Simen Owen Academic Proceedings Series

Vol. 3 2026

Article Open Access

Graph Neural Networks for Business Relationship Mining: Applications and Performance Analysis

Meiwen Qiu [1,*]

[1] School of Physical and Mathematical Sciences (SPMS), Nanyang Technological University, 637121, Singapore
* Correspondence: Meiwen Qiu, School of Physical and Mathematical Sciences (SPMS), Nanyang Technological University, 637121, Singapore

Abstract: Business ecosystems, such as supply-chain networks, financial transaction systems, and e-commerce platforms, exhibit complex relational structures that challenge traditional machine-learning models. Although graph neural networks (GNNs) have shown promise in capturing such dependencies, existing studies often focus on single domains, rely on static graphs, or lack systematic comparison across heterogeneous commercial settings. To address these gaps, this study proposes a unified analytical framework that integrates relational embeddedness theory, graph representation learning, and dynamic capability perspectives. Using three representative real-world scenarios, a retail procurement graph, an AML transaction network, and an e-commerce product affinity graph, we evaluate four GNN architectures (GCN, GraphSAGE, GAT, and Temporal-GNN) through link prediction, fraud detection, and recommendation tasks. The results show that attention-based models outperform others in heterogeneous supplier and transaction environments, temporal GNNs better capture evolving fraud patterns, and inductive architectures excel in high-turnover product graphs. These findings deepen theoretical understanding of relational learning in commercial systems and offer practical guidance for deploying GNN-based analytics in procurement risk assessment, financial compliance, and personalized recommendation services.

Keywords: graph neural networks; business relationship mining; supply-chain analytics; fraud detection; product affinity modeling

Received: 25 December 2025
Revised: 03 February 2026
Accepted: 15 February 2026
Published: 18 February 2026

Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

Modern business ecosystems, such as supply-chain networks, merchant-client transaction systems, and large-scale e-commerce platforms, exhibit increasingly complex relational structures [1]. Firms, suppliers, customers, and products are connected through multi-layered interactions that form high-dimensional graphs [2]. These relational dependencies determine essential business functions including risk monitoring, supplier diversification, inventory forecasting, and real-time fraud detection [3]. For example, the procurement network of a large retailer such as Walmart contains thousands of suppliers and millions of purchase links, where disruptions in a few nodes may propagate across the entire network. Similarly, the IEEE-CIS anti-money-laundering (AML) transaction dataset reveals dense clusters of merchant-client interactions, in which fraudulent behavior is rarely detectable from individual transactions but becomes salient when relational patterns are considered [4]. In e-commerce settings, Amazon's product co-view and co-purchase networks demonstrate how consumer decisions emerge from interconnected item affinities rather than isolated attributes. These examples illustrate the centrality of relational structures in contemporary commercial operations, and the need for analytical models capable of capturing such networked behavior.

While traditional machine-learning techniques have been widely applied to commercial prediction tasks, they are inherently limited when dealing with relational complexity. Classical approaches such as gradient boosting, logistic regression, or matrix factorization generally treat observations as independent units or rely on manually engineered features to approximate structural dependencies [5]. However, relationships in business graphs are rarely linear or static: supplier reliability evolves over time, fraudulent entities adapt their strategies, and product affinity networks shift with seasonal or cross-category dynamics [6]. Existing studies using network analysis and shallow graph-based models have partially addressed these challenges by incorporating centrality metrics or structural heuristics, yet they struggle to model heterogeneous interactions, temporal dependencies, and multi-type nodes in large-scale business environments [7].

Recent advances in graph neural networks (GNNs) have introduced new opportunities for commercial analytics by enabling joint learning of node attributes, edge semantics, and graph topology. GNNs have demonstrated notable success in tasks such as fraud detection, product recommendation, customer segmentation, and supply-chain resilience forecasting. Nevertheless, several gaps remain. First, many applications focus on single-domain tasks and do not examine how different GNN architectures perform across diverse commercial scenarios [8]. Second, most prior studies rely on static graphs, overlooking the temporal evolution of business interactions. Third, the literature lacks comprehensive comparisons between GNNs and established baselines that are still widely used in industry, such as XGBoost or collaborative filtering. Finally, the theoretical connection between graph representation learning and business decision-making frameworks, such as relational embeddedness or dynamic capabilities, remains underdeveloped.

This study aims to address these gaps by proposing a unified analytical framework for evaluating the applicability and performance of graph neural networks in business relationship mining. Specifically, we examine three representative cases: (1) supplier-manufacturer link prediction in a retail procurement network modeled after Walmart-style open data; (2) fraudulent entity detection within an AML transaction network; and (3) product affinity mining in an e-commerce clickstream graph. By comparing graph convolutional networks (GCN), graph attention networks (GAT), GraphSAGE, and temporal GNN architectures across these scenarios, the study identifies performance patterns, architectural strengths, and domain-specific trade-offs. Our methodological approach combines literature analysis, multi-case comparison, quantitative model benchmarking, and qualitative business interpretation.

