EAI Endorsed Transactions
on Scalable Information Systems

Research Article **EAI.EU**

Risk Propagation and Resilience Early Warning in Manufacturing Supply Chains Under Generative AI–Driven Public-Opinion Disturbances

Jing Liu[1*], Jin Jia[2] and Peng Liu[3]

[1]School of Liberal Arts and Sciences, North China Institute of Aerospace Engineering, LangFang, Hebei, 065000, China.
[2]Audit Office, China People's Police University, Langfang, Hebei, 065000, China.
[3]School of Computer Science, North China Institute of Aerospace Engineering, Langfang, Hebei, 065000, China.

Abstract

INTRODUCTION: To address three challenges arising from the large-scale generation and cross-platform dissemination of generative AI content—namely, the difficulty of quantifying external information disturbances in manufacturing supply chains, time-varying risk-propagation relationships, and the limited credibility of early warnings—we propose a Generative-AI-Aware Risk Transmission and Resilience Prediction Network GAI-RTPNet.
OBJECTIVES: The network constructs a node-level exogenous disturbance representation from semantic polarity, event relevance, propagation intensity, source credibility, content novelty, AI-generation probability, and semantic consistency. It updates the edge weights of supply relationships through disturbance-driven dynamic graph attention and jointly models risk evolution across nodes and time periods using a risk-propagation gate and temporal self-attention.
METHODS: Based on a shared spatiotemporal representation, GAI-RTPNet applies multitask learning to simultaneously predict future node risk, propagation probabilities, and system resilience. It further combines heteroscedastic and model uncertainty estimates to produce four-tier trustworthy early warnings.
RESULTS: Experiments on GDELT 2.0, SupplyGraph, and HC3 show that GAI-RTPNet achieves values of 0.048, 0.867, and 0.041 for Risk MAE, Propagation Accuracy, and Resilience MAE, respectively, with a mean warning lead time of 4.1 d. Its performance remains stable under noise, random missingness, and cross-scenario testing.
CONCLUSION: The results support the effectiveness of explicitly incorporating generative-AI-driven public-opinion disturbances into dynamic-graph risk propagation and resilience prediction, providing a unified data-driven approach to risk identification, risk-propagation inference, and proactive warning in manufacturing supply chains.

Keywords: generative AI; public-opinion disturbance; manufacturing supply chain; risk propagation; resilience prediction; uncertainty-aware early warning

Received on 25 August 2026, accepted on 11 September 2026, published on 22 September 2026

doi: 10.4108/eetsis.14733

1. Introduction

Supply chains in manufacturing involve interconnected processes in supply, production, storage facilities, logistics and demand in the market. The local anomalies are capable of spreading across various nodes through supply dependencies, order linkages and logistic connections. News, platform content and industry events are increasingly becoming significant external information sources to make decisions about operations as the digitalization and real-time functioning of supply networks progresses. According to a capability point of view, Jackson et al. [1], generative AI is changing the way supply chain data analytics works,

[*]Corresponding author. Email: jingliu0805@163.com

knowledge creation, and decision support. Nevertheless, the increased speed at which produced information spreads has introduced new risk factors of authenticity, timeliness and traceability of external information. The opportunities and resilience risks related to generative AI were also discussed by Boone et al. [2], and they focused on how information credibility and decision bias impact the quality of supply chain responses. Therefore, manufacturing supply chain risk analysis should not be limited to traditional operating-state predictions but rather must simultaneously consider external information disturbance modeling, network propagation, and predictive reliability.

Studies of supply chain resilience have yielded methods that are founded on predictive AI, deep learning, explainable AI, and dynamic adaptation. The roles played by AI and prescriptive analytics in risk identification and recovery choice were discussed by Smyth et al. [3] who pointed out that the predictions should be more tightly associated with the long-term resilience actions. In the context of network evolution, Ivanov [4] explained that the research on supply chain resilience is moving away towards recovery in a localized way to dynamic adaptation at various levels and actors. Deep learning, survival analysis and explainable AI were integrated by Li et al. [5] in end-to-end resilience management. Empirical studies that are related show that AI capabilities may enhance the effectiveness of risk sensing and response efficiency [6], but their conversion into resilience performance will depend on the dynamism of the supply chain [7]. A big-data analytics culture driven by AI may also help to build agility and resilience in complex supply chains [8]. These studies form a basis of data-driven resilience modeling, yet the majority of the current approaches are based mainly on the internal operational data of the companies.

