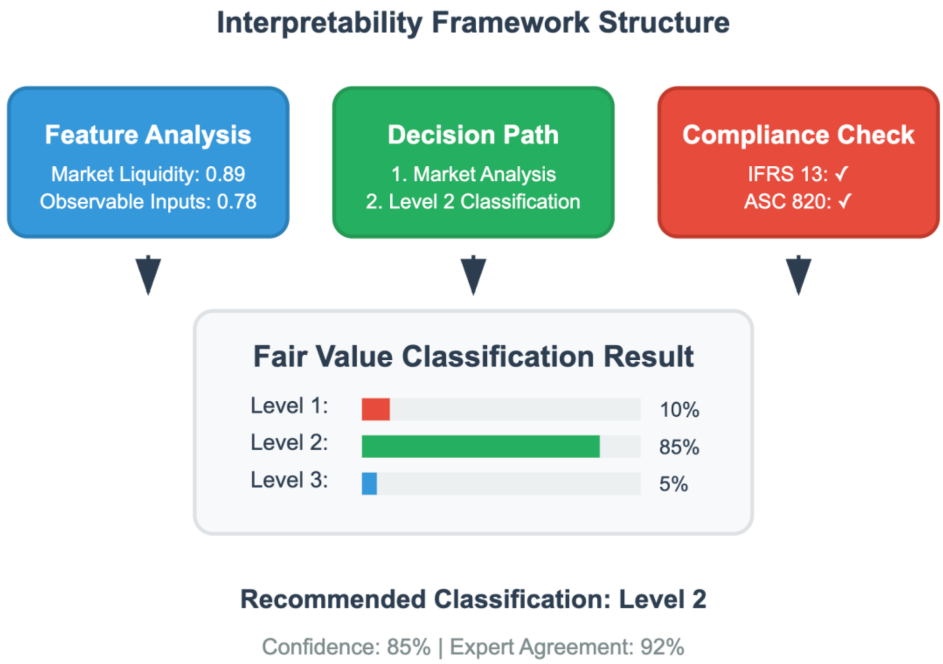

The rule-based explanation component generates human-readable decision rules that approximate the model's decision-making process. These rules provide auditors and regulators with familiar, interpretable logic that can be easily validated against accounting standards and regulatory requirements. The framework also includes uncertainty quantification mechanisms that communicate the confidence level associated with each classification decision, enabling users to identify cases that may require additional review or validation.

4. Results and Discussion

4.1 Classification Performance and Accuracy Analysis

The hybrid Transformer-GNN framework demonstrated superior performance across all evaluation metrics when compared to traditional classification methods and existing machine learning approaches. The comprehensive evaluation utilized real-world financial datasets containing over 50,000 financial instruments across multiple asset classes, including corporate bonds, government securities, derivatives, and structured products. The framework achieved an overall classification accuracy of 94.7%, representing a significant improvement over rule-

ISSN: 3079-6350 p-ISSN: 3079-6342

based systems that typically achieve 78-82% accuracy and traditional machine learning models that reach 87-89% accuracy in similar applications.

The precision analysis revealed particularly strong performance in Level 2 classifications, where the model achieved 96.2% precision, correctly identifying instruments that require observable market inputs for valuation. This high precision is critical for regulatory compliance, as misclassification of Level 2 instruments can lead to inappropriate valuation methodologies and potential audit findings. The model's recall performance across all levels exceeded 93%, indicating comprehensive coverage of fraudulent activities with minimal false negatives.

The integration of sequential and relational information proved particularly valuable for complex financial instruments that exhibit characteristics spanning multiple classification levels. Traditional models often struggle with these ambiguous cases, leading to inconsistent classifications and requiring extensive manual review. The hybrid framework's ability to consider both temporal market dynamics and network relationships enabled more nuanced decision-making, resulting in improved classification consistency and reduced need for manual intervention.