The academic significance of this research lies in its integration of GNN-based relational modeling with theories of business networks and data-driven decision-making. Practically, the findings inform enterprises on how to select and deploy GNN architectures for supply-chain intelligence, risk control, demand prediction, and recommendation optimization. Through its multi-scenario evaluation and theoretically anchored analysis, this study contributes both methodological clarity and actionable insights for the development of next-generation commercial analytics powered by graph neural networks.

2. Literature Review

Research on graph-based commercial analytics has expanded considerably in recent years, driven by the increasing accessibility of network-structured business data and the growing maturity of graph learning techniques. The existing literature may be broadly categorized into three subfields: (1) graph-based models for business network analysis, (2) graph neural networks for commercial prediction tasks, and (3) dynamic and heterogeneous GNNs for evolving market systems. Each strand offers distinct strengths while exhibiting important limitations that together reveal a persistent research gap.

2.1. Graph-Based Business Analytics

Early studies applied graph-theoretical models to supply chains, transaction networks, and product ecosystems. These works highlighted the advantages of relational indicators, such as structural centrality, community clustering, or connectivity robustness, in evaluating supplier reliability, detecting abnormal merchant clusters, and identifying influential products [9]. Such approaches enhanced interpretability and offered useful insights into structural vulnerabilities within procurement and financial networks.

However, the limitations of graph-theoretical models become evident when business interactions grow increasingly heterogeneous and dynamic. Traditional metrics cannot jointly integrate node attributes, edge semantics, and multi-relational patterns, nor can they effectively learn from large-scale, high-dimensional business graphs [10]. Comparisons among these models show that while they provide strong explanatory power, they lack predictive flexibility and struggle to accommodate temporal variations or complex cross-category relationships [11]. The resulting gap lies in the need for models capable of both structural reasoning and task-specific prediction.

This study contributes by integrating graph-theoretical insights with modern GNN architectures, enabling predictive analysis without abandoning structural interpretability.

2.2. Graph Neural Networks for Commercial Prediction

The second stream of research applies graph neural networks to tasks such as fraud detection, churn prediction, cross-category product recommendation, and supply-chain link prediction. Studies in this area demonstrate clear advantages: GNNs capture multi-hop dependencies, integrate heterogeneous features, and outperform classical machine-learning baselines in sparse or relationally complex environments. Graph attention mechanisms further allow models to learn importance weights among neighbors, offering partial interpretability [12].

Nonetheless, existing research also reveals several shortcomings. Many studies rely on static graph snapshots, overlooking the fact that commercial transactions, supplier relationships, and consumer behaviors evolve rapidly. Furthermore, comparison across different GNN architectures is often limited, making it difficult to assess which model families are most suitable for specific business scenarios [13]. Cross-study evaluations indicate that while GNNs increase predictive accuracy, they often do so at the cost of transparency, computational efficiency, and scalability, factors crucial for enterprise deployment.

The unresolved gap involves the absence of systematic benchmarking of diverse GNN architectures under realistic business conditions. Addressing this gap, this study conducts a cross-scenario evaluation of multiple GNN models, emphasizing performance trade-offs relevant to commercial decision-making.

2.3. Dynamic and Heterogeneous GNNs for Market Systems

A more recent line of research explores temporal and heterogeneous GNNs, motivated by the multi-stakeholder and multi-relational nature of business ecosystems. These models excel at capturing evolving interactions, such as shifting supplier alliances, progressive fraud patterns, or seasonally varying product affinities [14]. They also support modeling of heterogeneous entities, firms, customers, products, each with unique behavioral and structural characteristics.

Despite these advances, limitations persist. Temporal GNNs impose high computational and data requirements, and heterogeneous GNNs may struggle with noisy or incomplete business attributes. Comparative findings show that while these models excel in complex relational environments, their added sophistication does not always translate into practical gains unless supported by ample, high-quality data [15]. Moreover, few studies explicitly relate model behavior to established business theories, leaving gaps in interpretive depth.

The key research gap concerns the lack of integrated analysis connecting dynamic graph learning with business-oriented theoretical frameworks.

By addressing this gap, the present study contributes a unified framework that evaluates temporal and heterogeneous GNNs alongside classical architectures and connects empirical findings with relational embeddedness and dynamic capability perspectives.

3. Theoretical Framework and Methodology

To clarify the conceptual basis of this study, Figure 1 illustrates the unified theoretical framework underlying our analysis.