The shift in the focus of real-time risk identification is moving towards dynamic analysis of indicators rather than fixed analysis of heterogeneous data and temporal state inference. The technical trends in supply chain risk management include the dynamic data fusion, online risk identification, and real-time forecasting, according to the systematic review by Deiva Ganesh and Kalpana [9]. Real-time data mining research has also demonstrated that unstructured online data may add to order, inventory and logistics data and indicate impending disruption at an earlier stage [10]. Hybrid deep learning has been applied to detect disruptions in cognitive digital supply chain twins [11] in terms of model architecture, whereas intelligent digital twins can be used to test stress, resilience, and viability [12]. Natural-language data could also be transformed into structured inputs to a logistics simulation through AI [13]. All these works form the foundation of joint modeling of textual semantics, operational time series, and network relationships. Most models, however, do not use external information as their main input and therefore are not able to keep updating supply relationships or risk propagation states.

The mechanism of the occurrence of external information disturbances has been additionally changed by Generative AI. The massive amounts of automatically generated content spread quickly, are highly semantically repetitive, and can easily be repeated across platforms with little expense. If this type of content is about supplier closures, quality of products, logistics bottlenecks, recalls or predictions of demand, the resultant information shock can impact procurement, production, and market behavior before the underlying facts are even confirmed. Akhtar et al. [14] applied AI and machine learning in order to identify fake news and misinformation that might lead to supply chain disruptions, indicating that external textual data can act as an early warning signal of risks. It was later discovered that misinformation may enhance operational disturbances due to the actions of individuals who participate in the supply chain [15], whereas verification of information and credibility tools can reduce these risks [16]. Park and Nan [17] in their review of the concept of generative AI and misinformation also added that there are generative technologies that not only increase the capacity of content-production but also make it more difficult to detect and govern. Therefore, the use of a single sentiment score or event label would not be enough to describe the level of risk intensity and the potential of spreading of generative-AI-induced public-opinion disturbances. Compared with conventional human-written news, generative-AI content can be produced at much lower marginal cost, scaled rapidly, replicated across platforms, and accompanied by uncertain generation-detection scores; these properties create a disturbance regime in which volume, novelty, credibility, and generation probability must be interpreted jointly rather than as a single sentiment signal [17].

Current studies of generative AI in supply chains have been mostly concerned with coordination, productivity, and decision support. Li et al. [18] reported a correlation between the generative-AI-based supply chain management and coordination mechanisms, but the simulation on how generated content is introduced into risk-propagation channels is still restricted. The process of text-mining news, web data, and online data is able to derive risk indicators out of unstructured data and offer supplementary information to detect disruption [19]. Supply network capabilities need to consider structural connectivity, environmental change, and adaptation processes in order to be resilient system resiliency perspective [20]. Explainable AI may also determine important factors that contribute to resilience and enhance the clarity of model choices [21]. Therefore, external public-opinion characteristics ought not to have entered a predictor as fixed-risk scores following text analysis but should have coevolved with supply-network conditions and relationships. The supply chain disruptions display significant network propagation, and their size is determined by three factors: topology, tier dependency and operating conditions. The study by Habibi et al. [22] demonstrated that factoring the disruption spread in the evaluation of the network resilience significantly alters the assessments of risks. Subsequent studies have found that various levels and connectivity patterns have different influences on cascade size and recovery behavior [23]. Propagation characterization based on topological and operational viewpoints may represent network resilience better [24], whereas reachability after failure and functional survival are critical aspects of logistics-network resilience [25]. Distributionally robust optimization can be used to make strong mitigation and recovery choices when confronted with random perturbations [26]. Such methods, nevertheless, usually presuppose predetermined disruption situations or constant disturbance inputs. They cannot properly reflect the variations in relationship weights due to real-time generative-AI public opinion signals or their dynamic impact on propagation between nodes. Classical disruption-propagation and ripple-effect studies likewise show that multi-tier dependence and network structure govern how local shocks spread through supply chains [27,28].

