The American Journal of Management and Economics Innovations
ISSN 2693-0811 Volume 08 - 2026

Adaptive Graph-Temporal Deep Learning for Cascading Supply Chain Disruption Prediction and Resilience Optimization

Mohammad Kawsur Sharif
Washington University of Virginia, Department of Management Information Systems

Fariha Noor Nitu
Master of Science in Management Science & Supply Chain Management, Wichita State University

Asaduzzaman Anik
Master of Business Administration (MBA) in management, Stanton University, Los Angeles, California

Marjahan Risalat
Department of Business Administration and Management, Stanton University, USA

Md Yousuf
Master of Science in Data Analytics, College of Sciences and Technology, University of Houston-Downtown, USA.

Md Juweel Rana
Bachelor Of Business Administration, International American University, Los Angeles, California, USA

Received: 02 October 2026 | Received Revised Version: 01 November 2026 | Accepted: 09 November 2026 | Published: 19 November 2026

Volume 08 Issue 11 2026 | DOI:10.37547/tajmei/Volume08Issue11-02

Abstract

Supply-chain disruptions and shipment delays threaten logistics efficiency, operational continuity, and supply-chain resilience. Existing predictive approaches often examine transportation dependencies and temporal shipment patterns independently, limiting their ability to capture the complex relationships underlying operational delays. To address this limitation, we propose an Adaptive Graph-Temporal Resilience Network (AGTR-Net), an integrated deep-learning framework that combines graph neural networks, temporal sequence modeling, adaptive feature fusion, and self-supervised representation learning for shipment-delay prediction. The proposed framework is designed to capture interconnected transportation activities, identify evolving operational patterns, and provide explainable and uncertainty-aware risk assessments. We evaluate the framework using the Cargo 2000 Freight Tracking and Tracing dataset from the UCI Machine Learning Repository and compare its performance against conventional machine-learning algorithms, temporal deep-learning models, and graph neural network architectures. In the illustrative experimental results, AGTR-Net achieves an accuracy of 90.1%, precision of 87.2%, recall of 85.6%, and F1-score of 86.4%, with a ROC-AUC of 0.945 and PR-AUC of 0.901. It outperforms the strongest standalone graph-neural-network baseline, Graph Attention Network (GAT), which achieves an F1-score of 81.9% and PR-AUC of 0.851. Furthermore, AGTR-Net demonstrates improved probability calibration, with a Brier score of 0.098 and an expected calibration error of 0.026 in the illustrative evaluation. Ablation analysis further indicates the potential value of integrating graph representations, temporal learning, adaptive fusion, and self-supervised learning within a unified architecture. We also introduce a simulation-based recovery optimization component to explore how predictive risk assessments can support operational intervention and shipment recovery decisions. Overall, the proposed framework offers a unified approach to shipment-delay prediction and resilience-oriented logistics decision support, providing a foundation for more proactive and data-driven supply-chain operations. Further validation across diverse transportation networks and real-world disruption scenarios is necessary to establish its broader applicability.

Keywords: Supply Chain Resilience; Shipment Delay Prediction; Graph Neural Networks; Temporal Deep Learning; Adaptive Feature Fusion; Self-Supervised Learning; Explainable Artificial Intelligence; Uncertainty Quantification; Logistics Optimization.

© **2026** Fariha Noor Nitu,Asaduzzaman Anik,Marjahan Risalat,Md Yousuf,Md Juweel Rana. This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). The authors retain copyright and allow others to share, adapt, or redistribute the work with proper attribution.

Cite This Article: Mohammad Kawsur Sharif, Fariha Noor Nitu, Asaduzzaman Anik, Marjahan Risalat, Md Yousuf, & Md Juweel Rana. (2026). Adaptive Graph-Temporal Deep Learning for Cascading Supply Chain Disruption Prediction and Resilience Optimization. The American Journal of Management and Economics Innovations, 8(11), 16–46. https://doi.org/10.37547/tajmei/Volume08Issue11-02

1. Introduction

1.1. Background and Research Motivation

Global supply chains have become increasingly interconnected through international sourcing, manufacturing, transportation, warehousing, and distribution networks. Although this interconnectedness has improved operational efficiency and enabled organizations to access global markets, it has also increased their exposure to unexpected disruptions. Natural disasters, geopolitical conflicts, transportation interruptions, economic instability, public health emergencies, cyber incidents, and supplier failures can interrupt the movement of goods and materials, resulting in delivery delays, production interruptions, increased operational costs, and reduced customer satisfaction. Consequently, supply-chain resilience has become an important strategic capability for organizations seeking to maintain operational continuity under uncertain conditions.

Traditional supply-chain management approaches have primarily emphasized cost reduction, inventory efficiency, lean operations, and predictable delivery performance. However, these approaches may provide insufficient protection when disruptions affect multiple interconnected transportation and production activities. Christopher and Peck (2004) emphasized that resilient supply chains require organizations to develop capabilities for anticipating, responding to, and recovering from unexpected disturbances. Similarly, Tang (2006) demonstrated the importance of integrating risk-management strategies into supply-chain design and operational decision-making. These perspectives suggest that supply-chain resilience should not be considered exclusively as a recovery capability but also as an anticipatory capability that enables organizations to identify potential risks before they become significant operational problems.

The importance of supply-chain resilience has become particularly evident following the COVID-19 pandemic, which exposed vulnerabilities in global manufacturing, transportation, procurement, and distribution systems. Ivanov (2020) demonstrated how pandemic-related disruptions could propagate across interconnected supply networks and produce substantial operational consequences. These disruptions also highlighted the limitations of relying exclusively on historical averages, static risk assessments, and traditional forecasting methods. Organizations increasingly require predictive systems capable of processing operational information, identifying emerging risk patterns, and supporting timely decisions.

Within this context, artificial intelligence (AI) and deep learning have created new opportunities for improving supply-chain risk prediction. Machine-learning algorithms can identify nonlinear relationships among operational variables, while deep-learning architectures can learn complex temporal patterns from sequential data. Baryannis et al. (2019) identified artificial intelligence as an important research direction in supply-chain risk management, particularly for risk identification, assessment, and response. More recently, Zogaan et al. (2025) investigated the application of multiple deep-learning architectures to supply-chain risk prediction across critical industrial sectors, demonstrating the growing relevance of deep learning in resilience-oriented applications.

Despite these developments, accurately predicting operational disruptions remains challenging because supply-chain risks are not necessarily isolated events. A delay at one transportation stage may affect subsequent logistics activities, while disruptions involving highly connected transportation or supplier relationships may influence multiple downstream operations. Consequently, prediction models should consider both the sequential development of operational activities and the structural relationships among interconnected supply-chain entities.

We argue that combining graph-based learning with temporal deep learning provides a promising approach to addressing these challenges. Graph neural networks can

represent relationships among transportation nodes and operational entities, whereas temporal architectures can learn how operational conditions evolve over time. The integration of these approaches may provide a more comprehensive representation of logistics risk than models that rely exclusively on independent shipment records or static operational variables.

1.2. Supply-Chain Disruptions and the Need for Predictive Resilience

Supply-chain disruptions can be understood as unexpected events that interfere with the normal movement of products, materials, information, or resources across interconnected operational activities. Their consequences may extend beyond the location where the original disturbance occurs. For example, a transportation delay can affect subsequent shipment transfers, warehouse receiving activities, inventory replenishment, and ultimately customer delivery performance.

The propagation of disruption effects is frequently discussed through the concept of the ripple effect. Hosseini et al. (2019) examined quantitative approaches for evaluating supply-chain resilience and emphasized the importance of understanding vulnerability, response, and recovery capabilities. Hosseini et al. (2020) further investigated supplier-disruption propagation using an integrated Markov-chain and dynamic Bayesian-network approach. Their research demonstrated the importance of examining relationships among interconnected supply-chain entities rather than treating individual disruptions as independent operational events.

Ivanov and Dolgui (2021) also discussed the role of operations research in understanding and mitigating ripple effects in supply chains during the COVID-19 pandemic. Their work highlights the importance of evaluating interconnected operational dependencies and identifying strategies that can maintain supply-chain performance under unexpected disturbances.

However, many operational prediction systems continue to focus primarily on individual shipment characteristics, historical delivery performance, or static risk indicators. Although these approaches can identify shipment-level abnormalities, they may not adequately represent the relationships among transportation stages and their potential downstream consequences.

We therefore consider predictive resilience as a combination of early risk identification, operational dependency analysis, and recovery-oriented decision support. Our research focuses on developing a framework that can identify potential future shipment delays using historical operational information and transportation-network characteristics. We use operational delay as an observable proxy for logistics disruption because the selected public dataset provides planned and actual activity durations rather than comprehensive records of geopolitical, natural-disaster, or supplier-failure events.

1.3. Artificial Intelligence and Deep Learning in Supply-Chain Risk Prediction

Artificial intelligence has increasingly been applied to supply-chain risk management because of its ability to process large volumes of heterogeneous operational information and identify relationships that may not be easily captured through conventional statistical techniques. Baryannis et al. (2019) reviewed the application of artificial intelligence in supply-chain risk management and discussed its potential contribution to risk identification, assessment, and response.

Traditional machine-learning approaches, including Random Forest, Support Vector Machines, Gradient Boosting, XGBoost, and LightGBM, have demonstrated their usefulness in classification and prediction problems involving structured operational data. These approaches can identify nonlinear relationships among transportation characteristics, planned durations, shipment performance, and historical delays. However, conventional algorithms generally require carefully engineered features and may have limited ability to learn complex sequential representations or interconnected network relationships directly.

Deep-learning architectures provide additional opportunities for modeling complex operational patterns. Recurrent neural networks, particularly Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) models, can capture dependencies across sequential shipment activities. These architectures are relevant to logistics applications because transportation operations typically involve multiple consecutive milestones, including freight check-in, departure, arrival, transfer, and final delivery.

Temporal Convolutional Networks provide an alternative approach by learning temporal dependencies through convolutional operations. Transformer architectures further extend sequential learning through self-attention mechanisms, allowing models to assign different importance to individual observations within a sequence (Vaswani et al., 2017). This capability is particularly relevant when earlier shipment activities may provide important information about the risk associated with subsequent operational milestones.

Recent research has also examined deep-learning approaches specifically for supply-chain resilience. Zogaan et al. (2025) investigated recurrent neural networks, LSTM, GRU, convolutional neural networks, and artificial neural networks across pharmaceutical, automotive, agricultural, and energy-related applications. Their study illustrates the potential of deep-learning methods for identifying operational risks across different industrial contexts.

Nevertheless, we identify an important opportunity to extend conventional temporal learning by incorporating transportation-network dependencies. Shipment delays are not always determined by the operational history of an individual shipment. They may also be associated with the characteristics of connected transportation nodes, observed routes, and preceding logistics activities. This motivates our integration of graph-based and temporal representations within a unified predictive architecture.

1.4. Research Objectives and Contributions

The primary objective of this research is to develop an adaptive graph-temporal deep-learning framework for predicting logistics disruptions and supporting resilience-oriented decision-making. We propose the Adaptive Graph-Temporal Resilience Network (AGTR-Net), which combines graph neural networks, Transformer-based temporal representation learning, adaptive feature fusion, and uncertainty-aware prediction.

We pursue this objective by examining how transportation-network information and sequential operational features can be integrated to improve shipment-delay prediction. We also investigate whether self-supervised representation learning can improve the extraction of useful patterns from limited operational data. Furthermore, we examine how explainability techniques can identify influential operational variables and transportation relationships associated with elevated delay risk.

Our research also investigates the potential connection between predictive risk identification and recovery optimization. We propose a simulation-based reinforcement-learning component to evaluate how early-warning information could support operational decisions. This component is designed as a decision-support extension rather than a direct assumption that historical shipment records contain all information required to optimize real-world recovery actions.

We organize the research around the following objectives. We investigate the contribution of graph neural networks to representing transportation relationships and identifying network-related risk patterns. We examine the ability of temporal Transformers to capture sequential dependencies among shipment activities. We evaluate whether adaptive graph-temporal fusion improves predictive performance compared with standalone graph and temporal architectures. We also assess the contribution of self-supervised learning, uncertainty estimation, and explainability to the reliability of the proposed model.

Finally, we investigate the potential application of the proposed framework in U.S. airfreight, manufacturing, and distribution operations. We aim to establish a methodological foundation for integrating predictive analytics with operational resilience planning.

1.5. Research Novelty and Expected Contributions

We position the contribution of this research in the integration and adaptation of graph-temporal learning for logistics-process risk prediction rather than claiming that graph neural networks, Transformers, or reinforcement learning are individually new technologies.

First, we develop an adaptive graph-temporal architecture that integrates transportation-network representations with sequential shipment-process information. Unlike approaches that rely exclusively on structured shipment features or independent temporal sequences, the proposed architecture combines network context and operational history within a unified predictive representation.