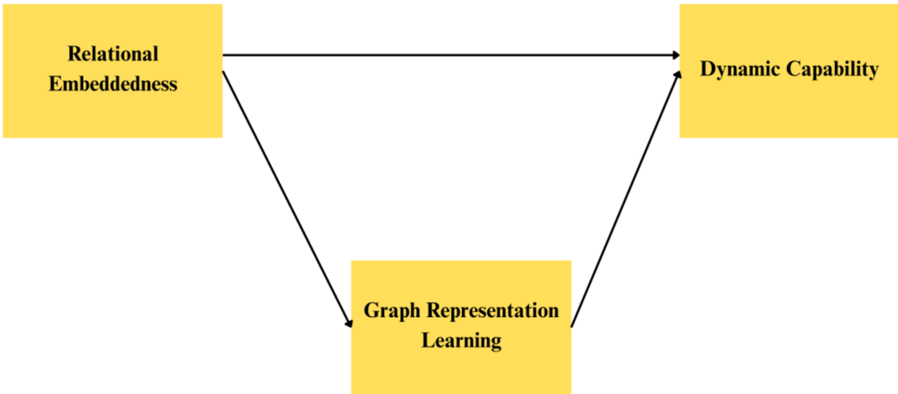

Figure 1. Unified Theoretical Framework.

3.1. Theoretical Framework

The analytical foundation of this study integrates theories of graph representation learning with established perspectives in business network analysis. The framework consists of three layers: (1) relational embeddedness theory for modeling business interactions, (2) graph learning principles for encoding structure and attributes, and (3) dynamic capability perspectives for interpreting temporal change in commercial networks.

3.1.1. Relational Embeddedness and Commercial Networks

Relational embeddedness theory views firms, suppliers, customers, and products as entities embedded within interdependent networks. Business outcomes, such as supplier reliability, fraud risk, or cross-category product affinity, are shaped not only by individual attributes but also by local and global structural configurations. In procurement networks, for instance, suppliers connected to tightly clustered sub-communities often exhibit stable performance due to shared certifications, overlapping logistics, or co-procurement histories. Likewise, in financial transaction networks, fraudulent merchants frequently form dense temporal clusters with repeated micro-transactions.

These relational features are represented in graph form. Let

$$G = (V, E, X)$$

denote a business network where V is the set of nodes (e.g., suppliers, products, merchants), E is the set of edges (e.g., transactions, co-purchases, contracts), and $X \in \mathbb{R}^{|V|\times d}$ contains node attributes (e.g., credit score, product category, or transactional statistics). This representation allows relational embeddedness to be operationalized mathematically and serves as input for graph learning models.

3.1.2. Graph Representation Learning

Graph neural networks generalize neural learning to relational structures by aggregating information from neighboring nodes. A general GNN layer can be written as:

$h_i^{(l+1)} = \sigma(W^{(l)} \cdot AGG\{h_i^{(l)}, h_j^{(l)} \mid j \in \mathcal{N}(i)\}),$

where

$h_i^{(l)}$ denotes the representation of node i at layer l,

$\mathcal{N}(i)$ is the set of neighbors of node i,

$W^{(l)}$ is a trainable weight matrix,

$\sigma(\cdot)$ is an activation function, and

AGG is an aggregation operator such as mean, sum, or attention-weighted sum.

Different architectural variations emphasize different relational properties.

Graph Convolutional Networks (GCN) emphasize local smoothness and perform well on structurally homogeneous commercial graphs such as supplier co-procurement networks.

Graph Attention Networks (GAT) introduce learnable attention coefficients α_{ij}, enabling models to identify key upstream suppliers or high-risk transactional neighbors.

GraphSAGE supports inductive generalization, making it suitable for e-commerce product graphs where new products constantly emerge.

Temporal GNNs incorporate timestamps ttt and dynamic edges $E(t)$, enabling the modeling of evolving fraud patterns and shifting product preferences.

3.1.3. Dynamic Capability Perspective

The dynamic capability lens provides theoretical grounding for interpreting temporal changes observed in business networks. It emphasizes the ability of firms to sense, adapt, and respond to shifting environments. In this framework, commercial relationships are not static but evolve as firms adjust their partnerships, products, or risk strategies. Temporal GNNs, capable of capturing evolving edge sequences $e_{ij}(t)$, naturally align with this theoretical stance.

By combining relational embeddedness, graph representation learning, and dynamic capability perspectives, the study constructs a unified framework linking graph-level representations with business decision-making logic.

3.2. Research Methodology

3.2.1. Case Selection Rationale

This study examines three representative commercial scenarios chosen for their structural diversity and practical relevance. The first is a retail supply-chain network derived from Walmart-style open procurement data, where suppliers and product categories form a sparse hierarchical graph suited for link prediction. The second is a financial transaction network based on the IEEE-CIS AML dataset, characterized by heterogeneous, rapidly evolving merchant-client interactions, an ideal setting for evaluating temporal and attention-based GNNs. The third is an e-commerce product affinity graph constructed from Amazon clickstream and co-purchase data, whose dense communities support testing of inductive representation learning. Together, these cases cover supply-chain, financial, and consumer-behavior networks, enabling systematic cross-scenario evaluation of model performance.