Altogether, the current approaches have three major shortcomings. To start with, there is no structural coupling between the analysis of the public opinion and the modeling of supply chain state. Semantic polarity, propagating intensity, credibility of sources, novelty of contents and generative-content features do not exist in a unified representation as dynamic node-level exogenous disturbance. Second, standard graph-temporal models acquire the relationship weights largely on fixed topologies or working conditions, thus it is hard to capture the process of increasing, decreasing or restructuring the effective channels of propagation due to the external shocks of information. Third, the prediction of node risk, the identification of propagation paths, the evaluation of system resilience, and the prediction reliability are usually trained with distinct goals that do not allow the emergence of a common warning message that could be directly used to provide proactive reactions. In essence, all these shortcomings occur due to the fact that the disturbances in the information space, risk propagation in the physical supply network, and warning reliability in the decision space have not been combined into one learning model.

In order to overcome these constraints, we present GAI-RTPNet. The model starts with a Generative-AI Disturbance Encoder that simultaneously encodes semantic, propagation, source, and generation properties to node-level extrinsic disturbances. This is followed by Disturbance Driven Dynamic Risk Propagation, whereby external information shocks can be directly used in the computation of dynamic edge weights estimation, risk propagation gating, and node state updates. On the basis of a common spatiotemporal representation, a multi-horizon risk-resilience joint prediction module provides a simultaneous output of the future node risk, propagation probabilities, and system resilience. Lastly, heteroscedastic and model uncertainty estimates are combined and independently used to estimate warning severity and confidence. In contrast to traditional models that add the concept of public opinion to supply networks at the last step, GAI-RTPNet considers the signal of generative-AI-driven public opinion to be the dynamic factor of the evolution of supply network states. It thus combines the concepts of disturbance perception, risk propagation, resilience prediction and reliable early warning into an end-to-end learning pipeline. The contribution is therefore the coupling mechanism rather than the individual use of standard neural modules: node-level generative-AI disturbance evidence is allowed to modulate relation strength and message gating before joint risk-resilience prediction and confidence-aware warning.

2. Problem Formulation and Supply Chain Risk Representation Under Generative AI–Driven Public-Opinion Disturbances

2.1. Dynamic Manufacturing Supply Chain Representation

As illustrated in Fig. 1, the manufacturing supply chain risk issue in the generative-AI-induced public-opinion disturbances is defined in three layers, which include external information disturbances, the dynamic supply network, and risk-resilience states. This disturbance is expressed as external public-opinion information using semantic features, event associations, scale of propagation, reliability of source, and generative-content features. This mapping of the representation is then done through event relevance and node structural exposure to node level exogenous shocks. The manufacturing supply chain is also represented as a dynamic attributed graph in which the node operating pressure, network topology and external information shocks are all part of a single state space. In such a representation, the concept of risk in a node implies the deterioration of the local operations in the prediction horizon, whereas the concept of system resilience denotes the capacity of the network to function, mitigate risks, and restore itself after being disrupted. Future node risk, risk-propagation relationships, critical propagation paths, and system resilience are specified as joint prediction targets. This formulation does not remove the difference between the information and operation domains but prevents directly comparing the strength of public opinion with the real loss of operations.

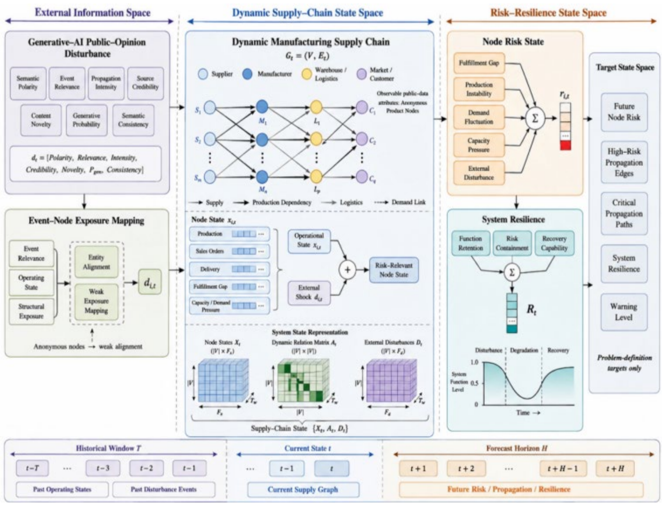

Figure 1. Problem formulation and Generative-AI-aware risk–resilience representation of dynamic manufacturing supply chains