Second, we investigate self-supervised representation learning as a mechanism for improving feature learning from limited shipment-process observations. This is relevant because publicly available logistics datasets may contain fewer independent shipment processes than the volume of activity records initially suggests.

Third, we introduce an adaptive feature-fusion mechanism that allows the model to combine graph and temporal representations according to their predictive contribution. We further investigate uncertainty-aware prediction to assess the reliability of estimated shipment-delay probabilities.

Fourth, we incorporate explainability into the predictive architecture to identify influential operational characteristics and transportation relationships. This can improve the interpretability of predictions and provide additional information for logistics planners.

Finally, we investigate the integration of predictive delay estimation with simulation-based recovery optimization. This connects operational risk identification with potential intervention strategies and provides a foundation for

further research into predictive and prescriptive supply-chain analytics.

2. Literature Review

2.1. Supply-Chain Risk Management and Resilience

Supply-chain risk management has developed into an important research area because organizations increasingly operate through interconnected sourcing, production, transportation, and distribution systems. These systems are exposed to uncertainties that may affect the availability of materials, the reliability of suppliers, transportation performance, and customer service.

Tang (2006) provided an influential review of supply-chain risk-management perspectives and discussed strategies for managing risks associated with uncertain demand, supply, and operational conditions. The research emphasized the importance of designing supply-chain strategies that balance operational efficiency with risk mitigation. Christopher and Peck (2004) similarly argued that supply-chain resilience requires organizations to develop capabilities that enable them to anticipate disturbances, respond effectively, and recover operational performance.

Hosseini et al. (2019) reviewed quantitative methods for supply-chain resilience analysis and classified approaches according to their ability to evaluate resilience capacity. Their work demonstrates that resilience involves more than identifying the probability of a disruption. It also involves understanding how supply chains withstand disturbances, adapt to changing conditions, and recover their operational capabilities.

Ivanov (2020) extended the discussion by examining supply-chain viability under the COVID-19 pandemic. The study highlighted the importance of adaptability, resilience, and sustainability in maintaining the long-term viability of supply networks. This perspective is relevant to our research because it suggests that predictive systems should support not only immediate risk identification but also broader operational adaptation.

We identify predictive analytics as an important extension of traditional risk-management approaches. Conventional resilience frameworks often emphasize mitigation strategies and recovery planning, whereas predictive models can provide information about emerging operational risks before they become critical. Our research builds on this perspective by developing a data-driven framework for identifying shipment-delay risk through transportation-network and temporal information.

2.2. Disruption Propagation and the Supply-Chain Ripple Effect

The ripple effect describes how a disruption originating in one part of a supply chain may spread through interconnected relationships and influence downstream operational performance. The concept is particularly important in complex supply networks where suppliers, manufacturers, logistics providers, and distribution facilities depend on one another.

Ivanov (2020) demonstrated the significance of disruption propagation during the COVID-19 pandemic and discussed how disturbances could affect interconnected supply-chain activities. Hosseini et al. (2020) developed an integrated Markov-chain and dynamic Bayesian-network approach to examine supplier-disruption propagation. Their research provided a quantitative perspective on how disruptions may affect multiple supply-chain stages and how vulnerability and recovery characteristics influence the overall consequences.

Ivanov and Dolgui (2021) further examined operations-research methods for addressing ripple effects during the pandemic. Their work emphasized the importance of evaluating system-wide consequences and designing response strategies that account for interdependencies among supply-chain activities.

More recent research has increasingly explored network-oriented approaches to disruption analysis. Ma et al. (2025) reviewed supply-chain resilience from a network-modeling perspective and emphasized the importance of understanding structural relationships and network vulnerabilities. An et al. (2026) also examined supply-chain risk propagation through a graph-neural-network perspective, highlighting the relevance of graph-based approaches for analyzing interconnected risk mechanisms.

We consider these findings important because transportation and logistics systems naturally contain network structures. Airports, ports, warehouses, distribution centers, and transportation routes are connected through operational relationships. A delay affecting one transportation stage may influence subsequent shipment activities, creating an opportunity to use network-based learning for operational risk prediction.

However, we distinguish between the general supplier-disruption ripple effect and the transportation-process dependencies available in our empirical dataset. The Cargo 2000 dataset supports the construction of a transportation-process graph from observed shipment movements, but it does not provide a complete multi-tier supplier network. Accordingly, we use graph learning to

investigate transportation-related dependencies rather than claiming that our primary experiment directly measures supplier-level disruption propagation.

2.3. Machine Learning for Supply-Chain Disruption Prediction

Machine learning has been increasingly applied to supply-chain risk identification, demand forecasting, supplier assessment, and logistics-performance prediction. Compared with traditional statistical approaches, machine-learning algorithms can identify complex nonlinear relationships among operational variables and provide flexible prediction models.

Baryannis et al. (2019) reviewed the application of artificial intelligence in supply-chain risk management and classified existing approaches according to risk identification, assessment, and response. Their findings indicate that machine learning and related AI techniques provide opportunities for processing multidimensional information and supporting more adaptive decision-making.

Tree-based algorithms such as Random Forest, XGBoost, LightGBM, and CatBoost are particularly relevant for structured logistics datasets. They can identify nonlinear relationships among planned durations, transportation characteristics, activity types, and historical performance indicators. They also provide useful baseline models because they can perform effectively without requiring extensive sequential or graph-based architecture development.

Nevertheless, conventional machine-learning algorithms generally depend on manually constructed features. They may not directly learn complex relationships between operational sequences and transportation-network structures. Their predictive performance can also depend heavily on the quality of feature engineering and the representativeness of the training observations.

We therefore use conventional machine-learning methods as comparative benchmarks rather than assuming that deep learning will automatically outperform them. This comparison enables us to determine whether the additional complexity of AGTR-Net provides measurable predictive benefits.

2.4. Deep Learning and Temporal Sequence Modeling

Deep learning has become increasingly relevant to operational forecasting because it can learn hierarchical representations from large and complex datasets. In logistics systems, operational events typically occur sequentially, and the characteristics of earlier activities may influence the performance of subsequent activities.

LSTM networks were introduced by Hochreiter and Schmidhuber (1997) to address limitations associated with learning long-range dependencies in recurrent neural networks. Their memory-cell architecture enables models to preserve relevant information across sequential observations. GRU networks provide an alternative recurrent architecture that uses gating mechanisms to control information flow.

These models are relevant to shipment-process analysis because transportation operations consist of ordered activities. For example, delays during freight check-in or departure may influence subsequent arrival and delivery milestones. By learning from previous operational observations, recurrent architectures can estimate the likelihood of future delays.

Transformer architectures provide another approach to sequential learning. Vaswani et al. (2017) introduced a self-attention-based architecture that enables models to learn relationships between sequence elements without relying exclusively on recurrent processing. This mechanism allows the model to assign different attention weights to historical activities according to their relevance to the prediction task.

We consider Transformer-based temporal learning particularly appropriate for the proposed research because shipment-process sequences may contain operational activities with different predictive importance. However, the limited number of independent shipment-process instances in the selected dataset requires careful model regularization and experimental evaluation.

We therefore compare LSTM, GRU, TCN, and Transformer architectures to investigate the contribution of different temporal-learning mechanisms. We subsequently integrate the selected temporal representation into the graph-temporal architecture.

2.5. Graph Neural Networks in Supply-Chain and Logistics Modeling

Supply chains are naturally represented as networks because their operational structure consists of interconnected entities and relationships. Traditional machine-learning models frequently represent each observation as an independent feature vector, which may result in the loss of useful information about relationships among transportation nodes and operational entities.

Graph neural networks address this limitation by learning representations directly from graph-structured data. Graph

convolutional networks aggregate information from neighboring nodes, allowing the representation of a node to incorporate information from its surrounding network structure.

Graph Attention Networks provide a further development by assigning different importance weights to neighboring nodes. Veličković et al. (2018) introduced attention-based graph learning to enable models to learn which neighboring nodes contribute more strongly to a node representation. This is relevant to transportation networks because different routes and connected transportation facilities may have different operational characteristics.

Recent studies have increasingly investigated graph neural networks for supply-chain risk identification and resilience analysis. Kayalvizhi et al. (2025) proposed a graph-neural-network-based framework for modeling supply-chain resilience under multimodal disruptions. An et al. (2026) examined graph-neural-network perspectives on supply-chain risk propagation, illustrating the increasing interest in network-aware risk modeling.

These studies provide a foundation for investigating graph-based learning in supply-chain applications. However, the construction of a graph alone does not guarantee better predictive performance. The usefulness of graph learning depends on the accuracy of the underlying relationships, the availability of meaningful node and edge attributes, and the suitability of the prediction task.

In our research, we construct a directed transportation graph from the observed airport identifiers and shipment movements available in the Cargo 2000 dataset. We use graph neural networks to learn transportation-network representations and investigate whether these representations improve shipment-delay prediction when combined with temporal operational information.

2.6. Hybrid Graph-Temporal Deep Learning

Graph-based learning and temporal deep learning address different but complementary aspects of supply-chain risk prediction. Graph neural networks represent relationships among connected entities, whereas temporal models learn how operational information changes over sequential observations.

A graph-only architecture may identify important transportation relationships but may not adequately capture the temporal progression of individual shipment activities. Conversely, a temporal-only architecture may learn shipment-level patterns while overlooking relationships among transportation nodes.

Combining these representations provides an opportunity to learn both structural and sequential dependencies. This approach is especially relevant to transportation operations, where the progression of a shipment is influenced by its previous activities and the characteristics of its route.

Recent research has also explored graph-based risk inference and ensemble forecasting for logistics disruption prediction. Zhao et al. (2026) investigated an intelligent expert system integrating graph-based risk inference with ensemble forecasting. This work reflects the growing interest in combining network information and predictive analytics for logistics disruption management.

We extend this research direction by proposing an adaptive graph-temporal fusion architecture. Rather than treating graph and temporal representations as independent inputs, AGTR-Net uses a trainable fusion mechanism to adjust their relative contributions to the final prediction.

We evaluate this approach against standalone graph models, standalone temporal models, and conventional feature-concatenation architectures. Through these comparisons, we aim to establish whether adaptive integration provides measurable benefits for shipment-delay prediction.

2.7. Self-Supervised Learning and Limited Operational Data

One of the challenges associated with deep learning in supply-chain applications is the limited availability of labeled disruption observations. Operational datasets may contain a large number of recorded activities but a substantially smaller number of independent shipment processes or verified disruption events.

Self-supervised learning provides an opportunity to learn useful representations from unlabeled or partially labeled data. Rather than depending exclusively on manually assigned target labels, self-supervised approaches construct learning objectives from the underlying structure of the data.

For sequential logistics data, masked-feature reconstruction can train a model to recover selected operational attributes from incomplete sequences. Temporal contrastive learning can encourage representations of related shipment observations to remain similar while distinguishing unrelated observations.

For graph-structured information, graph reconstruction and link-prediction objectives can be used to learn structural representations from observed relationships.

These representations can subsequently initialize supervised prediction models.

We consider self-supervised learning particularly relevant to the Cargo 2000 dataset because the number of independent shipment processes is limited relative to the number of recorded operational activities. We therefore investigate whether pretraining can improve the quality of graph-temporal representations before supervised delay prediction.

We also recognize that self-supervised learning does not automatically solve data scarcity or guarantee improved generalization. Its contribution must be evaluated through controlled experiments that compare pretrained and non-pretrained versions of the same architecture.

2.8. Explainable Artificial Intelligence and Uncertainty-Aware Prediction

Although deep-learning models can provide strong predictive capabilities, their internal complexity can make it difficult for operational managers to understand why a shipment has been classified as high risk. This limitation can reduce the practical usefulness of predictive systems, particularly when predictions influence transportation planning, shipment prioritization, or additional operational expenditure.

Explainable artificial intelligence provides techniques for examining the variables and relationships that contribute to model predictions. SHAP and Integrated Gradients can be used to investigate the contribution of individual operational features, while graph-explanation methods can identify influential nodes and relationships within a graph-based prediction.

We incorporate explainability to examine whether variables such as cumulative delay, recent activity performance, transportation-leg complexity, and network connectivity contribute to elevated shipment-delay predictions. We also investigate the consistency of explanations across different observations.

Uncertainty estimation provides an additional dimension of prediction reliability. A model may generate a high predicted probability even when its input information is incomplete or substantially different from the training observations. Probability calibration and uncertainty-aware prediction can help distinguish reliable estimates from predictions that require additional operational investigation.

We therefore include uncertainty-aware prediction and explainability as supporting components of AGTR-Net. We use them to improve the interpretability and operational reliability of the proposed framework rather than treating predictive accuracy as the only measure of model quality.

2.9. Reinforcement Learning for Supply-Chain Recovery and Resilience Optimization

Predictive analytics can identify potential operational risks, but organizations also require strategies for responding to predicted disruptions. This creates a connection between predictive risk estimation and prescriptive decision-making.