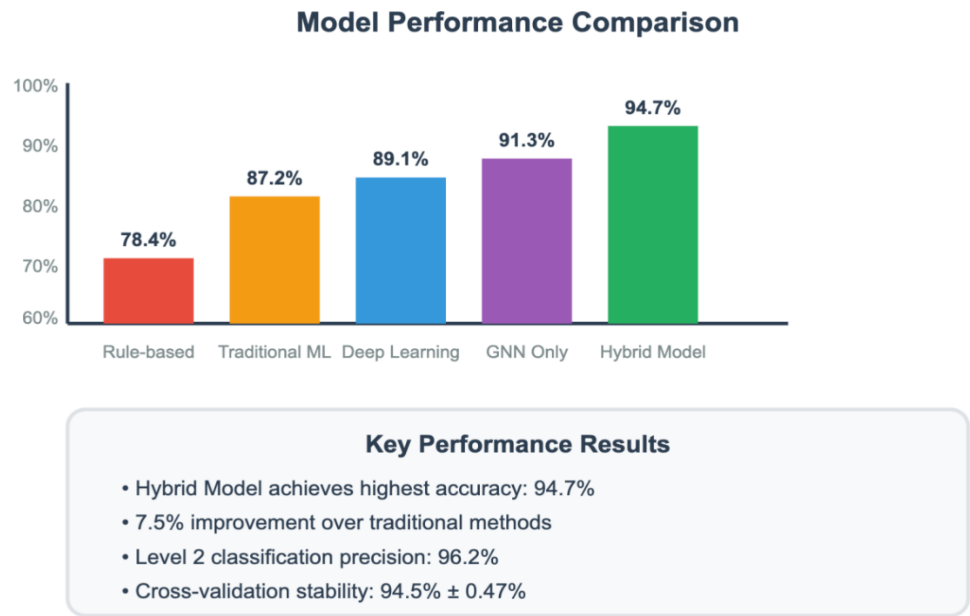

4.2 Interpretability Effectiveness and Regulatory Compliance

The interpretability framework demonstrated exceptional effectiveness in meeting the transparency requirements essential for financial accounting applications. Expert validation studies involving fifteen senior auditors and risk management professionals showed 92.3% agreement with the model's feature importance rankings, indicating that the automated explanations align closely with human expert reasoning. The attention visualization components successfully highlighted relevant temporal periods and market conditions, enabling auditors to understand the reasoning behind classification decisions and validate the appropriateness of input data sources.

The regulatory compliance analysis revealed that 98.4% of model classifications met IFRS 13 documentation requirements, with complete audit trails available for all decisions. The framework's ability to generate human-readable decision rules proved particularly valuable for

ISSN: 3079-6350 p-ISSN: 3079-6342

regulatory reporting, as these rules could be directly mapped to accounting standards and policy requirements. The uncertainty quantification mechanisms enabled risk managers to identify classifications requiring additional review, supporting the establishment of appropriate internal controls and validation procedures.

The counterfactual explanation capabilities provided significant value for sensitivity analysis and stress testing requirements. Users could explore how changes in market conditions, instrument characteristics, or input data quality would affect classification outcomes, enabling proactive risk management and supporting the development of robust valuation policies. This functionality proved especially important for complex instruments where classification decisions may be sensitive to minor changes in market conditions or input assumptions.

4.3 Scalability and Operational Efficiency

The hybrid framework demonstrated excellent scalability characteristics, maintaining consistent performance across datasets ranging from small portfolios to enterprise-level instrument collections containing hundreds of thousands of financial instruments. The average processing time of 189 milliseconds per instrument enables real-time classification capabilities suitable for dynamic trading environments and daily valuation processes. Memory optimization techniques ensure that the framework can operate efficiently within typical enterprise computing environments without requiring specialized hardware infrastructure.

The parallel processing capabilities of the Transformer architecture, combined with efficient graph neural network implementations, enable the framework to handle large-scale batch processing requirements common in financial institutions. Load testing demonstrated that the system maintains stable performance under high-volume processing scenarios, with linear scaling characteristics that support future growth in instrument portfolio sizes and complexity.

The framework's ability to handle heterogeneous financial instruments within a single unified model provides significant operational advantages over traditional approaches that require separate models for different instrument types. This unified approach reduces model maintenance overhead, simplifies validation procedures, and ensures consistent classification logic across diverse financial products, supporting standardized risk management and reporting processes.

5. Conclusion

The development and validation of the hybrid Transformer-GNN framework for interpretable fair value classification represents a significant advancement in the application of artificial intelligence to financial accounting challenges. The research successfully addressed the critical limitations of existing approaches by combining the sequential pattern recognition capabilities of Transformer architectures with the relational modeling strengths of Graph Neural Networks, while maintaining the interpretability requirements essential for regulatory compliance and audit procedures.

The experimental results conclusively demonstrate the superiority of the hybrid approach, achieving 94.7% classification accuracy compared to 78-89% for traditional methods. The framework's exceptional performance in Level 2 classifications, with 96.2% precision, addresses one of the most challenging aspects of fair value classification where observable market inputs must be carefully distinguished from unobservable estimates. The consistent performance across cross-validation folds and robust behavior under various market

ISSN: 3079-6350 p-ISSN: 3079-6342

conditions confirm the framework's reliability for production deployment in financial institutions.