The manufacturing supply chain at time t can be portrayed using a dynamic attributed graph. Theoretical models can have nodes that represent suppliers, manufacturers, warehouse/logistics entities and market/customer-side entities and relational edges that describe supply dependencies, production coordination, logistics links, and demand connections. Empirical validation is done through the anonymized product-relation network in SupplyGraph. Consequently, nodes in the experimental dynamic graph will denote observable product-level entities and their supply chain relationships. Since the anonymized features cannot be used to identify real companies, the relevant interpretations are limited to the states of the nodes and the network structure. This approach clearly separates the theoretical entity model and the observatory granularity of the public data maintaining the verifiability of the dynamic-graph propagation mechanism.

$$G_t=(V,E_t,X_t,A_t) \tag{1}$$

where V is the node set; E_t denotes time-varying supply relations, production dependencies, logistics connections, and demand associations; $X_t \in \mathrm{R}^{N\times F}$ is the node-state matrix; and A_t is the dynamic relation-weight matrix. Node states comprise variables such as production volume, sales orders, distributor deliveries, factory dispatches, fulfillment gaps, capacity pressure, and demand volatility. To reduce the influence of different measurement scales on graph message passing, continuous operational variables are standardized using statistics calculated from the training interval, and the same statistics are applied to the validation and test sets. Categorical relations participate in edge-weight calculation through relation-type embeddings. Thus, node attributes characterize local operating conditions, the relational structure constrains potential propagation paths, and external disturbances represent time-varying information shocks.

$$X_t^H=\{G_{t\text{-}T+1},\ldots,G_t\},\quad Y_t^F=\{G_{t+1},\ldots,G_{t+H}\} \tag{2}$$

Given a historical window of length T and a prediction window of length H, the model simultaneously outputs future node risks, future high-risk propagation edges, critical propagation paths, system resilience, and warning levels. This multiobjective formulation maintains consistency between local risk prediction and network-level functional states: node risk determines anomaly severity, propagation probability identifies the relational channels across which risk may spread, system resilience measures changes in overall network functionality, and the warning level further incorporates predictive uncertainty to provide risk stratification.

$$O_t=\{\widehat{R}_{t+1:t+H},\widehat{P}_{t+1:t+H},\widehat{C}_{t+1:t+H},\widehat{\mathrm{Res}}_{t+1:t+H},\widehat{W}_{t+1:t+H}\} \tag{3}$$

2.2. Representation of Generative-AI-Driven Public-Opinion Disturbances

The external disturbances of public opinion are mutually caused by semantics, events related associations, the magnitude of propagation, the source of information, and attributes of generative content. Generative Probability is considered as one type of conditional evidence and modeled together with Semantic Polarity, Event Relevance, Propagation Intensity, Source Credibility, Content Novelty, and Semantic Consistency. The idea is to evaluate the possible contribution of an information item to supply chain node status changes not to understand how a text was produced. Hence, the representation of exogenous disturbance due to the Generative Probability is only raised when the event is relevant to the supply chain, its propagation intensity is at a high enough level and is accompanied by other semantic evidence, which also decreases the role of a single detection score in risk assessment.

$$u_{k,t}=\left[s_{k,t},r_{k,t},p_{k,t},c_{k,t},n_{k,t},g_{k,t},q_{k,t}\right] \tag{4}$$

where s, r, p, c, n, g, and q denote Semantic Polarity, Event Relevance, Propagation Intensity, Source Credibility, Content Novelty, Generative Probability, and Semantic Consistency, respectively. Before entering the fusion module, all components are normalized while retaining their directional and reliability information. Negative semantic polarity characterizes event orientation, propagation intensity reflects dissemination scale, source credibility adjusts evidence reliability, novelty distinguishes repeated narratives from new information, and semantic consistency constrains the stability of generative content across multiple sources.

$$\alpha_{k,t}=\mathrm{softmax}_k\left(w_\alpha^\top \tanh\left(W_u u_{k,t}+W_e e_{k,t}\right)\right),\quad d_t=\sum_k \alpha_{k,t}\, W_d u_{k,t} \tag{5}$$