Reinforcement learning provides a framework for sequential decision-making in which an agent learns to select actions based on observed states and the rewards associated with those actions. Schulman et al. (2017) introduced Proximal Policy Optimization, a policy-gradient algorithm designed to provide a practical balance between optimization stability and computational efficiency.

In supply-chain applications, reinforcement learning can be investigated for decisions involving transportation alternatives, inventory allocation, production scheduling, and recovery planning. However, meaningful implementation requires an environment that represents operational constraints, available actions, costs, and the consequences of decisions.

We propose using PPO as a simulation-based recovery optimization component. The predictive output from AGTR-Net would provide risk information to the simulated decision environment, while the recovery agent would select among feasible operational actions.

We recognize that historical shipment datasets do not necessarily contain the counterfactual information required to determine what would have happened if a different route or recovery action had been selected. Consequently, we use simulation-based evaluation and distinguish simulated recovery performance from verified real-world operational outcomes.

2.10. Public Datasets and Empirical Research in Logistics Prediction

The availability of open-source datasets is important for reproducible research in logistics and supply-chain analytics. Public datasets enable researchers to compare prediction algorithms, investigate operational patterns, and develop methods that can subsequently be validated using industry-specific information.

The Cargo 2000 Freight Tracking and Tracing dataset, available through the UCI Machine Learning Repository, provides anonymized airfreight process information

covering approximately five months of business execution. It contains 3,942 shipment-process instances, 7,932 transportation legs, and 56,082 activities, with 98 features (Metzger, 2015).The dataset is particularly relevant to temporal logistics analysis because it contains sequential operational activities and planned-versus-effective duration information. These characteristics provide an opportunity to investigate operational delay prediction and transportation-process dependencies.

However, the dataset also has limitations. It does not provide comprehensive calendar-dated disruption histories, verified multi-tier supplier relationships, complete inventory information, or detailed alternative-route costs. Therefore, the dataset is more appropriate for studying shipment-process delay prediction than for directly validating all forms of global supply-chain disruption prediction.We address this limitation by defining an observable operational-delay target and constructing a transportation-process graph from the relationships supported by the original records. We also clearly distinguish between empirical observations and assumptions introduced into the simulation-based recovery component.

2.11. Research Gaps and Positioning of the Proposed Framework

Our review of the literature identifies several research opportunities at the intersection of supply-chain resilience, graph learning, temporal deep learning, and operational decision support.

First, although traditional machine-learning and temporal deep-learning approaches have been investigated for supply-chain risk prediction, there remains a need to evaluate how transportation-network information can be integrated with sequential operational data. Many shipment-level prediction approaches focus on individual records and engineered variables, while network-based methods emphasize structural relationships. We investigate whether combining these perspectives improves operational-delay prediction.Second, although graph neural networks provide an effective mechanism for learning from interconnected data, their integration with temporal shipment-process representations requires careful architectural design. We address this opportunity through an adaptive graph-temporal fusion mechanism that combines network embeddings and temporal representations.

Third, limited labeled operational data creates challenges for training complex deep-learning models. We therefore investigate self-supervised representation learning as a supporting mechanism for extracting useful information from shipment sequences and transportation structures.Fourth, predictive accuracy alone may be insufficient for operational deployment. Logistics managers also require reliable probability estimates and explanations of the variables contributing to elevated risk. We therefore integrate uncertainty estimation and explainability into the proposed framework.

Fifth, predictive risk identification does not automatically produce an operational recovery strategy. We investigate a simulation-based reinforcement-learning component to explore how risk predictions may inform feasible recovery decisions.Based on these research opportunities, we propose AGTR-Net as an integrated graph-temporal framework for shipment-delay prediction and resilience-oriented decision support. We position the framework as an extension of existing graph-learning and temporal-learning methods rather than claiming that the individual algorithms are new.

The principal distinction of our research is the evaluation of adaptive graph-temporal representation learning, self-supervised pretraining, uncertainty-aware prediction, and simulation-based recovery optimization within a unified logistics-process research framework. We use the UCI Cargo 2000 dataset to establish an empirical foundation for shipment-delay prediction and identify the additional data requirements for extending the framework to supplier-level and industry-wide disruption analysis.

Methodology

Research Design

We develop an adaptive graph-temporal deep-learning framework to predict supply-chain disruptions and assess their potential operational consequences. The proposed framework integrates sequential shipment information, transportation-network relationships, operational performance indicators, and engineered risk variables to identify potential disruptions before the completion of the shipment process. Our methodological approach consists of data collection, data preprocessing, feature extraction, feature engineering, model development, and model evaluation.

We adopt a quantitative and experimental research design based on publicly available logistics data. The primary objective is to investigate whether combining graph-based learning with temporal deep-learning techniques can improve disruption prediction compared with conventional machine-learning and standalone deep-learning approaches. We further examine whether the proposed framework can identify critical transportation

nodes and operational activities associated with disruption propagation.

We organize the methodology around a proposed Adaptive Graph-Temporal Resilience Network (AGTR-Net). The framework combines a graph neural network for learning transportation-network dependencies, a temporal encoder for capturing sequential operational patterns, and an adaptive feature-fusion mechanism for integrating heterogeneous information. We also develop an uncertainty-aware prediction component to quantify the reliability of disruption predictions. Finally, we investigate a simulation-based recovery optimization module to evaluate potential resilience strategies.

Data Collection and Dataset Description

We obtain our primary dataset from the UCI Machine Learning Repository, which provides publicly accessible datasets for academic research and machine-learning experimentation. Specifically, we select the Cargo 2000 Freight Tracking and Tracing dataset because it contains actual, anonymized airfreight tracking and tracing information derived from logistics business processes. The dataset covers approximately five months of business execution and provides detailed information about shipment processes, transportation legs, planned activity durations, actual activity durations, and airport-related transportation information.

The dataset contains 3,942 overall shipment-process instances, 7,932 transportation legs, and 56,082 recorded activities. The repository provides 98 features describing the operational and transportation characteristics of the shipment processes. The dataset is multivariate and sequential, making it particularly relevant to our objective of modeling operational dependencies and identifying abnormal transportation performance. The dataset also contains missing values, which we address during preprocessing.

We select this dataset because its process-oriented structure allows us to examine how deviations in individual transportation activities may influence subsequent activities and the overall shipment-delivery process. The planned and effective durations provide an opportunity to identify operational delays, while the transportation-leg identifiers and anonymized airport identifiers enable us to construct a graph-based representation of the observed logistics network.

Table: Description of the primary research dataset

Dataset characteristics	Description
Dataset name	Cargo 2000 Freight Tracking and Tracing
Repository	UCI Machine Learning Repository
Dataset creator	Andreas Metzger
Year	2015
DOI	10.24432/C5559K
Data availability	Publicly available
License	Creative Commons Attribution 4.0 (CC BY 4.0)
Business domain	Airfreight transportation and logistics
Observation period	Approximately five months
Shipment-process instances	3,942
Transportation legs	7,932
Recorded activities	56,082
Number of features	98
Data characteristics	Multivariate and sequential
Data types	Primarily integer-valued operational and categorical identifiers

Major information	Planned activity durations, effective activity durations, transportation-leg identifiers, airport identifiers, and shipment-process structure
Research application	Shipment-delay prediction, process-risk identification, transportation-network analysis, and operational resilience assessment
Download source	https://archive.ics.uci.edu/dataset/382/cargo+2000

Source: UCI Machine Learning Repository, Cargo 2000 Freight Tracking and Tracing (Metzger, 2015).

We retrieve the dataset through the UCI repository and preserve its original identifiers and variable definitions during the initial data-loading stage. We use the original dataset documentation to distinguish planned durations from effective durations and to identify the operational meaning of individual activities. For reproducibility, we retain the original dataset version, document the retrieval date, and record all subsequent data-transformation procedures.

The dataset contains planned and actual activity durations for incoming and outgoing transportation legs. For example, variables associated with RCS represent freight check-in activities, DEP represents departure activities, RCF represents arrival activities, and DLV represents delivery activities. We use these operational milestones to reconstruct the shipment-process sequence and identify deviations from the planned logistics schedule.

We acknowledge that the dataset does not contain a complete supplier-to-supplier relationship database, explicit supplier-tier information, or comprehensive calendar-dated disruption histories. Consequently, we construct a transportation-process graph rather than assuming that the observed airport relationships represent a complete global supplier network. We define operational delay as a measurable proxy for disruption and interpret the resulting predictions within the logistics-process context.

Data Preprocessing

We perform data preprocessing to improve data quality, maintain consistency across transportation activities, and prepare the dataset for graph-temporal learning. Initially, we examine the dataset structure, variable types, missing-value distribution, duplicate records, and statistical characteristics of numerical variables. We also investigate the distribution of planned and effective durations to identify unusually large values, negative durations, and potentially invalid observations.

We first remove duplicate shipment-process records where duplication can be confirmed using the original process identifiers. We retain legitimate multiple transportation legs and repeated activities because these represent the actual structure of the shipment process rather than duplicate observations. We also verify the uniqueness and consistency of process identifiers and transportation-leg identifiers before constructing the network representation.

We then address missing values according to their operational meaning. For numerical variables, we use median imputation when the missingness can reasonably be treated as an incomplete measurement and when sufficient information is available to support imputation. For categorical or identifier variables, we preserve missingness as an explicit unknown category when appropriate. We do not automatically impute missing airport identifiers as though they represented known locations. Where an essential identifier is unavailable, we retain the observation for applicable non-network analyses but exclude it from graph relationships that cannot be reliably established.

We examine extreme values in planned and effective activity durations using distributional analysis and interquartile-range diagnostics. We investigate unusual observations before removing them because unusually long durations may represent genuine logistics disruptions. We correct demonstrably invalid values and retain plausible extreme delays to ensure that the model learns from operationally important events. When numerical transformations are necessary, we apply robust scaling or logarithmic transformations to reduce the influence of highly skewed distributions.

We subsequently reconstruct the shipment-process sequence by organizing operational activities according to their transportation-leg identifiers and process structure. We distinguish incoming transportation legs from outgoing transportation legs and preserve their relationships within each shipment process. This step is essential because the proposed temporal model must learn the progression of operational activities rather than treating each shipment as an unrelated static record.

We create a consistent representation of transportation activities using their planned durations, effective durations, activity types, and corresponding transportation identifiers. We ensure that the model receives only information available at the prediction point. In particular, we prevent the actual duration of a future activity from being used to predict whether that same activity will experience a delay. We use completed activities as historical information and reserve future activity outcomes for target construction.

We encode categorical variables using suitable categorical embeddings or integer mappings. We standardize numerical variables using parameters calculated exclusively from the training data. The same transformations are subsequently applied to validation and testing observations without recalculating normalization parameters. This procedure prevents information leakage between the training and evaluation stages.

Finally, we organize the processed dataset into shipment-level sequences, transportation-leg records, and graph-compatible node and edge structures. This representation provides a common input format for conventional machine-learning algorithms, temporal deep-learning models, and graph-based architectures.

Feature Extraction

We extract meaningful operational, temporal, and network-related representations from the preprocessed shipment records. Our objective is to capture not only individual shipment characteristics but also the relationships between transportation activities and the potential influence of earlier operational deviations on subsequent shipment performance.

We begin by extracting operational features from the original planned and effective activity-duration variables. Planned duration represents the expected time required to complete an activity, whereas effective duration represents the observed time required for its completion. We use these variables to characterize operational performance and identify deviations from expected logistics schedules.

We extract activity-level information associated with freight check-in, departure, arrival, and delivery processes. We also extract transportation-leg identifiers, anonymized airport identifiers, and the number of transportation hops. These variables provide information about the structure and complexity of individual shipment routes.

We subsequently extract sequential representations from the operational activity records. We arrange the activities in their corresponding process order and construct sequences that represent the progression of shipment operations. The temporal encoder receives these sequences to learn dependencies between preceding and subsequent activities. We use recurrent or attention-based mechanisms to identify operational patterns that may indicate increasing delay risk.

For graph-based feature extraction, we represent the transportation system as a graph:

[

G=(V,E,X)

]

where (V) represents transportation-related nodes, (E) represents observed transportation relationships, and (X) represents node and edge attributes.

We construct airport nodes using the anonymized airport identifiers available in the dataset. We represent observed transportation movements between airports as directed edges when the available route information supports such relationships. We also maintain shipment and transportation-leg identifiers to associate process activities with their corresponding network movements.

We extract graph features such as node connectivity, degree centrality, route frequency, and the number of observed transportation relationships associated with each airport. These features help us identify highly connected transportation nodes that may be associated with a greater number of shipment processes. We also calculate route-level characteristics to examine whether operational delays are concentrated in particular transportation relationships.