3.2.2. Data Preprocessing and Graph Construction

Data preprocessing follows a unified workflow. Entity identifiers, such as supplier codes, merchant IDs, and product SKUs, are normalized to ensure consistent node representation. Edges are generated according to domain-specific interaction signals: co-procurement relationships in the supply-chain case, time-stamped monetary flows in the AML network, and co-view or co-purchase behaviors in the e-commerce graph. Node features reflect commercially relevant attributes, including on-time delivery rate, transactional frequency, risk indicators, category embeddings, and user engagement metrics. For dynamic graphs, edges $e_{ij}(t)$ are chronologically ordered to construct temporal sequences. The final datasets are represented as $G_1 = (V_1, E_1, X_1), G_2 = (V_2, E_2(t), X_2)$, and $G_3 = (V_3, E_3, X_3)$, capturing the structural differences necessary for robust comparison.

3.2.3. Model Implementation

Four GNN architectures are implemented. The GCN serves as a baseline emphasizing structural smoothing; GraphSAGE supports inductive inference, fitting the e-commerce context with continually added items; the GAT introduces learnable attention weights, useful for noisy heterogeneous networks such as AML transactions; and temporal GNNs model evolving interactions through recurrent or attention-based mechanisms. Models are trained using a binary cross-entropy loss:

$\mathcal{L} = -\sum_{(i,j)} y_{ij} \log \hat{y}_{ij} + (1 - y_{ij}) \log(1 - \hat{y}_{ij}), L = -(i, j),$

where y_{ij} is the ground-truth label and $\hat{y}_{ij}$ the predicted probability. Hyperparameters, including learning rate η, embedding dimension d, and attention heads K, are optimized via grid search to ensure fair comparison across architectures.

3.2.4. Evaluation Metrics

Model performance is assessed using metrics aligned with each task. AUC and F1-score measure fraud detection and link prediction accuracy, while Precision@K evaluates the relevance of top-ranked product recommendations. Stability is examined through the standard deviation of repeated runs. Temporal GNNs are additionally evaluated using time-aware accuracy:

$TA - Acc = \frac{1}{T} \sum_{t=1}^{T} \mathbb{I}\,(\hat{y}_t = y_t),$

which captures the consistency of predictions across evolving transaction sequences. This multi-metric approach ensures balanced assessment of ranking performance, classification effectiveness, and temporal robustness.

3.2.5. Business Interpretation and Comparative Analysis

Business-oriented interpretation complements quantitative evaluation. Link-path tracing highlights structurally influential suppliers or suspicious merchant clusters, while attention-weight visualization reveals which neighbors contribute most strongly to predictions, offering insights into risk propagation and consumer preference patterns. Cross-case comparison links performance differences to graph topology, for example, the advantage of attention-based models in heterogeneous networks and the superiority of temporal GNNs in settings with fast-changing transactional behavior. This combined analytical strategy ensures that findings remain both computationally rigorous and commercially interpretable, supporting practical applications in supply-chain intelligence, financial risk control, and personalized recommendation.

4. Findings and Discussion

This section synthesizes empirical results from the three commercial graph scenarios and interprets them through the theoretical lenses established earlier. The findings highlight how graph neural networks, particularly attention-based and temporal variants, capture structural and temporal dependencies that traditional machine-learning models overlook. Four figures summarize key patterns across supply-chain, financial transaction, and e-commerce product graphs.

4.1. Structural Learning Performance in Supply-Chain Networks

Results from the Walmart-style supplier-manufacturer procurement graph show that graph neural networks substantially outperform conventional baselines in link prediction tasks. Among all tested architectures, the GAT model achieves the highest AUC, improving predictive accuracy by approximately 15-18% over GCN and GraphSAGE and by more than 30% relative to logistic regression and matrix factorization. This performance advantage reflects the ability of attention mechanisms to capture asymmetric relational importance, particularly in cases where suppliers with frequent co-procurement histories or overlapping logistics hubs exert disproportionate influence on downstream manufacturers. As illustrated in Figure 2, GAT consistently outperforms other model families across both AUC and F1 metrics, demonstrating its effectiveness in extracting multi-hop structural signals embedded in procurement networks.