For news reports with publicly available entity identifiers, candidate node associations can be established using organizational entities and event semantics. For the anonymized SupplyGraph data, weak alignment is constructed using temporal information, manufacturing/FMCG sector topics, event relevance, node operating states, and structural exposure. This mechanism is used only to generate node-level exogenous disturbance inputs and does not establish firm-level causal correspondences. The learned propagation results therefore describe risk dependencies and state evolution on the observed graph rather than inferring the identities or causal relationships of the actual commercial entities represented by anonymized nodes. Because anonymization removes entity-level ground truth, alignment accuracy cannot be directly estimated from these data; accordingly, the model does not claim firm-level event-to-node matching or causal attribution. Alignment noise can dilute event-relevance estimates, assign exogenous evidence to structurally exposed but semantically mismatched nodes, and thereby attenuate downstream risk and propagation predictions. The reported noise/missingness robustness results are therefore interpreted only as indirect evidence of tolerance to imperfect inputs, not as a direct benchmark of entity-alignment accuracy.

$$m_{i,k,t}=\sigma\left(w_m^\top\left[e_{k,t}\Vert \mathrm{h}_{i,t}^{\mathrm{op}}\Vert \mathrm{h}_i^{\mathrm{str}}\right]\right),\quad D_{i,t}=\sum_k m_{i,k,t}\,\alpha_{k,t} d_{k,t} \tag{6}$$

2.3. Definitions of Risk State and Supply Chain Resilience

Node risk is a continuous measure of the level of service failure and operational decline during the prediction window. Risk supervision is constructed from fulfillment gaps, production instability, demand volatility, capacity pressure, and factory-side flow imbalance in the publicly available operational trajectories, resulting in a reproducible relative risk scale. It is a measure of the degree of anomaly in the network state, and it is not used as a label to describe real corporate events. In this continuous formulation, the gradient information is retained when the deviation is small and up until the severe anomaly, and is presented as a common foundation of multi-horizon regression, risk-growth, and threshold calibration of warnings.

$$R_{i,t}=\sigma\left(w_r^{\top}\phi\left(SD_{i,t},IP_{i,t},DD_{i,t},DV_{i,t},CS_{i,t},D_{i,t}\right)+b_r\right) \tag{7}$$

The characteristic of system resilience is the network capability to remain functional, reduce risk and recover during disturbances. Within the graph framework, the model synthesizes node serviceability, risk concentration, and recovery tendencies in order to form a continuous resilience goal, enabling it to demonstrate the effect of system-level differences on different propagation structures at the same node-risk conditions. The given definition highlights the complementarity of risk and resilience: risk is a negative local or propagating state, and resilience is the ability of the network to absorb shocks and regain its functionality. Both of them define the severity and the duration of the warning.

To make the supervision targets reproducible at the data level, the risk, propagation, and resilience labels are generated directly from the SupplyGraph trajectories before model training. For node i on day t, five operational stress components are calculated from the four observed time series as follows:

$$FG_{i,t}=\frac{\max\left(\text{SalesOrder}_{i,t}-\text{Delivery}_{i,t},0\right)}{\left|\text{SalesOrder}_{i,t}\right|+\varepsilon}. \tag{8}$$

$$PI_{i,t}=\frac{\left|\Delta\text{Production}_{i,t}\right|}{\left|\text{Production}_{i,t-1}\right|+\varepsilon}. \tag{9}$$

$$DV_{i,t}=\frac{SD_{14}\left(\text{SalesOrder}_{i,t}\right)}{\left|Mean_{14}\left(\text{SalesOrder}_{i,t}\right)\right|+\varepsilon}. \tag{10}$$

$$CP_{i,t}=\frac{\max\left(\text{SalesOrder}_{i,t}-\text{Production}_{i,t},0\right)}{\left|\text{SalesOrder}_{i,t}\right|+\varepsilon}. \tag{11}$$

$$FI_{i,t}=\frac{\left|\text{Production}_{i,t}-\text{FactoryIssue}_{i,t}\right|}{\left|\text{Production}_{i,t}\right|+\varepsilon}. \tag{12}$$

Each component is standardized using the mean and standard deviation estimated from the training interval and then mapped to [0,1] by a sigmoid transformation. The continuous node-risk label for horizon h ∈ {1, 3, 7} d is computed as the mean of the five normalized stress components over days t+1 through t+h:

$$R_{i,t}^{(h)}=\frac{1}{5h}\sum_{\tau=1}^{h}\left(\widetilde{FG}_{i,t+\tau}+\widetilde{PI}_{i,t+\tau}+\widetilde{DV}_{i,t+\tau}+\widetilde{CP}_{i,t+\tau}+\widetilde{FI}_{i,t+\tau}\right),\quad h\in\{1,3,7\}. \tag{13}$$

Public-opinion variables are not included in this target and remain exogenous model inputs, thereby avoiding circular supervision.