We apply graph neural networks to transform the original node and edge information into higher-dimensional representations. Through neighborhood aggregation, the graph encoder learns information from directly connected transportation nodes and their surrounding network structures. This enables the model to incorporate network dependencies into shipment-risk prediction.

We also investigate whether the combination of operational, sequential, and graph-derived representations provides additional predictive information compared with using individual feature groups independently. This allows us to evaluate the contribution of each information source within the proposed architecture.

Feature Engineering

Following feature extraction, we perform feature engineering to generate additional variables that describe

operational deviations, transportation complexity, and potential disruption patterns. We design these features to improve the ability of the model to recognize early signs of shipment delays and identify activities that may contribute to subsequent operational problems.

We calculate activity-level delay using the difference between effective duration and planned duration:

[

Delay_{i} = EffectiveDuration_{i}-PlannedDuration_{i}

]

A positive value indicates that an activity required more time than originally planned, whereas a negative value indicates that the activity was completed faster than expected.

We also calculate the relative delay ratio:

[

DelayRatio_{i} =

\frac{EffectiveDuration_{i}-PlannedDuration_{i}}

{\max(PlannedDuration_{i},\epsilon)}

]

where (\epsilon) represents a small positive constant used to prevent division by zero.

We construct a binary activity-delay indicator using a predefined tolerance threshold:

[

D_i =

\begin{cases}

1, & EffectiveDuration_i > PlannedDuration_i+\tau\

0, & \text{otherwise}

\end{cases}

]

where (\tau) represents a tolerance value determined through domain-informed sensitivity analysis. We select and document this threshold before final model evaluation. We also examine alternative thresholds to determine whether the prediction results remain stable under different definitions of operational delay.

We calculate cumulative delay across completed activities within each shipment process. This feature captures the accumulated operational deviation before the prediction point:

[

CumulativeDelay_t=\sum_{i=1}^{t}Delay_i

]

We additionally calculate the proportion of completed activities that experienced delays:

[

DelayFrequency_t=

\frac{\sum_{i=1}^{t}D_i}{t}

]

These variables allow us to distinguish isolated activity delays from repeated operational problems within the same shipment process.

We engineer route-complexity features using the number of transportation hops, the number of observed airport connections, and the number of transportation legs associated with a shipment. We also calculate graph-based indicators, including node degree, connectivity, and route frequency. These indicators provide additional information about the structural characteristics of transportation movements.

We construct process-stage indicators to identify whether a shipment is currently associated with freight check-in, departure, arrival, or delivery activities. We also develop interaction features between planned duration, accumulated delay, route complexity, and network connectivity. These interaction variables allow the model to identify whether transportation complexity or previous operational delays are associated with increased future delay risk.

To represent the evolving operational condition of each shipment, we create rolling and cumulative features from previously completed activities. These include recent average delay, maximum observed delay, variation in activity duration, and the difference between recent and historical operational performance. We calculate these features using only information available before the prediction point.

We construct the final target variable from future shipment-process outcomes. At a given prediction point, we define the target as whether a subsequent activity or the overall shipment process experiences a delay exceeding the predefined tolerance. We maintain separate targets for activity-level delay and shipment-level delay where the data structure supports reliable reconstruction.

Importantly, we do not use future effective durations, future delay indicators, or completed delivery outcomes as input features when predicting the same future event. We use these variables only for target construction and subsequent evaluation. This separation is necessary to

avoid target leakage and to ensure that the proposed model represents a realistic early-warning system.

Model Development

We develop the proposed Adaptive Graph-Temporal Resilience Network (AGTR-Net) to integrate transportation-network dependencies with sequential operational information. The architecture consists of a graph representation module, a temporal feature-learning module, an adaptive fusion mechanism, and multiple prediction heads. We design the model to capture both the structural relationships between transportation nodes and the sequential development of operational delays.

Transportation Graph Representation

We first construct a directed transportation graph from the observed shipment routes. Each airport identifier represents a transportation node, while each observed movement between two airport identifiers represents a directed relationship. We associate the relevant transportation-leg and shipment-process information with the corresponding graph edges.

We represent each node using its available network characteristics, including connectivity and observed route frequency. We represent each edge using available transportation-related features and operational information associated with the corresponding movement. We avoid introducing unsupported supplier relationships or external geographical characteristics that cannot be verified from the original dataset.

The resulting graph provides the structural input for the graph encoder. We use a graph neural network to learn relationships between connected transportation nodes and identify network structures associated with higher operational risk.

Graph Neural Network Encoder

We develop the graph encoder using a relational graph attention mechanism. The graph attention layer assigns different importance weights to neighboring nodes according to their contribution to the representation of the target node. This allows the model to learn that certain transportation relationships may be more informative than others when predicting operational delays.

The graph representation can be expressed as:

[

h_v^{(l+1)} =

\sigma\left(

\sum_{u\in \mathcal{N}(v)}

\alpha_{vu}^{(l)}W^{(l)}h_u^{(l)}

\right)

]

where (h_v^{(l)}) represents the embedding of node (v) at layer (l), (\mathcal{N}(v)) represents its neighboring nodes, (W^{(l)}) represents a trainable transformation matrix, (\alpha_{vu}^{(l)}) represents the attention weight between connected nodes, and (\sigma) represents a nonlinear activation function.

We use the resulting graph embeddings to represent transportation-network dependencies. Where the available process records support a reliable sequence of network states, we update graph representations as additional completed transportation information becomes available. Otherwise, we use the observed transportation graph as a static structural representation and update shipment-level operational features over process progression.

Temporal Feature-Learning Module

We develop the temporal encoder to learn sequential dependencies among shipment activities. We investigate Transformer-based temporal encoding as the principal approach because its attention mechanism can identify relationships between activities that are separated within the shipment process.

The temporal encoder receives the sequence of completed operational activities, including planned durations, observed effective durations, delay ratios, activity types, and process-stage indicators. We apply positional or process-order embeddings to preserve the sequence of operational activities.

We represent the temporal sequence as:

[

H_t=Transformer(X_1,X_2,\ldots,X_t)

]

where (X_t) represents the engineered operational feature vector at process step (t), and (H_t) represents the learned temporal representation.

We also evaluate a Temporal Convolutional Network and Long Short-Term Memory network as alternative temporal architectures. These comparisons allow us to determine whether attention-based sequence learning provides a measurable improvement over convolutional and recurrent approaches.

Adaptive Graph-Temporal Fusion

We combine the graph representation and temporal representation through an adaptive feature-fusion

mechanism. The purpose of this component is to allow the model to assign different importance to transportation-network information and operational sequence information depending on the characteristics of each shipment.

We calculate an adaptive fusion gate as:

[

g_t=\sigma(W_g[h_t^{G}\Vert h_t^{T}]+b_g)

]

where (h_t^{G}) represents the graph embedding, (h_t^{T}) represents the temporal embedding, (W_g) and (b_g) represent trainable parameters, and (\Vert) denotes vector concatenation.

We then obtain the fused representation:

[

z_t=g_t\odot h_t^{G}+(1-g_t)\odot h_t^{T}

]

where (\odot) represents element-wise multiplication. This mechanism allows the model to adaptively integrate graph-based and temporal information rather than assigning fixed importance to either representation.

We further investigate an uncertainty-aware fusion mechanism in which the model estimates the confidence associated with different feature representations. We use this component to reduce the influence of unreliable or incomplete information during prediction.

Multi-Task Prediction Module

We develop a multi-task prediction component to estimate different aspects of shipment-process risk. The primary prediction task is binary operational-delay classification. The model estimates the probability that a future activity or shipment process will experience a delay exceeding the predefined threshold.

The primary prediction function is:

[

\hat{P}(D_{t+h}=1)=

\sigma(W_dz_t+b_d)

]

where (D_{t+h}) represents the future delay outcome, (z_t) represents the fused graph-temporal representation, and (W_d) and (b_d) represent trainable parameters.

We additionally investigate regression-based prediction of future delay magnitude and classification of delay severity. These auxiliary tasks enable the model to distinguish between the probability of a delay and its potential operational magnitude.

Where the reconstructed activity sequences contain sufficient information, we also investigate time-to-delivery-delay prediction. We do not claim that this task represents conventional calendar-time survival prediction because the dataset covers a limited observation period and does not provide a complete calendar-dated history for all activities.

We jointly optimize the prediction tasks using a weighted loss function:

[

\mathcal{L}{total}=

\lambda_1\mathcal{L}{classification}

+\lambda_2\mathcal{L}{regression}

+\lambda_3\mathcal{L}{severity}

+\lambda_4\mathcal{L}_{regularization}

]

where the (\lambda) parameters control the contribution of the individual learning objectives. We determine these weights through validation experiments rather than assuming that all tasks contribute equally.

Self-Supervised Representation Learning

We investigate self-supervised learning to improve the representation of shipment sequences and transportation-network structures. This component is particularly relevant because the dataset contains a limited number of shipment-process instances compared with the number of parameters that may be required by a complex deep-learning architecture.

We apply masked-feature reconstruction to train the model to recover selected operational attributes from partially masked shipment sequences. We also investigate temporal contrastive learning, in which different representations of the same shipment process are encouraged to remain similar while representations of unrelated processes are distinguished.

For graph representation learning, we investigate graph reconstruction and link-prediction objectives using relationships supported by the observed transportation data. We use the resulting pretrained representations to initialize the graph-temporal model before supervised disruption prediction.

We evaluate whether self-supervised pretraining improves predictive performance compared with training the same architecture directly using the available supervised labels.

Explainability and Risk Interpretation

We incorporate an explainability component to identify the operational and network-related variables that contribute most strongly to individual predictions. We use feature-attribution techniques, such as Integrated Gradients and SHAP, to examine the influence of operational variables on predicted delay risk.

For graph-based predictions, we investigate graph explanation techniques to identify influential transportation nodes and relationships. We use these explanations to examine whether highly connected nodes, repeated route delays, or specific operational activities contribute to elevated shipment-risk predictions.

We interpret these explanations as model-based associations rather than definitive evidence of real-world causal relationships. This distinction is important because the dataset does not contain a controlled intervention design.

Resilience and Recovery Optimization

We extend the predictive framework with a simulation-based recovery optimization component. The purpose of this module is to investigate how early-warning predictions could support operational decisions when transportation delays are anticipated.

We formulate the recovery problem as a sequential decision-making task in which the simulated logistics environment provides an operational state, an available recovery action, and a corresponding reward. Potential actions include selecting an alternative observed route, prioritizing shipments, adjusting shipment sequencing, or allocating available operational resources where the required capacity information is available.

We investigate Proximal Policy Optimization as a candidate reinforcement-learning algorithm. We define the reward function to penalize accumulated delay, additional transportation cost, and unsuccessful recovery actions:

[

R_t=-(\alpha C_t+\beta D_t+\gamma F_t)

]

where (C_t) represents additional operational cost, (D_t) represents delay-related penalties, (F_t) represents recovery failure or infeasibility, and (\alpha), (\beta), and (\gamma) represent the corresponding weighting parameters.

We train the recovery policy within a simulated environment rather than directly applying reinforcement learning to static historical records. We calibrate the simulation using observable operational relationships from the dataset and clearly distinguish observed information from assumptions introduced by the simulator. Because the UCI dataset does not provide complete alternative-route capacity, inventory, or supplier-cost information, we restrict the recovery experiment to actions and constraints that can be reasonably parameterized. We treat any additional simulated costs or capacity values as experimental assumptions and perform sensitivity analysis to examine their influence on the results.

Model Evaluation

We evaluate the proposed AGTR-Net through a systematic experimental framework designed to assess predictive accuracy, generalization, calibration, network-risk identification, and recovery performance. We compare the proposed architecture against conventional machine-learning algorithms, standalone temporal deep-learning models, and graph-based models.

We first establish baseline performance using Random Forest, XGBoost, LightGBM, and CatBoost. These algorithms provide strong conventional machine-learning benchmarks for structured operational data. We subsequently evaluate Long Short-Term Memory, Gated Recurrent Unit, Temporal Convolutional Network, and Transformer-based models to assess their ability to learn sequential shipment patterns.

We also evaluate Graph Convolutional Networks, GraphSAGE, and Graph Attention Networks to determine the contribution of graph-based information. We compare these models with the proposed graph-temporal architecture to investigate whether integrating transportation-network dependencies and process sequences improves prediction performance.

We construct the training, validation, and testing partitions at the shipment-process level. We ensure that activities belonging to the same shipment-process instance do not appear across multiple partitions in a way that could cause information leakage. Where a reliable chronological ordering of prediction examples can be reconstructed, we additionally perform time-ordered evaluation. We use the training partition to estimate model parameters, the validation partition to tune hyperparameters, and the testing partition only for final performance assessment.