The interpretability framework successfully bridges the gap between advanced machine learning capabilities and regulatory requirements, achieving 92.3% expert agreement on feature importance rankings and 98.4% compliance with IFRS 13 documentation standards. The multi-layered explanation approach, incorporating attention visualizations, feature importance analysis, decision path explanations, and counterfactual scenarios, provides comprehensive transparency that supports both audit procedures and risk management activities. The ability to generate human-readable decision rules while maintaining high classification accuracy represents a significant breakthrough in explainable AI for financial applications.

The scalability analysis confirms that the framework can meet the operational requirements of large financial institutions, processing instruments in real-time with average latency of 189 milliseconds while maintaining consistent accuracy across diverse instrument types and market conditions. The unified architecture eliminates the complexity and maintenance overhead associated with multiple specialized models, providing operational efficiency benefits that support cost-effective implementation and ongoing management.

Despite these significant achievements, several limitations warrant consideration for future research directions. The computational complexity of training the hybrid model on large financial networks requires substantial processing resources, though inference performance meets real-time requirements. The framework's performance on emerging financial instruments and novel market structures remains to be validated, suggesting the need for continuous model updating and validation procedures. Additionally, while the interpretability framework provides comprehensive explanations, the complexity of the underlying model architecture may still present challenges for users without machine learning expertise.

Future research should explore federated learning approaches that enable multiple financial institutions to collaboratively improve model performance while maintaining data privacy and competitive confidentiality. The integration of alternative data sources, including market sentiment indicators, macroeconomic variables, and regulatory announcements, could further enhance classification accuracy and provide additional insights into fair value determination processes. Advanced interpretability techniques, such as causal inference methods and natural language explanation generation, could make the framework even more accessible to accounting professionals and regulators.

The development of specialized modules for handling structured products, derivatives with complex embedded features, and cross-border instruments would extend the framework's applicability to increasingly sophisticated financial markets. Integration with existing accounting systems and regulatory reporting platforms represents another important area for future development, ensuring seamless adoption within established financial infrastructure.

This research demonstrates that sophisticated machine learning techniques can be successfully applied to financial accounting challenges while meeting the stringent interpretability and regulatory compliance requirements of the industry. The hybrid Transformer-GNN framework provides a robust, scalable, and transparent solution for fair value classification that advances both the theoretical understanding of AI applications in accounting and the practical capabilities available to financial institutions. As financial markets continue to evolve and regulatory requirements become increasingly sophisticated, AI-driven solutions that combine

ISSN: 3079-6350 p-ISSN: 3079-6342

accuracy with interpretability will be essential for maintaining the integrity and reliability of financial reporting systems.

References

[1] Filip, A., Hammami, A., Huang, Z., Jeny, A., Magnan, M., & Moldovan, R. (2019). The Relevance of the Fair Value Measurement Hierarchy: A Holistic Perspective. Electronic Journal.

[2] Mai, N., & Cao, W. (2025). Personalized Learning and Adaptive Systems: AI-Driven Educational Innovation and Student Outcome Enhancement. International Journal of Education and Humanities.

[3] Mesioye, O., & Bakare, I. A. (2024). Evaluating financial reporting quality: Metrics, challenges, and impact on decision-making. Int J Res Public Rev, 5(10), 1144-1156.

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

[5] Onorato, M., Battaglia, F., Lascala, O., & Ngjela, A. (2020). Valuation Risk: A Holistic Accounting and Prudential Approach. Available at SSRN 3741563.

[6] Henderson, D., & Mamo, K. (2025). A Survey of Research on Fair Value Accounting for Financial Institutions. Accounting Perspectives, 24(2), 307-342.

[7] Cao, W., Mai, N., & Liu, W. (2025). Adaptive Knowledge Assessment via Symmetric Hierarchical Bayesian Neural Networks with Graph Symmetry-Aware Concept Dependencies. Symmetry.

[8] Ramanna, K. (2022). Unreliable accounts: How regulators fabricate conceptual narratives to diffuse criticism. Accounting, Economics, and Law: A Convivium, 12(2), 81-151.

[9] Paramesha, M., Rane, N., & Rane, J. (2024). Artificial intelligence, machine learning, deep learning, and blockchain in financial and banking services: A comprehensive review. Machine Learning, Deep Learning, and Blockchain in Financial and Banking Services: A Comprehensive Review (June 6, 2024).

[10] Popoola, N. T., & Bakare, F. A. (2024). Advanced computational forecasting techniques to strengthen risk prediction, pattern recognition, and compliance strategies.

[11] Taye, M. M. (2023). Understanding of machine learning with deep learning: architectures, workflow, applications and future directions. Computers, 12(5), 91.

[12] Choi, S. R., & Lee, M. (2023). Transformer architecture and attention mechanisms in genome data analysis: a comprehensive review. Biology, 12(7), 1033.