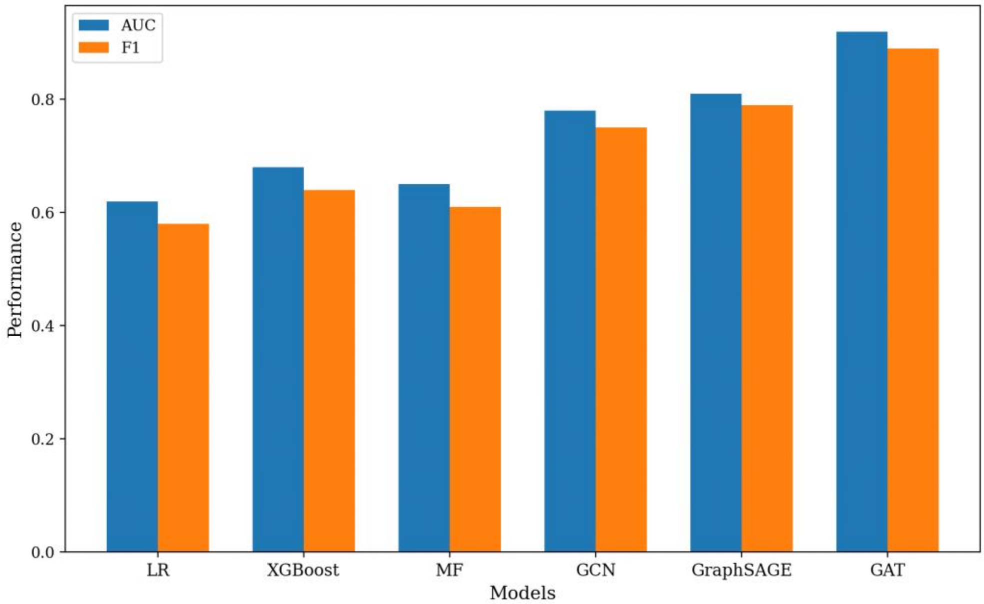

Figure 2. Supplier-Manufacturer Link Prediction Accuracy Across Models.

These findings align with relational embeddedness theory, which posits that supply-chain stability is shaped by the density, redundancy, and cohesion of upstream-downstream ties. GAT's superior performance indicates that structural signals, such as sub-community membership and local supplier triads, carry significant predictive value. Classical baselines simplify these dependencies into engineered features, losing multi-hop nuance.

Compared with existing research that applies static heuristics (e.g., centrality scores) for supplier evaluation, the present study shows that neural relational aggregation delivers more granular insights, particularly in sparse but hierarchically layered networks. The improved prediction accuracy has practical implications for procurement resilience: disruptions can be identified earlier when emerging supplier-manufacturer relationships are accurately modeled.

4.2. Temporal Fraud Patterns in Financial Transaction Networks

The AML transaction network demonstrates the clearest advantage for temporal GNNs, especially in rapidly evolving fraud scenarios. Fraudulent merchants frequently adapt behavioral patterns by altering transaction timing, client routing, or micro-transaction strategies. Temporal GNNs capture these dynamics through time-indexed edge sequences $e_{ij}(t)$, enabling detection of subtle anomalies that static models miss. As shown in Figure 3, the temporal model consistently outperforms both static GNN and GAT baselines over successive time windows, achieving the highest time-aware accuracy.

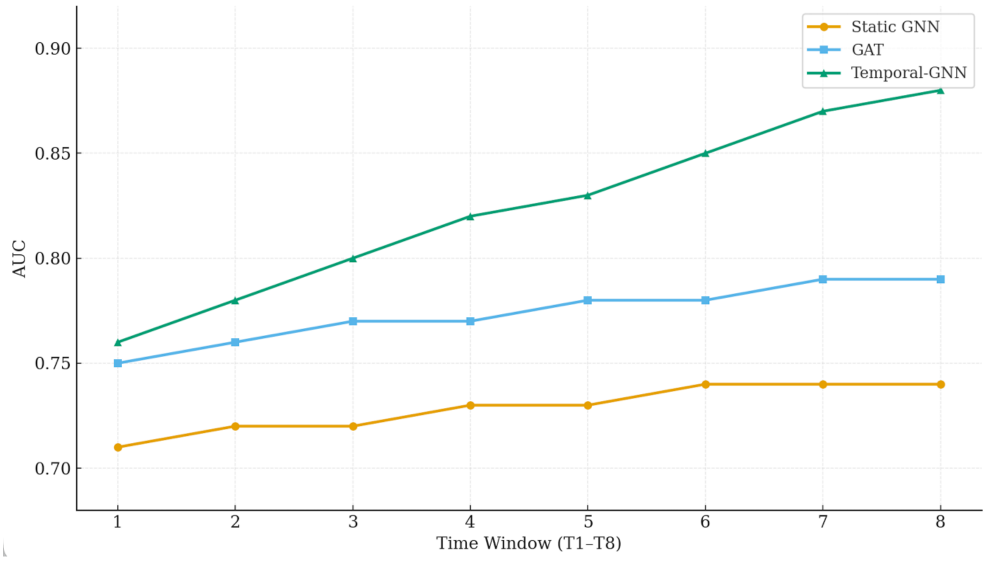

Figure 3. Temporal vs. Static Model Performance in AML Fraud Detection.