We evaluate binary delay prediction using several complementary classification metrics. Accuracy measures the proportion of correctly classified observations, while precision measures the proportion of predicted delay

events that are actual delay events. Recall measures the proportion of actual delay events successfully identified by the model. We calculate the F1-score to provide a combined assessment of precision and recall.

We place particular emphasis on the area under the Precision-Recall curve because operational delay prediction may involve unequal class distributions. We also calculate the area under the Receiver Operating Characteristic curve to evaluate the model's ability to distinguish between delayed and non-delayed observations across different classification thresholds.

We express the classification metrics as follows:

[
Precision=\frac{TP}{TP+FP}
]
[
Recall=\frac{TP}{TP+FN}
]
[
F1=2\times
\frac{Precision\times Recall}
{Precision+Recall}
]

where (TP), (FP), and (FN) represent true positives, false positives, and false negatives, respectively.

We evaluate probability calibration using the Brier score and Expected Calibration Error. These metrics help us determine whether predicted disruption probabilities correspond to observed event frequencies. This assessment is important because a resilience-oriented decision system should provide reliable risk estimates rather than only high classification accuracy.

Where we formulate delay magnitude as a regression task, we evaluate performance using Mean Absolute Error and Root Mean Squared Error. These metrics measure the difference between predicted and observed delay durations. We also evaluate the performance of the model across different delay thresholds to determine whether its predictions remain stable under alternative operational definitions of disruption.

We examine early-warning performance by evaluating whether the model can identify future delays using only completed activities available at the prediction point. We measure the number of correctly identified future delay events and the amount of process progression available before the delayed activity occurs. We do not report calendar-day warning lead times unless the required calendar timestamps are available and verified.

We evaluate the graph component by examining its ability to identify transportation nodes and relationships associated with delayed shipment processes. We investigate node-level risk-ranking performance and the consistency of influential-node identification across different experimental partitions. We also examine whether graph-based information improves prediction performance compared with models that use only shipment-level operational features.

We evaluate the recovery optimization module separately from the supervised prediction models. We compare the PPO-based policy against predefined baseline strategies, such as no intervention, simple delay-priority rules, and feasible alternative-route selection. We assess the resulting performance using simulated accumulated delay, additional operational cost, recovery success rate, and the frequency of infeasible actions. We report these outcomes as simulation-based findings rather than as verified improvements in actual logistics operations.

We conduct ablation experiments to determine the individual contribution of the proposed architecture components. We compare the complete AGTR-Net against versions without graph attention, without temporal encoding, without adaptive fusion, and without self-supervised pretraining. We also evaluate the effect of removing uncertainty estimation and individual feature groups.

We perform robustness experiments by introducing controlled levels of missing operational information, varying delay thresholds, and evaluating performance across different transportation-process structures. We examine whether the model remains stable when selected features are unavailable and whether its performance deteriorates under changes in the operational characteristics of the shipment processes.

Finally, we use repeated training runs with controlled random seeds to assess the stability of model performance. We report mean performance and standard deviation across repeated experiments and, where appropriate, calculate confidence intervals using shipment-level resampling. We select the final model based on validation performance and report its performance on the held-out test set without further parameter tuning.

Reproducibility and Experimental Implementation

We implement the proposed framework using Python and established open-source machine-learning libraries. We use Pandas or Polars for data processing, Scikit-learn for preprocessing and baseline models, PyTorch for deep-learning development, and PyTorch Geometric for graph representation learning. We use appropriate temporal-learning implementations for sequence modeling and Stable-Baselines3 for the PPO-based recovery experiment.

We document the model architecture, hyperparameters, feature definitions, delay-label thresholds, data-partitioning strategy, and random seeds to support reproducibility. We also retain the original UCI dataset citation and document all transformations applied to the source data.

We do not assume that the proposed model will outperform every baseline. Instead, we evaluate whether the combination of graph learning, temporal representation learning, and adaptive feature fusion provides statistically and operationally meaningful improvements. We consider the proposed framework successful only if the experimental results demonstrate improvements under leakage-controlled testing and if the model's predictions remain interpretable and sufficiently calibrated.

Methodological Scope and Limitations

We recognize that the selected dataset primarily represents airfreight shipment processes over a relatively short observation period. Therefore, our empirical findings are most directly applicable to operational delay prediction and transportation-process resilience rather than to every form of global supply-chain disruption.

We also acknowledge that the anonymized airport identifiers restrict our ability to connect the observed transportation network to verified geographical locations. Similarly, the dataset does not provide comprehensive supplier identities, inventory levels, procurement costs, or multi-tier supplier relationships. We therefore avoid interpreting the constructed transportation graph as a complete supplier network.

We address these limitations by defining disruption through an explicit operational-delay proxy, preserving the distinction between observed and simulated information, and evaluating the model within the scope of the available data. We consider the proposed framework a methodological foundation that can subsequently be extended using additional open-source supplier-network, trade-flow, or external disruption-event datasets.

Overall, our methodology provides a reproducible approach to integrating graph neural networks, temporal deep learning, adaptive feature fusion, and simulation-based recovery optimization for logistics disruption prediction. Through comparative evaluation and component-level experimentation, we aim to establish whether the proposed AGTR-Net can provide more informative and reliable early-warning predictions than conventional machine-learning and standalone deep-learning approaches.

4. Results and Discussion

4.1. Dataset Characteristics and Descriptive Analysis

We evaluate the proposed Adaptive Graph-Temporal Resilience Network (AGTR-Net) using the Cargo 2000 Freight Tracking and Tracing dataset obtained from the UCI Machine Learning Repository. The original dataset contains 3,942 shipment-process instances, 7,932 transportation legs, and 56,082 recorded activities, with 98 original features. We use the shipment-process structure to investigate operational delays and transportation-related dependencies.

Following data preprocessing, we construct shipment-level and activity-level observations and derive operational delay indicators using planned and effective activity durations. For the purpose of this illustrative experiment, we assume that 3,816 shipment-process observations remain after data-quality screening and that 3,412 observations contain sufficient information for the primary supervised prediction task. We use a hypothetical delay-label threshold of 30 minutes to distinguish delayed and non-delayed operational activities, assuming that the duration units have been verified as minutes in the source documentation.

Table 1. Illustrative dataset preprocessing and distribution results

Dataset characteristics	Sample result
Original shipment-process instances	3,942
Original transportation legs	7,932
Original recorded activities	56,082

Original features	98
Shipment processes retained after preprocessing	3,816
Shipment processes excluded	126
Valid supervised observations	3,412
Delayed observations	1,128
Non-delayed observations	2,284
Delayed observations (%)	33.06%
Non-delayed observations (%)	66.94%
Engineered operational and network features	34

We observe in this illustrative distribution that approximately one-third of the supervised observations are associated with operational delays. This class distribution demonstrates why accuracy alone may not provide a sufficient assessment of predictive performance. We therefore emphasize recall, F1-score, and Precision-Recall AUC when comparing the proposed model with conventional machine-learning and deep-learning approaches.

4.2. Comparative Performance of Machine-Learning and Deep-Learning Models

We compare the proposed AGTR-Net against conventional machine-learning algorithms, temporal deep-learning models, and graph-based neural networks. We use Random Forest, XGBoost, LightGBM, and CatBoost as structured-data benchmarks. We further evaluate LSTM, GRU, Temporal Convolutional Network (TCN), and Transformer architectures to examine their ability to capture sequential shipment-process dependencies.

We also evaluate Graph Convolutional Network (GCN), GraphSAGE, and Graph Attention Network (GAT) architectures to investigate the contribution of transportation-network information. Finally, we assess AGTR-Net, which integrates graph representation learning, temporal attention, and adaptive feature fusion.

Table 2. comparative performance of predictive models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)	ROC-AUC	PR-AUC
Random Forest	82.4	76.8	70.2	73.3	0.861	0.774
XGBoost	85.1	80.5	76.4	78.4	0.891	0.817
LightGBM	84.7	79.8	75.9	77.8	0.886	0.809
CatBoost	84.3	79.1	75.1	77.1	0.882	0.803
LSTM	84.8	80.2	76.9	78.5	0.889	0.815
GRU	84.2	79.4	75.8	77.6	0.883	0.807
TCN	85.5	81.3	78.1	79.7	0.898	0.826
Transformer	86.2	82.1	79.5	80.8	0.907	0.839
GCN	85.0	80.6	77.3	78.9	0.893	0.820
GraphSAGE	86.1	82.0	79.2	80.6	0.905	0.836
GAT	87.0	83.1	80.7	81.9	0.914	0.851
AGTR-Net	**90.1**	**87.2**	**85.6**	**86.4**	**0.945**	**0.901**

The results indicate that AGTR-Net achieves the highest predictive performance among the evaluated architectures. It produces an accuracy of 90.1%, precision of 87.2%, recall of 85.6%, and F1-score of 86.4%. Its

ROC-AUC and PR-AUC values are 0.945 and 0.901, respectively.

We observe that the Transformer model achieves an F1-score of 80.8%, demonstrating the potential contribution of temporal attention to shipment-delay prediction. The GAT model achieves an F1-score of 81.9%, suggesting that transportation-network information may provide additional predictive value. However, the proposed AGTR-Net achieves a higher illustrative F1-score than both individual architectures.

Compared with XGBoost, AGTR-Net improves the illustrative F1-score from 78.4% to 86.4%, representing an absolute improvement of 8.0 percentage points. Its illustrative PR-AUC increases from 0.817 to 0.901. These differences suggest that integrating temporal and graph representations may improve the identification of delayed shipment processes in this hypothetical experiment.We emphasize that these differences are descriptive comparisons of sample values, not statistically verified improvements. Actual statistical significance must be assessed using repeated experiments and appropriate paired testing on the same held-out observations.

4.3. Analysis of Classification Performance

We further investigate the predictive performance of the proposed architecture by examining its ability to identify delayed and non-delayed shipment observations. We use precision, recall, and F1-score to evaluate the model's performance across both classes.

The illustrative results indicate that AGTR-Net provides a more balanced relationship between precision and recall than the conventional machine-learning benchmarks. This is important because a delay-prediction model that identifies only a small proportion of delayed shipments may fail to provide sufficient early-warning value, even if its overall accuracy appears high.

Table 3. Illustrative confusion matrix for AGTR-Net

Actual class	Predicted delayed	Predicted non-delayed	Total
Delayed	193	33	226
Non-delayed	31	426	457
Total	224	459	683

Based on this illustrative confusion matrix, AGTR-Net correctly identifies 193 delayed observations out of 226 actual delayed observations. The corresponding recall is approximately 85.4%. The model also correctly classifies 426 non-delayed observations, resulting in a specificity of approximately 93.2%.

We observe that the model generates 31 false-positive predictions and 33 false-negative predictions in this hypothetical test sample. False positives may lead to unnecessary operational interventions, whereas false negatives may result in delayed shipments remaining unidentified. We therefore consider the balance between these errors essential when selecting a practical operating threshold.

The model should not be deployed using a fixed probability threshold without considering the relative costs of missed disruptions and unnecessary alerts. We would determine the final threshold using validation data and operational requirements rather than optimizing it on the test set.

4.4. Ablation Study of the Proposed AGTR-Net Architecture

We conduct an illustrative ablation analysis to examine the contribution of the graph encoder, temporal encoder, adaptive fusion mechanism, and self-supervised learning component.

We compare the complete architecture with alternative configurations in which individual components are removed. This experiment is designed to establish whether the proposed model's improvement is associated with the integration of multiple information sources rather than the complexity of the architecture alone.

Table 4. results of AGTR-Net

Model configuration	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)	PR-AUC
Temporal-only model	84.8	80.2	76.9	78.5	0.815
Graph-only model	85.0	80.6	77.3	78.9	0.820
Graph + temporal model	87.1	83.2	81.0	82.1	0.858
Graph + temporal + adaptive fusion	88.4	85.0	83.0	84.0	0.879
Graph + temporal + self-supervised learning	89.0	85.8	83.8	84.8	0.887
Complete AGTR-Net	**90.1**	**87.2**	**85.6**	**86.4**	**0.901**

The illustrative ablation results demonstrate that the graph-temporal combination performs better than either the temporal-only or graph-only configuration. The graph-temporal architecture achieves an F1-score of 82.1%, compared with 78.5% for the temporal-only model and 78.9% for the graph-only model.

The addition of adaptive fusion increases the illustrative F1-score to 84.0%. We interpret this as evidence that adaptive integration may be more effective than treating graph and temporal representations as independent information sources.

The complete AGTR-Net achieves the highest illustrative F1-score of 86.4%. The difference between the complete model and the graph-temporal architecture without additional components is 4.3 percentage points. These results suggest that adaptive fusion and self-supervised representation learning may provide complementary contributions to operational-delay prediction.

4.5. Feature Importance and Explainability Analysis

We investigate the influence of operational and network-related variables on the predicted probability of shipment delay. We use SHAP-based feature attribution for structured inputs and graph-explanation methods to examine the contribution of transportation-network relationships.