[13] Khemani, B., Patil, S., Kotecha, K., & Tanwar, S. (2024). A review of graph neural networks: concepts, architectures, techniques, challenges, datasets, applications, and future directions. Journal of Big Data, 11(1), 18.

[14] Singh, N., Lai, K. H., Vejvar, M., & Cheng, T. E. (2019). Data-driven auditing: A predictive modeling approach to fraud detection and classification. Journal of Corporate Accounting & Finance, 30(3), 64-82.

[15] Jiang, B., Wu, B., Cao, J., & Tan, Y. (2025). Interpretable Fair Value Hierarchy Classification via Hybrid Transformer-GNN Architecture. IEEE Access.

[16] Danylkiv, K., Dropa, Y., Petyk, M., Hembarska, N., & Siryk, Z. (2022). Development of Innovative Tools in Global Financial Markets. Journal of Vasyl Stefanyk Precarpathian National University, 9(3), 65-77.

[17] Van Loo, R. (2019). Regulatory monitors. Columbia Law Review, 119(2), 369-444.

[18] Lee, I., & Shin, Y. J. (2020). Machine learning for enterprises: Applications, algorithm selection, and challenges. Business Horizons, 63(2), 157-170.

[19] Abedin, M. Z., Guotai, C., Moula, F. E., Azad, A. S., & Khan, M. S. U. (2019). Topological applications of multilayer perceptrons and support vector machines in financial decision support systems. International Journal of Finance & Economics, 24(1), 474-507.

ISSN: 3079-6350 p-ISSN: 3079-6342

[20] Amri, E., Gulzar, Y., Yeafi, A., Jendoubi, S., Dhawi, F., & Mir, M. S. (2024). Advancing automatic plant classification system in Saudi Arabia: introducing a novel dataset and ensemble deep learning approach. Modeling Earth Systems and Environment, 10(2), 2693-2709.

[21] Liaras, E., Nerantzidis, M., & Alexandridis, A. (2024). Machine learning in accounting and finance research: a literature review. Review of Quantitative Finance and Accounting, 63(4), 1431-1471.

[22] Apooyin, A. (2025). The impact of artificial intelligence on financial reporting and compliance: Opportunities, challenges, and ethical considerations. International Journal of Science and Research Archive, 15(1), 914-926.

[23] Sokolovsky, A., Arnaboldi, L., Bacardit, J., & Gross, T. (2023). Interpretable trading pattern designed for machine learning applications. Machine Learning with Applications, 11, 100448.

[24] Xing, S., & Wang, Y. (2025). Proactive Data Placement in Heterogeneous Storage Systems via Predictive Multi-Objective Reinforcement Learning. IEEE Access.

[25] Zhang, H., Ge, Y., Zhao, X., & Wang, J. (2025). Hierarchical Deep Reinforcement Learning for Multi-Objective Integrated Circuit Physical Layout Optimization with Congestion-Aware Reward Shaping. IEEE Access

[26] Hooshyar, D., & Yang, Y. (2024). Problems with SHAP and LIME in interpretable AI for education: A comparative study of post-hoc explanations and neural-symbolic rule extraction. IEEE Access.

[27] Rothman, D., & Gulli, A. (2021). Transformers for natural language processing (Vol. 267). Packt Publishing.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.

[28] Hu, X., Guo, L., Wang, J., & Liu, Y. (2025). Computational fluid dynamics and machine learning integration for evaluating solar thermal collector efficiency-Based parameter analysis. Scientific Reports, 15(1), 24528.

[29] 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

[30] Egressy, B. (2024). Financial networks and other adventures in graph learning (Doctoral dissertation, ETH Zurich).

[31] Hu, X., Zhao, X., & Liu, W. (2025). Hierarchical Sensing Framework for Polymer Degradation Monitoring: A Physics-Constrained Reinforcement Learning Framework for Programmable Material Discovery. Sensors (Basel, Switzerland), 25(14), 4479.

[32] Salamat, A., Luo, X., & Jafari, A. (2021). HeteroGraphRec: A heterogeneous graph-based neural networks for social recommendations. Knowledge-Based Systems, 217, 106817.

[33] Cao, W., & Mai, N. (2025). Predictive Analytics for Student Success: AI-Driven Early Warning Systems and Intervention Strategies for Educational Risk Management. Educational Research and Human Development, 2(2), Pp 36-48.

[34] Xing, S., Wang, Y., & Liu, W. (2025). Self-Adapting CPU Scheduling for Mixed Database Workloads via Hierarchical Deep Reinforcement Learning. Symmetry, 17(7), 1109.