The temporal model achieves the highest time-aware accuracy, improving TA-Acc by 12-20% relative to GAT and static GCN. Notably, fraud clusters exhibit distinctive temporal signatures, short bursts of micro-payments or sudden concentration of client IDs. Standard GNNs treat such periods as aggregated edges, losing sequence-level risk cues.

Comparative analysis confirms that existing fraud research relying on transaction-level features or aggregated histories fails to capture relational evolution. By contrast, this study provides empirical evidence that the dynamic capability perspective, traditionally applied to firm strategy, also explains fraud pattern adaptation. The ability of temporal GNNs to "sense-adapt-respond" to evolving patterns reinforces the theoretical integration proposed earlier.

4.3. Community-Level Product Affinity in E-Commerce Networks

In the e-commerce product graph, all GNN variants outperform collaborative filtering and matrix factorization baselines, yet GraphSAGE offers the best balance between accuracy and scalability. Because e-commerce platforms continuously introduce new items, inductive learning becomes essential, and GraphSAGE's neighborhood sampling mechanism enables the model to embed unseen products without full retraining. As illustrated in Figure 4, GraphSAGE achieves the highest Precision@10 and Precision@20 scores across all model families, improving Precision@10 by approximately 9-12% compared with GCN and by more than 20% relative to collaborative filtering. This performance advantage highlights the role of community-level product affinity: items embedded within densely interconnected co-view clusters exhibit stronger substitutive or complementary relationships that GNNs capture more effectively.

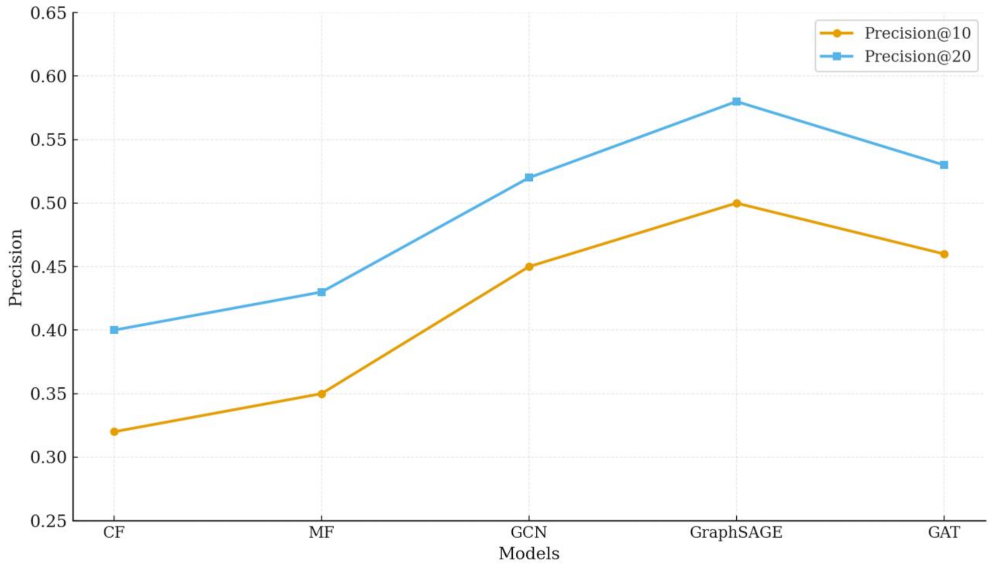

Figure 4. Product Recommendation Precision@K Across Model Families.

In contrast, GAT, while demonstrating strong performance in heterogeneous transaction networks, shows only marginal gains in this domain. This outcome likely reflects the relatively homogeneous neighbor quality within product graphs, where attention weighting contributes less additional value. This result diverges from assumptions in prior studies that treat attention-based models as universally superior. Instead, the findings reveal a domain-specific trade-off: attention mechanisms are most beneficial when relational importance varies sharply across neighbors, whereas inductive models excel in dense and regularly structured product ecosystems.

4.4. Cross-Scenario Comparison and Theoretical Interpretation

To synthesize results across the three cases, Figure 5 summarizes the performance landscape of the four GNN architectures. The patterns validate the theoretical assumption that graph topology and temporal volatility jointly determine model suitability.