Table 5. feature-importance ranking

Rank	Feature	Relative importance (%)	Interpretation
1	Cumulative operational delay	18.6	Captures accumulated process deviation
2	Recent activity delay	15.2	Identifies recent operational deterioration
3	Planned activity duration	12.4	Represents expected operational time
4	Transportation-leg complexity	10.8	Captures route complexity
5	Previous activity performance	9.7	Represents historical process performance
6	Transportation-node connectivity	8.9	Describes network structure

7	Delay frequency	7.6	Captures repeated operational deviations
8	Number of transportation hops	6.8	Represents transportation complexity
9	Activity type	5.4	Identifies the operational stage
10	Other engineered features	4.6	Captures remaining information

Note: Hypothetical normalized importance values for demonstration; they are not actual SHAP results.

The illustrative feature-importance results indicate that cumulative operational delay and recent activity delay are the most influential variables. Together, these features account for 33.8% of the illustrative normalized importance. This suggests that the previous operational condition of a shipment may provide important information when estimating its future delay risk.

We also observe that transportation-leg complexity, transportation-node connectivity, and the number of transportation hops contribute to the illustrative prediction results. These features indicate the potential importance of incorporating network structure into shipment-delay modeling.

We interpret these feature contributions as predictive associations. We do not conclude that a highly connected transportation node necessarily causes shipment delays, because additional operational and external factors may influence the observed outcomes.

4.6. Model Calibration and Robustness Analysis

We evaluate the reliability of predicted probabilities using the Brier score and Expected Calibration Error. These measures are particularly relevant when predicted probabilities are used to prioritize shipments or initiate operational interventions.

Table 6. Illustrative model calibration results

Model	Brier score	Expected Calibration Error
XGBoost	0.142	0.061
LSTM	0.151	0.068
Transformer	0.132	0.047
GAT	0.128	0.043
AGTR-Net	0.098	0.026

Note: Illustrative hypothetical results. Lower values indicate better calibration.

The illustrative results indicate that AGTR-Net achieves a lower Brier score and Expected Calibration Error than the comparison models. This suggests that the proposed architecture may provide more reliable probability estimates in addition to improved classification performance.

We also conduct hypothetical robustness experiments by introducing controlled missing-feature conditions. The purpose is to evaluate how predictive performance may change when shipment information is incomplete.

Table 7. robustness analysis under missing operational information

Missing-feature proportion	AGTR-Net accuracy (%)	AGTR-Net F1-score (%)
0%	90.1	86.4
5%	89.2	85.3
10%	87.9	83.8
15%	86.2	81.9
20%	83.8	79.1

Note: Illustrative hypothetical robustness results.

The illustrative results indicate a gradual reduction in predictive performance as the proportion of missing information increases. This emphasizes the importance of reliable operational data integration for practical deployment. We also consider uncertainty estimation and missing-data indicators important components of a production-oriented prediction system.

4.7. Recovery Optimization Results

We investigate the potential operational value of integrating shipment-delay predictions with a simulation-based recovery optimization module. We use Proximal Policy Optimization (PPO) as the candidate reinforcement-learning algorithm and compare it with predefined operational strategies.

We formulate the simulation environment to represent shipment states, delay-risk predictions, and feasible recovery actions. We evaluate the approaches using simulated accumulated delay, additional operational cost, and recovery success rate.

Table 8. Illustrative recovery optimization comparison

Recovery strategy	Average simulated delay (hours)	Additional cost per shipment (USD)	Recovery success rate (%)
No intervention	18.4	0.00	0.0
Delay-priority rule	15.9	8.50	38.2
Alternative-route rule	13.7	15.20	54.6
PPO-based recovery policy	10.8	18.70	71.4

The illustrative results suggest that the PPO-based recovery policy achieves the lowest simulated average delay of 10.8 hours and the highest recovery success rate of 71.4%. Compared with the no-intervention strategy, the simulated average delay decreases by approximately 41.3%.

However, the PPO strategy also produces the highest additional cost per shipment among the evaluated approaches. This indicates a trade-off between operational delay reduction and recovery expenditure. We therefore consider cost-sensitive optimization essential when selecting recovery actions.We emphasize that these results depend on the assumptions used to construct the simulated environment. The Cargo 2000 dataset does not provide complete alternative-route capacity, contractual costs, or inventory information. Consequently, the simulated results cannot be interpreted as proof of actual savings in commercial freight operations.

4.8. Overall Comparative Discussion

The illustrative experimental results indicate that AGTR-Net has the potential to improve operational-delay prediction by combining transportation-network representations with sequential shipment-process information.

The comparison between XGBoost, Transformer, GAT, and AGTR-Net highlights the potential benefits of combining structured-data learning, temporal attention, and graph representation learning. XGBoost provides a strong conventional benchmark, while Transformer and GAT demonstrate the potential value of temporal and network information independently.The proposed AGTR-Net produces the strongest illustrative performance across

accuracy, precision, recall, F1-score, ROC-AUC, and PR-AUC. Its ablation results further suggest that adaptive fusion and self-supervised representation learning may contribute to improved predictive performance.

We also observe that the proposed model's potential value extends beyond binary delay classification. Probability calibration, graph explainability, and simulation-based recovery optimization provide additional information that could support operational decision-making.Nevertheless, we recognize that a more complex model does not automatically provide greater practical value. The model must demonstrate stable performance on independent shipment processes, provide useful warning time, and operate within the computational and data-quality constraints of real logistics systems.

5. Implementation of the Best-Performing Model in U.S. Industries

5.1. Selection of the Best Model for Industrial Implementation

Based on the illustrative results, we select AGTR-Net as the preferred predictive architecture because it achieves the highest F1-score of 86.4%, ROC-AUC of 0.945, and PR-AUC of 0.901 among the evaluated models. Its illustrative performance also indicates stronger calibration than the comparison models.

We do not recommend deploying the complete architecture immediately across all industrial operations. Instead, we propose beginning with a controlled pilot in the U.S. air-cargo and freight-forwarding sector, where shipment tracking, planned transportation activities, and actual operational events can be monitored systematically.

We consider this sector an appropriate starting point because the original UCI dataset represents airfreight shipment processes. Once the framework demonstrates reliable performance using real operational data, we can investigate extensions to manufacturing, retail distribution, and other logistics-intensive industries.

5.2. Application in the U.S. Air-Cargo Industry

We propose that U.S. freight-forwarding companies integrate AGTR-Net with their existing transportation-management systems and shipment-tracking platforms. The system would receive operational events such as shipment check-in, departure, arrival, transfer, and delivery.

The model would process completed operational events and estimate the probability that the next shipment activity or delivery process will experience a delay. When the predicted probability exceeds an operationally defined threshold, the system would generate an early-warning notification.

For example, a freight-forwarding company could identify a shipment with a predicted 85% probability of a delay at its next transportation milestone. The system could examine the shipment's previous operational performance, route complexity, and observed transportation relationships to explain the elevated risk.

A logistics planner could then investigate the shipment, communicate with the transportation provider, prioritize the shipment, or consider a feasible alternative route. The final operational decision would remain under human supervision.

5.3. Application in U.S. Manufacturing and Distribution

We extend the framework to U.S. manufacturing by incorporating additional information from enterprise-resource-planning systems, supplier-management systems, warehouse-management systems, and transportation-management platforms.The extended graph would represent suppliers, manufacturing facilities, distribution centers, transportation providers, and observed material movements. The temporal model would process supplier delivery records, inventory movements, production schedules, and shipment events.

For example, if an inbound shipment containing a critical production component is predicted to arrive late, the system could estimate the potential operational consequences using available inventory and production information. It could then provide an early warning to procurement and production-planning teams.In retail distribution, the same approach could help identify inbound shipments that may affect replenishment schedules and product availability. The model could support decisions about inventory prioritization, distribution-center allocation, and alternative transportation arrangements. We recognize that these applications require additional data and sector-specific validation. We therefore treat them as proposed extensions rather than outcomes demonstrated by the primary airfreight experiment.

5.4. Proposed Industrial Implementation Performance Indicators

We propose evaluating the real-world deployment using operational measures rather than relying exclusively on predictive classification metrics.

Table 9. Proposed U.S. industrial deployment targets

Performance indicator	Illustrative pilot target	Operational objective
Delay prediction F1-score	At least 0.85	Maintain balanced prediction quality
PR-AUC	At least 0.85	Improve delayed-shipment identification
Prediction calibration	ECE below 0.05	Improve probability reliability
Early-warning lead time	At least 2 hours where data permits	Provide time for intervention
False-alert rate	Below 15% of generated alerts	Limit unnecessary interventions
Shipment on-time performance	Improve by 5% relative to baseline	Improve delivery reliability
Average operational delay	Reduce by 10% relative to baseline	Improve logistics efficiency
Model availability	At least 99% during pilot operations	Maintain system availability

We would establish these targets before deployment and evaluate them against a documented baseline. We would also compare operational performance across equivalent shipment categories and account for changes in shipment volume and external operating conditions.

5.5. Implementation Architecture and Deployment Strategy

We propose deploying AGTR-Net through a cloud-based or enterprise-hosted architecture. Operational data would be collected from company information systems and transformed into standardized shipment-event records.The data-processing layer would validate identifiers, timestamps, planned durations, and actual activity information. The graph-construction layer would update observed transportation relationships, while the temporal-processing layer would maintain the sequence of completed shipment act

6. Conclusion

In this study, we proposed an Adaptive Graph-Temporal Resilience Network (AGTR-Net) as an integrated deep-learning framework for predicting operational shipment delays and supporting supply-chain resilience. We developed the framework around the premise that shipment delays cannot always be understood through isolated operational events or historical performance indicators. Instead, logistics activities are interconnected through transportation dependencies and evolve over time. By combining graph-based representation learning with temporal sequence modeling, our proposed approach aims to capture both the structural relationships among transportation activities and the temporal patterns associated with changing shipment conditions. Through adaptive feature fusion, self-supervised representation learning, explainable prediction, and uncertainty-aware assessment, AGTR-Net provides a comprehensive architecture for transforming operational logistics data into predictive information that can support proactive decision-making.

We positioned the proposed framework within the broader development of artificial intelligence applications for supply-chain risk management and resilience. Conventional machine-learning methods can identify patterns in structured operational data, while temporal deep-learning models can capture sequential dependencies and graph neural networks can represent relationships among interconnected transportation activities. However, considering these dimensions independently may limit the ability to recognize complex operational dependencies. Our proposed architecture addresses this methodological challenge by integrating graph and temporal representations within a unified learning framework. Furthermore, its adaptive fusion mechanism is designed to adjust the contribution of different representations according to the available operational information, while self-supervised learning offers a potential strategy for extracting useful patterns

from large volumes of operational records where labeled delay information may be limited.

For empirical evaluation, we selected the Cargo 2000 Freight Tracking and Tracing dataset from the UCI Machine Learning Repository. Its shipment-process records, transportation legs, and activity-level information provide a suitable foundation for investigating operational delay prediction and temporal logistics patterns. The dataset also offers an opportunity to construct transportation-activity graphs from observed relationships rather than relying on assumed supplier or network connections. Nevertheless, we recognize that the dataset represents a specific airfreight operating environment and does not provide comprehensive information about supplier-tier dependencies, geopolitical disruptions, natural disasters, inventory constraints, or alternative-route recovery costs. Consequently, the empirical scope of this study is centered on shipment-delay prediction and operational transportation dependencies. Broader claims concerning global supply-chain disruption propagation and supplier-level resilience require additional datasets and independent validation.

Beyond prediction, we conceptualized a simulation-based recovery optimization component to investigate how predictive risk information could support operational intervention decisions. This component provides a pathway for evaluating alternative recovery strategies, including delay-priority rules, alternative-route selection, and reinforcement-learning-based policies. By separating historical prediction from simulated intervention, we acknowledge that observational shipment records alone cannot establish the causal effects of recovery actions. Accordingly, the proposed optimization component should be interpreted as a decision-support extension whose effectiveness depends on the realism of the simulation environment, the validity of its operational assumptions, and subsequent validation using real-world intervention data.

The proposed framework has several potential implications for logistics operators and supply-chain managers. More reliable advance identification of shipment delays could help organizations prioritize shipments requiring attention, allocate operational resources more effectively, coordinate transportation activities, and evaluate recovery alternatives before delays escalate. Explainability and uncertainty estimation may further improve the practical usefulness of predictive systems by helping decision-makers understand the factors contributing to a risk assessment and distinguish confident predictions from uncertain cases. In this respect, we view AGTR-Net not as a replacement for human operational judgment, but as a decision-support architecture that can complement existing transportation management and supply-chain monitoring systems.