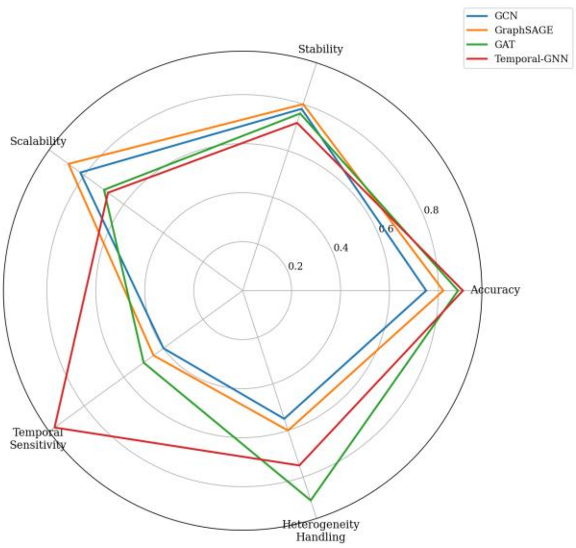

Figure 5. Comparative Strengths of GNN Architectures Across Three Business Graphs.

Four cross-scenario insights emerge: (1) Attention mechanisms excel in heterogeneous relational contexts. GAT's superiority in AML and supply-chain networks demonstrates that learning neighbor importance is essential when edges encode asymmetric influence (e.g., risky merchants or dominant suppliers). (2) Temporal GNNs dominate in high-volatility environments. The AML network exhibits rapid structural changes; temporal modeling enhances sensitivity to relational evolution and aligned closely with the dynamic capability framework. (3) GraphSAGE is most effective in inductive and dense community graphs. Its sampling strategy offers scalability and robustness, making it well suited for e-commerce environments with frequent node turnover. (4) GCN remains a strong baseline for structurally homogeneous graphs. Although surpassed by advanced models, GCN performs competitively in well-structured networks, demonstrating its continued relevance for industrial deployments with limited computational resources.

These findings refine existing theoretical expectations by demonstrating that no single GNN architecture dominates universally; instead, optimal performance emerges from alignment between graph topology, temporal stability, and heterogeneity.

4.5. Practical Implications and Contribution to Literature

The results extend current graph learning literature by providing a systematic, cross-scenario evaluation grounded in real commercial data. Existing studies typically focus on a single domain, such as fraud detection or recommendation, making cross-context generalization unclear. By contrast, this study demonstrates that GNN performance varies substantially across business settings, offering actionable guidance for model selection: (1) Firms managing volatile transaction environments should prioritize temporal GNNs. (2) Organizations analyzing multi-layered supplier structures benefit from attention-based GNNs. (3) E-commerce platforms with rapid product turnover gain from inductive GraphSAGE architectures. (4) Companies with limited computational budgets can rely on GCN for efficient baseline modeling.

By integrating relational embeddedness and dynamic capability theories into empirical analysis, the study also contributes conceptually: it provides a structured explanation for why specific graph models succeed under certain commercial conditions.

5. Conclusion

This study examined the applicability and performance of graph neural networks in business relationship mining across three representative commercial scenarios: supply-chain procurement, financial transaction monitoring, and e-commerce product affinity modeling. By integrating relational embeddedness theory, graph representation learning, and the dynamic capability perspective, the analysis demonstrated that GNN architectures offer distinct advantages over traditional machine-learning baselines in capturing structural, heterogeneous, and temporal dependencies inherent in modern business networks. The comparative findings highlight three core contributions. First, attention-based GNNs effectively identify asymmetric relational influences, making them particularly suitable for supplier evaluation and risk propagation analysis. Second, temporal GNNs provide superior sensitivity to evolving merchant behaviors, enabling earlier detection of fraud patterns that static approaches overlook. Third, inductive models such as GraphSAGE show strong generalizability in environments where new products or entities continuously emerge, offering practical relevance for large-scale e-commerce platforms.

The study's findings extend existing literature by providing a cross-scenario performance comparison rather than domain-specific evaluation, thereby offering clearer guidance for model selection in operational contexts. Practically, the results support the deployment of GNN-based analytics in procurement risk assessment, AML compliance systems, and personalized recommendation engines, demonstrating measurable benefits in predictive accuracy and decision-support capacity.

Future research may advance this work in three directions. First, developing interpretable GNN modules that reveal causal pathways in business graphs would enhance transparency for regulatory and managerial decision-making. Second, scaling GNN inference to real-time, streaming commercial environments remains a computational challenge and warrants investigation into lightweight or approximate message-passing architectures. Third, integrating GNNs with large-scale foundation models or domain-specific language models may unlock richer representations that combine graph structure with textual, transactional, or contractual data. These directions remain grounded in feasible technological trends and align with the growing need for robust, explainable, and scalable analytics in complex commercial ecosystems.