Several limitations provide important directions for future research. First, the use of a single historical airfreight dataset restricts the generalizability of the proposed framework across transportation modes, industries, and geographical regions. Future studies should evaluate the approach using larger and more diverse datasets covering road, rail, maritime, and multimodal transportation networks. Second, additional information concerning weather conditions, port and airport congestion, supplier performance, inventory availability, geopolitical events, and transportation costs could improve the representation of external disruption factors. Third, future research should examine the framework under distribution shifts, changing network structures, incomplete observations, and previously unseen disruption scenarios. Fourth, the simulation-based recovery component should be validated through realistic operational constraints and, where possible, controlled pilots with logistics organizations. Finally, further investigation into computational efficiency, model interpretability, probability calibration, and deployment scalability will be necessary before the framework can be adopted in time-sensitive industrial environments.

In conclusion, we establish AGTR-Net as a proposed research framework for integrating graph-based dependency modeling, temporal deep learning, adaptive representation fusion, and uncertainty-aware decision support within operational shipment-delay prediction. Its central contribution lies in treating logistics delay prediction as a problem involving both interconnected transportation relationships and evolving operational sequences. By connecting predictive analytics with a simulation-based pathway for recovery planning, this study provides a foundation for developing more proactive, data-driven, and resilient logistics operations. Future empirical validation across diverse transportation networks will be essential to determine the framework's predictive advantages, operational value, and applicability to broader supply-chain resilience challenges.

References

1. Mia, M. M., Roy, M. K., YASSAR, I. S., Mottalib, M. Y., Yezdani, S., Nijhum, A. M., ... & Uddin, M. K. (2025). Integrating Blockchain Security and Machine Learning for Fraud Detection in the US Banking System. Emerging Frontiers Library for The

American Journal of Engineering and Technology, 7(11), 65-76.

2. An, X., Qin, K., Xing, Z., & Quan, P. (2026). Review of supply chain risk propagation mechanisms: A graph neural network perspective. Procedia Computer Science, 287, 1048–1055. https://doi.org/10.1016/j.procs.2026.08.146
3. Baryannis, G., Validi, S., Dani, S., & Antoniou, G. (2019). Supply chain risk management and artificial intelligence: State of the art and future research directions. International Journal of Production Research, 57(7), 2179–2202. https://doi.org/10.1080/00207543.2018.1530476
4. Christopher, M., & Peck, H. (2004). Building the resilient supply chain. The International Journal of Logistics Management, 15(2), 1–14. https://doi.org/10.1108/09574090410700275
5. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735
6. Hosseini, S., Ivanov, D., & Dolgui, A. (2019). Review of quantitative methods for supply chain resilience analysis. Transportation Research Part E: Logistics and Transportation Review, 125, 285–307. https://doi.org/10.1016/j.tre.2019.03.001
7. Hosseini, S., Ivanov, D., & Dolgui, A. (2020). Ripple effect modelling of supplier disruption: Integrated Markov chain and dynamic Bayesian network approach. International Journal of Production Research, 58(11), 3284–3303. https://doi.org/10.1080/00207543.2019.1661538
8. Ivanov, D. (2020). Predicting the impacts of epidemic outbreaks on global supply chains: A simulation-based analysis on the coronavirus outbreak (COVID-19/SARS-CoV-2) case. Transportation Research Part E: Logistics and Transportation Review, 136, Article 101922. https://doi.org/10.1016/j.tre.2020.101922
9. Ivanov, D. (2022). Viable supply chain model: Integrating agility, resilience and sustainability perspectives—Lessons from and thinking beyond the COVID-19 pandemic. Annals of Operations Research, 319, 1411–1431. https://doi.org/10.1007/s10479-020-03640-6
10. Ivanov, D., & Dolgui, A. (2021). OR-methods for coping with the ripple effect in supply chains during COVID-19 pandemic: Managerial insights and research implications. International Journal of Production Economics, 232, Article 107921. https://doi.org/10.1016/j.ijpe.2020.107921
11. Kayalvizhi, S., Saral Jeeva Jothi, D., Devi, T., & Brijet, Z. (2025). Graph neural network-based predictive modeling for enhanced supply chain resilience against multi-modal disruptions. Journal of Information Systems Engineering and Management, 10(33s), 628–642. https://doi.org/10.52783/jisem.v10i33s.5571
12. Ma, C., Zhang, L., You, L., & Tian, W. (2025). A review of supply chain resilience: A network modeling perspective. Applied Sciences, 15(1), Article 265. https://doi.org/10.3390/app15010265
13. Metzger, A. (2015). Cargo 2000 freight tracking and tracing [Data set]. UCI Machine Learning Repository. https://doi.org/10.24432/C5559K
14. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. arXiv. https://doi.org/10.48550/arXiv.1707.06347
15. Tang, C. S. (2006). Perspectives in supply chain risk management. International Journal of Production Economics, 103(2), 451–488. https://doi.org/10.1016/j.ijpe.2005.12.006
16. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In Advances in Neural Information Processing Systems (Vol. 30, pp. 5998–6008).
17. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph attention networks. In International Conference on Learning Representations.
18. Zhao, T., Chen, G., & Busababodhin, P. (2026). An intelligent expert system for logistics disruption prediction and mitigation in global supply chains: Integrating graph-based risk inference with ensemble forecasting. Scientific Reports. https://doi.org/10.1038/s41598-026-61617-0
19. Zogaan, W. A., Ajabnoor, N., & Salamai, A. A. (2025). Leveraging deep learning for risk prediction and resilience in supply chains: Insights from critical industries. Journal of Big Data, 12, Article 94. https://doi.org/10.1186/s40537-025-01143-4
20. YASSAR, I. S. (2026). TRACEABLE AND EXPLAINABLE AI LINEAGE ARCHITECTURES FOR REGULATORY-COMPLIANT DECISION INTELLIGENCE SYSTEMS. Computers and Education Letters, 3(01), 1-18.
21. Khatun, P., Umam, S., Razzak, R. B., Shamsuddin, I. B., & Salma, N. (2025). A study on the effectiveness

of machine learning models for hepatitis prediction. Scientific Reports, 15(1), 30659.

22. Umam, S., Razzak, R. B., Munni, M. Y., & Rahman, A. (2025). Exploring the non-linear association of daily cigarette consumption behavior and food security-An application of CMP GAM regression. PLoS One, 20(7), e0328109.

23. Razzak, R. B., & Umam, S. (2025, November). The Psychological Burden Behind the Urgent Care Visit: A Mediation Analysis of Smoking, Anxiety, and Healthcare Utilization Using National Health Interview Survey (NHIS), 2023. In APHA 2025 Annual Meeting and Expo. APHA.

24. Nguyen, A. T. P., Shak, M. S., & Al-Imran, M. (2024). Advancing early skin cancer detection: A comparative analysis of machine learning algorithms for melanoma diagnosis using dermoscopic images. International Journal of Medical Science and Public Health Research, 5(12), 119-133.

25. Mahmud, F., Das, A. C., Shak, M. S., Rahman, N., Eva, A. A., Mridha, M. F., & Hossen, M. J. (2025). HybridTabNet-QC: A transformer-based clinical feature fusion framework for heart disease risk prediction. IEEE Open Journal of the Computer Society, 7, 1-13.

26. Das, A. C., Shak, M. S., Rahman, N., Mahmud, F., Shoaib, H. A., & Hossen, M. J. (2026). Cardiovascular disease prediction using variational recurrent autoencoders with uncertainty estimation. Scientific Reports.

27. Das, A. C., Shak, M. S., Rahman, N., Mahmud, F., Shoaib, H. A., & Hossen, M. J. (2026). Cardiovascular disease prediction using variational recurrent autoencoders with uncertainty estimation. Scientific Reports.

28. Das, A. C., Shak, M. S., Rahman, N., Mahmud, F., Eva, A. A., & Hasan, M. N. (2025, May). Self-Supervised Contrastive Learning for Disease Trajectory Prediction. In 2025 5th International Conference on Pervasive Computing and Social Networking (ICPCSN) (pp. 732-738). IEEE.

29. Mozumder, M. A. S., Mahmud, F., Shak, M. S., Sultana, N., Rodrigues, G. N., Al Rafi, M., ... & Bhuiyan, M. S. M. (2024). Optimizing customer segmentation in the banking sector: a comparative analysis of machine learning algorithms. Journal of Computer Science and Technology Studies, 6(4), 01-07.

30. Rahman, M. H., Das, A. C., Shak, M. S., Uddin, M. K., Alam, M. I., Anjum, N., ... & Alam, M. (2024). Transforming customer retention in fintech industry through predictive analytics and machine learning. The American Journal of Engineering and Technology, 6(10), 150-163.

31. Bhuiyan, R. J., Akter, S., Uddin, A., Shak, M. S., Islam, M. R., Rishad, S. S. I., ... & Hasan-Or-Rashid, M. (2024). Sentiment analysis of customer feedback in the banking sector: A comparative study of machine learning models. The American Journal of Engineering and Technology, 6(10), 54-66.

32. Razzak, R. B., & Umam, S. (2025, November). Health Equity in Action: Utilizing PRECEDE-PROCEED Model to Address Gun Violence and associated PTSD in Shaw Community, Saint Louis, Missouri. In APHA 2025 Annual Meeting and Expo. APHA.

33. Khatun, P., Umam, S., Razzak, R. B., Shamsuddin, I. B., & Salma, N. (2025). A study on the effectiveness of machine learning models for hepatitis prediction. Scientific reports, 15(1), 30659.

34. Umam, S., Razzak, R. B., Munni, M. Y., & Rahman, A. (2025). Exploring the non-linear association of daily cigarette consumption behavior and food security-An application of CMP GAM regression. Plos one, 20(7), e0328109.

35. Adams, R., Grellner, S., Umam, S., & Shacham, E. (2023, November). Using google searching to identify where sexually transmitted infections services are needed. In APHA 2023 Annual Meeting and Expo. APHA.

36. Umam, S., Adams, R., Shacham, E., & Charles, D. L. (2024, October). Predictors of weapon carrying for high school students. In APHA 2024 Annual Meeting and Expo. APHA.

37. YASSAR, I. S. (2023). SCALABLE SDN-BASED ARCHITECTURE FOR LARGE-SCALE ENTERPRISE NETWORK MANAGEMENT. Insights Sustainable Engineering Practices, 1(01), 115-130.

38. Sayed, M. A., Badruddowza, M. S. U. S., Al Mamun, A., Nabi, N., Mahmud, F., Alam, M. K., ... & Choudhury, M. Z. M. E. (2024). Comparative analysis of machine learning algorithms for predicting cybersecurity attack success: A performance evaluation. The American Journal of Engineering and Technology, 6(09), 81-91.

39. MAMUN, A., Nath, A., Dey, S. K., Nath, P., RAHMAN, M., SHORNA, J., & Anjum, N. (2025). Real-time malware detection in cloud infrastructures using convolutional neural networks: a deep learning

framework for enhanced cybersecurity. INTERNATIONAL JOURNAL OF COMPUTER SCIENCE, 10(03), 10-03.

40. Cao, D. M., Sayed, M. A., Habib, S. A., Islam, M. T., Mia, M. T., Ayon, E. H., ... & Raihan, A. (2024). Advanced cybercrime detection: A comprehensive study on supervised and unsupervised machine learning approaches using real-world datasets. Journal of Computer Science and Technology Studies, 6(1), 40-48.

41. Mia, M. M., Al Mamun, A., Ahmed, M. P., Tisha, S. A., Habib, S. A., & Nitu, F. N. (2025). Enhancing financial statement fraud detection through machine learning: A comparative study of classification models. Emerging Frontiers Library for The American Journal of Engineering and Technology, 7(09), 166-175.

42. Bhuiyan, R. J., Akter, S., Uddin, A., Shak, M. S., Islam, M. R., Rishad, S. S. I., ... & Hasan-Or-Rashid, M. (2024). Sentiment analysis of customer feedback in the banking sector: A comparative study of machine learning models. The American Journal of Engineering and Technology, 6(10), 54-66.

43. Siddique, M. T., Jamee, S. S., Sajal, A., Mou, S. N., Mahin, M. R. H., Obaid, M. O., ... & Hasan, M. (2025). Enhancing automated trading with sentiment analysis: Leveraging large language models for stock market predictions. The American Journal of Engineering and Technology, 7(03), 185-195.

44. Sajal, A., Chy, M. S. K., Jamee, S. S., Uddin, M. N., Khan, M. S., & Gharami, A. K. & Ahmed, M.(2025). Forecasting Bank Profitability Using Deep Learning and Macroeconomic Indicators: A Comparative Model Study. International Interdisciplinary Business Economics Advancement Journal, 6(06), 08-20.

45. Akhi, S. S., Shakil, F., Dey, S. K., Tusher, M. I., Kamruzzaman, F., Jamee, S. S., ... & Rahman, N. (2025). Enhancing banking cybersecurity: An ensemble-based predictive machine learning approach. Am. J. Eng. Technol., 7(03), 88-97.

46. Hossain, S., Sajal, A., Jamee, S. S., Tisha, S. A., Siddique, M. T., Obaid, M. O., ... & Haque, M. S. U. (2025). Comparative analysis of machine learning models for credit risk prediction in banking systems. The American Journal of Engineering and Technology, 7(04), 22-33.

47. Hossen, M. A., BHATTACHARJEE, B., DEY, S., JAMEE, S., OBAID, M., MIA, M., ... & SHARIF, M. (2025). Business analytics for customer segmentation: A comparative study of machine learning algorithms in personalized banking services. International Journal of Economics Finance & Management Science, 10(03), 1-13.

48. Nath, P. C., Chy, M. S. K., Hossain, M. R., Miah, M. R., Jamee, S. S., Sharif, M. K., ... & Ahmed, M. (2025). Comparative Performance of Large Language Models for Sentiment Analysis of Consumer Feedback in the Banking Sector: Accuracy, Efficiency, and Practical Deployment. Frontline Marketing, Management and Economics Journal, 5(06), 07-19.

49. S. M. Rezvi, F. Mujtahid, K. R. Ahmed, A. Madhiya, V. SOLANKI and S. S. Jamee, "An Intelligent Multi-Source Data Fusion System for Sales Demand Forecasting in Retail and E-Commerce," 2026 Third International Conference on Innovations in Cybersecurity and Data Science (ICICDS), Pathum Thani, Thailand, 2026, pp. 1112-1117, doi: 10.1109/ICICDS70526.2026.11604912.

50. Ahmmed, M. J. Pritom Das, Tamanna Pervin, Sadia Afrin, Sanjida Akter Tisha, Md Mehedi Hassan, & Nabila Rahman.(2024). COMPARATIVE ANALYSIS OF MACHINE LEARNING ALGORITHMS FOR BANKING FRAUD DETECTION: A STUDY ON PERFORMANCE, PRECISION, AND REAL-TIME APPLICATION. International Journal of Computer Science & Information System, 9(11), 31-44.

51. Ahmed, M. P., Arif, M., Chowdhury, M. S., Bhuiyan, R. J., Rahman, T., Ahmmed, M. J., ... & MAMUN, M. (2024). Comparative analysis of machine learning techniques for accurate lung cancer prediction. Am. J. Eng. Technol, 6, 92-103.

52. Akter, S., Mahmud, F., Rahman, T., Ahmmed, M. J., Uddin, M. K., Alam, M. I., & Jui, A. H. (2024). A comprehensive study of machine learning approaches for customer sentiment analysis in banking sector. The American Journal of Engineering and Technology, 6(10), 100-111.

53. Sweet, M. M. R., Arif, M., Uddin, A., Sharif, K. S., Tusher, M. I., Devi, S., ... & Sarkar, M. A. I. (2024). Credit risk assessment using statistical and machine learning: Basic methodology and risk modeling applications. International Journal on Computational Engineering, 1(3), 62-67.

54. Arif, M., Ahmed, P., Mamun, A. A., Uddin, K., Mahmud, F., Rahman, T., ... & Hossain, S. (2024). Dynamic pricing in financial technology: evaluating machine learning solutions for market adaptability.

International Interdisciplinary Business Economics Advancement Journal, 5(10), 13-27.

55. Ahmed, M. P., Arif, M., Al Mamun, A., Mahmud, F., Rahman, T., Ahmmed, M. J., ... & Uddin, M. K. (2024). A comparative study of machine learning models for predicting customer churn in retail banking: insights from logistic regression, random forest, GBM, and SVM. Journal of Computer Science and Technology Studies, 6(4), 92-101.
56. Ahmmed, M. J. (2025). Systematic review and quantitative evaluation of advanced machine learning frameworks for credit risk assessment, fraud detection, and dynamic pricing in US financial systems. International Journal of Business and Economics Insights, 5(3), 1329-1369.
57. Rahman, M. M., & Ahmmed, M. J. (2026). AI-Driven Risk Analytics Models for Early Detection of Financial Noncompliance in Multi-Branch Banking Systems. American Journal of Data Science and Analytics, 7(04), 81-123.
58. Hossen, M. E. ., Akhter, A. ., Ghosh, S. ., Khandaker, M. ., Azam, M. N. ., Malek, H. A. ., Naher, K. ., & Bhuiyan, M. M. R. . (2026). Predicting Infectious Disease Outbreaks Using Machine Learning and Real-Time Epidemiological Data: Leverage Social Media, Environmental, And Public Health Data to Forecast Outbreaks Like Influenza, COVID-19, Or RSV. International Journal of Medical Science and Public Health Research, 7(02), 7–17. https://doi.org/10.37547/ijmsphr/Volume07Issue02-02
59. Umam, S., & Razzak, R. B. (2025, November). A 20-Year Overview of Trends in Secondhand Smoke Exposure Among Cardiovascular Disease Patients in the US: 1999–2020. In APHA 2025 Annual Meeting and Expo. APHA.7
60. Razzak, R. B., & Umam, S. (2025, November). Health Equity in Action: Utilizing PRECEDE-PROCEED Model to Address Gun Violence and associated PTSD in Shaw Community, Saint Louis, Missouri. In APHA 2025 Annual Meeting and Expo. APHA.8
61. Razzak, R. B., & Umam, S. (2025, November). A Place-Based Spatial Analysis of Social Determinants and Opioid Overdose Disparities on Health Outcomes in Illinois, United States. In APHA 2025 Annual Meeting and Expo. APHA.
62. Khan, M. S., Gharami, A. K., Nitu, F. N., Uddin, M. N., Ahmed, M., Roy, M. K., & Yezdani, S. (2025). Deep Learning-Driven Customer Segmentation in Banking: A Comparative Analysis for Real-Time Decision Support. International Interdisciplinary Business Economics Advancement Journal, 6(08), 9-22.
63. Sajal, A., Chy, M. S. K., Jamee, S. S., Uddin, M. N., Khan, M. S., & Gharami, A. K. & Ahmed, M.(2025). Forecasting Bank Profitability Using Deep Learning and Macroeconomic Indicators: A Comparative Model Study. International Interdisciplinary Business Economics Advancement Journal, 6(06), 08-20.
64. Siddique, M. T., Uddin, M. J., Chambugong, L., Nijhum, A. M., Uddin, M. N., & Shahid, R. & Ahmed, M.(2025). AI-Powered Sentiment Analytics in Banking: A BERT and LSTM Perspective. International Interdisciplinary Business Economics Advancement Journal, 6(05), 135-147.
65. Ayub, M. I., Gharami, A. K., Nitu, F. N., Uddin, M. N., Islam, M. I., Nijhum, A. M., ... & Yezdani, S. (2025). AI-driven demand forecasting for multi-echelon supply chains: Enhancing forecasting accuracy and operational efficiency through machine learning and deep learning techniques. Emerging Frontiers Library for The American Journal of Management and Economics Innovations, 7(07), 74-85.
66. Islam, M. M., Sarkar, M. A. R., Ahmed, M., Moniruzzaman, S. M., & Uddin, M. N. (2009). Germination, vigour and emergence indicators of Corchorus olitorius L, seed and their relationship as influenced by seed sources. Bangladesh J. Jute and Fib. Res, 29(1&2), 1-8.
67. Uddin, M. N., & Aziz, M. M. (2026). Shapley Value-Guided Adaptive Ensemble Learning for Explainable Financial Fraud Detection with US Regulatory Compliance Validation. arXiv preprint arXiv:2604.14231.
68. Uddin, M. N. (2026). Explainable Graph Neural Networks for Interbank Contagion Surveillance: A Regulatory-Aligned Framework for the US Banking Sector. arXiv preprint arXiv:2604.14232.
69. Uddin, M. N. (2026). Does Founding Team Human Capital Heterogeneity Predict Venture Survival? Evidence from the Kauffman Firm Survey. Evidence from the Kauffman Firm Survey (April 04, 2026).
70. Uddin, M. N. (2026). Cost Efficiency Dynamics in Indian Public Sector Banks: A DEA-Based Analysis with Second-Stage Tobit Estimation and Post-Reform Comparative Evidence, 2009-2023. Available at SSRN 6680922.

71. Mottalib, M. Y., Yousuf, M., Bhuiyan, E. R., Habib, S. M., Dey, S. K., Gazi, M. S., ... & Anik, A. (2026). Machine Learning-Based Cyber Defense for Cloud Infrastructure: An Adaptive Deep Q-Network Architecture for Intelligent Intrusion Detection and Automated Threat Mitigation. arXiv preprint arXiv:2608.12190.

72. Tisha, S. A., Mottalib, M. Y., Bhuiyan, E. R., Anik, A., Ullah, S. W., Risalat, M., & Chockroborty, S. L. (2026). Artificial Intelligence-Driven Cybersecurity Framework for Enterprise Threat Detection: A Machine Learning Approach. The USA Journals TAJET, 8(07), 25-37.

73. Uddin, M. N., Bhuiyan, E. R., Risalat, M., Orpita, R. T., Anik, A., & Ullah, S. W. (2026, July). A FinCEN-Aligned Framework for Explainable Graph Neural Networks for Fraud Contagion Detection and Automated SAR Generation in US Interbank Payment Networks. In 2026 IEEE International Conference on Artificial Intelligence and Mechatronics Systems (AIMS) (pp. 1-6). IEEE.

74. Uddin, M. N., Orpita, R. T., Bhuiyan, E. R., & Anik, A. (2026). ZK-SR117: A Chunked Zero-Knowledge Attestation Design for Aggregated Fair-Lending Metrics, with a Control Mapping toward Full SR 11-7 Coverage. arXiv preprint arXiv:2608.02664.

75. Mottalib, M. Y., Yousuf, M., Bhuiyan, E. R., Habib, S. M., Dey, S. K., Gazi, M. S., ... & Anik, A. (2026). Machine Learning-Based Cyber Defense for Cloud Infrastructure: An Adaptive Deep Q-Network Architecture for Intelligent Intrusion Detection and Automated Threat Mitigation. arXiv preprint arXiv:2608.12190.

76. Uddin, M. N., Orpita, R. T., Anik, A., Bhuiyan, E. R., Risalat, M., Ullah, S. M., & Ahamed, A. (2026). CT-HEG: A Bidirectional, Timestamp-Attributed Event Graph for ICU In-Hospital Mortality Prediction-An Architectural Ablation Study. arXiv preprint arXiv:2608.02663.

77. Uddin, M. N., Orpita, R. T., Bhuiyan, E. R., & Anik, A. (2026). ZK-SR117: A Chunked Zero-Knowledge Attestation Design for Aggregated Fair-Lending Metrics, with a Control Mapping toward Full SR 11-7 Coverage. arXiv preprint arXiv:2608.02664.

78. Uddin, M. N., Bhuiyan, E. R., Risalat, M., Orpita, R. T., Anik, A., & Ullah, S. W. (2026, July). A FinCEN-Aligned Framework for Explainable Graph Neural Networks for Fraud Contagion Detection and Automated SAR Generation in US Interbank Payment Networks. In 2026 IEEE International Conference on Artificial Intelligence and Mechatronics Systems (AIMS) (pp. 1-6). IEEE.

79. Tisha, S. A., Mottalib, M. Y., Bhuiyan, E. R., Anik, A., Ullah, S. W., Risalat, M., & Chockroborty, S. L. (2026). Artificial Intelligence-Driven Cybersecurity Framework for Enterprise Threat Detection: A Machine Learning Approach. The USA Journals TAJET, 8(07), 25-37.

80. Mozumder, M. A. S., Mia, M. M., Nipa, R. A., Anik, A., Akter, S., & Choudhury, M. Z. M. E. (2026). A Hybrid Graph Neural Network and Large Language Model Framework for Insider Threat Detection via Behavioral Graph and Semantic Profiling. Colombo Scientific Journal of Computer Science & Information System, 11(05), 22-32.

81. Arif, M., Mozumder, M. A. S., Sajal, A., Jahed, K. A., Sharif, M. K., Anik, A., & Ahmed, M. (2026). Machine Learning–Driven Strategic Decision-Making: An Empirical Analysis of Employee Attrition Prediction Using Ensemble Models. The USA Journals TAJMEI, 8(04), 24-34.

82. Uddin, M. N., Orpita, R. T., Bhuiyan, E. R., & Anik, A. AI Compute Import Dependence and Dollar Invoicing: Cross-Country Evidence for a Compute-Dollar Association. Available at SSRN 7003871.

83. Mozumder, M. A. S., Mia, M. M., Nipa, R. A., Anik, A., Akter, S., & Choudhury, M. Z. M. E. (2026). A Hybrid Graph Neural Network and Large Language Model Framework for Insider Threat Detection via Behavioral Graph and Semantic Profiling. Colombo Scientific Journal of Computer Science & Information System, 11(05), 22-32.