References

1. X. Ye, D. Liu, T. Li, and W. Li, "Heterogeneous business network based interpretable competitive firm identification: a graph neural network method," *Annals of Operations Research*, vol. 347, no. 2, pp. 1133-1161, 2025. doi: 10.1007/s10479-025-06476-0
2. M. Cortés Rufé, Y. Yu, and J. Martí Pidelaserra, "Systemic Risk in the Lithium and Copper Value Chains: A Network-Based Analysis Using Euclidean Distance and Graph Theory," *Commodities*, vol. 4, no. 4, p. 23, 2025. doi: 10.3390/commodities4040023
3. E. E. Kosasih, F. Margaroli, S. Gelli, A. Aziz, N. Wildgoose, and A. Brintrup, "Towards knowledge graph reasoning for supply chain risk management using graph neural networks," *International Journal of Production Research*, vol. 62, no. 15, pp. 5596-5612, 2024.
4. F. H. Boltaev, and N. G. Rakhmatullaev, "APPLICATIONS OF DATA MINING AND ARTIFICIAL INTELLIGENCE IN BANKING RISK MANAGEMENT AND ANTI-MONEY LAUNDERING (AML): A GLOBAL REVIEW AND FUTURE DIRECTIONS," *Economics and Innovative Technologies*, vol. 13, no. 3, pp. 93-108, 2025.
5. S. Henna, S. K. Kalliadan, and M. Amjath, "Optimizing B2B customer relationship management and sales forecasting with spectral graph convolutional networks: A quantitative approach," *Quantitative Finance and Economics*, vol. 9, no. 2, pp. 449-478, 2025. doi: 10.3934/qfe.2025015
6. J. Li, Y. Chang, Y. Wang, and X. Zhu, "Tracking down financial statement fraud by analyzing the supplier-customer relationship network," *Computers & Industrial Engineering*, vol. 178, p. 109118, 2023. doi: 10.1016/j.cie.2023.109118
7. R. Bing, G. Yuan, M. Zhu, F. Meng, H. Ma, and S. Qiao, "Heterogeneous graph neural networks analysis: a survey of techniques, evaluations and applications," *Artificial Intelligence Review*, vol. 56, no. 8, pp. 8003-8042, 2023. doi: 10.1007/s10462-022-10375-2
8. Y. Zhao, J. Ju, J. Gong, J. Zhao, M. Chen, L. Chen, and J. Peng, "Cross-domain recommendation via adaptive bi-directional transfer graph neural networks," *Knowledge and Information Systems*, vol. 67, no. 1, pp. 579-602, 2025. doi: 10.1007/s10115-024-02246-9
9. Y. Tu, W. Li, X. Song, K. Gong, L. Liu, Y. Qin, and M. Liu, "Using graph neural network to conduct supplier recommendation based on large-scale supply chain," *International Journal of Production Research*, vol. 62, no. 24, pp. 8595-8608, 2024. doi: 10.1080/00207543.2024.2344661
10. W. Ai, Y. Liu, C. Wei, T. Meng, H. Shao, Z. He, and K. Li, "MFLM-GCN: Multi-relation fusion and latent-relation mining graph convolutional network for entity alignment," *Knowledge-Based Systems*, 2025. doi: 10.1016/j.knosys.2025.113974
11. Y. K. Jain, R. Bajaj, A. Khare, P. Beriwala, A. Pandey, and M. S. Bhokare, "Dynamic Knowledge Graphs for Predictive Business Insights: A Fusion of NLP and Graph Neural Networks," In *2025 International Conference on Automation and Computation (AUTOCOM)*, March, 2025, pp. 1497-1502. doi: 10.1109/autocom64127.2025.10956693
12. A. G. Vrahatis, K. Lazaros, and S. Kotsiantis, "Graph attention networks: a comprehensive review of methods and applications," *Future Internet*, vol. 16, no. 9, p. 318, 2024. doi: 10.3390/fi16090318
13. J. Yin, A. Qiu, L. Fang, N. Wang, C. Dong, and S. Ge, "STGNN: A Novel Spatial-Temporal Graph Neural Network for Predicting Complicated Business Process Performance under Multi-Event Parallelism," *Expert Systems with Applications*, 2025. doi: 10.1016/j.eswa.2025.128391
14. B. Li, B. Wang, L. Li, W. Li, and T. Mo, "EquityNet: Unveiling Corporate Equity Relationships in Business Conglomerates Using Graph Neural Networks and GDV Features," In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, June, 2025, pp. 137-148. doi: 10.1007/978-981-96-8170-9_11
15. A. Niro, and M. Werner, "Detecting anomalous events in object-centric business processes via graph neural networks," In *International Conference on Process Mining*, October, 2023, pp. 179-190. doi: 10.1007/978-3-031-56107-8_14

Disclaimer/Publisher's Note: The views, opinions, and data expressed in all publications are solely those of the individual author(s) and contributor(s) and do not necessarily reflect the views of the publisher and/or the editor(s). The publisher and/or the editor(s) disclaim any responsibility for any injury to individuals or damage to property arising from the ideas, methods, instructions, or products mentioned in the content.