The propagation label is constructed only on candidate relations contained in the SupplyGraph topology. Let Q80 denote the 80th percentile of the node-risk labels in the training interval and QΔ75 the 75th percentile of positive one-day risk increments. For a candidate edge (i,j) and horizon h, the edge-level propagation label is defined as:

$$y_{ij,t,h}^{\text{prop}}=\begin{cases}1, & R_{i,t}\geq Q_{80},\ R_{j,t}<Q_{80},\ \text{and}\exists\tau\in[1,h]:\left(R_{j,t+\tau}\geq Q_{80}\ \text{or}\Delta R_{j,t+\tau}\geq Q_{\Delta 75}\right),\\ 0, & \text{otherwise.}\end{cases} \tag{14}$$

This temporal-order rule produces edge-level weak supervision for propagation identification and does not imply firm-level causal transmission. The condition that the target node is below Q80 at time t is intentionally an onset rule: it labels newly activated propagation into a previously non-high-risk target. Persistent high-risk-to-high-risk transmission is not counted as a new positive propagation event and remains represented by continuous node states and dynamic edge weights. Such temporal-order weak labels can still contain false-positive or false-negative edges when co-movement is not true transmission. This label noise may reduce propagation-head calibration or shift attention toward correlated rather than causal paths; topology restriction and the shared node-risk/resilience objectives can limit, but cannot eliminate, this effect.

The resilience label is constructed at graph level from three observable quantities over the same prediction horizon: service retention (SR), risk containment (RC), and recovery tendency (RT). Their data-level definitions and aggregation are:

$$SR_{t,h}=\text{Mean}_{i,\tau}\left[\min\left(\frac{\text{Delivery}_{i,t+\tau}}{\text{SalesOrder}_{i,t+\tau}+\varepsilon},1\right)\right]. \tag{15}$$

$$RC_{t,h}=1-\text{Mean}_{i,\tau}\left[\mathbb{I}\left(R_{i,t+\tau}\geq Q_{80}\right)\right]. \tag{16}$$

$$RT_{t,h}=\text{Norm}_{\text{train}}\left[\max\left(\max_{1\leq\tau\leq h}\bar{R}_{t+\tau}-\bar{R}_{t+h},0\right)\right]. \tag{17}$$

$$Res_{t,h}^{\text{label}}=\frac{SR_{t,h}+RC_{t,h}+RT_{t,h}}{3}. \tag{18}$$

Here, service retention is the network-average delivery fulfillment ratio; risk containment is one minus the proportion of high-risk nodes; and recovery tendency is the normalized positive reduction in aggregate node risk from its maximum within the horizon to the end-of-horizon value. After each term is scaled to [0,1] using training-interval statistics, their arithmetic mean forms the graph-level resilience label, with larger values indicating stronger functional retention, containment, and recovery. Specifically, the aggregate node risk used in recovery tendency is the arithmetic mean of all node-risk labels at a given day. Norm_train denotes min-max scaling of the recovery precursor using its training-interval minimum and maximum, with an epsilon denominator safeguard and clipping to [0,1].

All normalization statistics, percentile thresholds, and rolling-window quantities used in label construction are estimated from the training interval only and then frozen for validation and testing. Thus, the three supervision targets are generated from observable operational trajectories under a common rule, while GDELT- and HC3-derived variables are used only as external disturbance evidence or auxiliary generative-probability inputs.

$$\mathrm{Res}_t=\beta_1 A_t^{\mathrm{abs}}+\beta_2 P_t^{\mathrm{ope}}+\beta_3 S_t^{\mathrm{sup}}+\beta_4 C_t^{\mathrm{rec}}, \quad \sum_j \beta_j=1 \tag{19}$$

3. GAI-RTPNet Model for Risk Propagation and Resilience Early Warning

The GAI-RTPNet views generative-AI-based public opinion perturbation as an exogenous factor and it represents supply chain conditions by using disturbance sensing, dynamic-graph risk propagation, multi-horizon joint risk-resilience prediction and early warning with awareness of uncertainties. The main processing route contains event representation, node-disturbance estimation, dynamic edge updating, risk-gated message passing, temporal-dependency modeling, multitask prediction, and reliable early warning. In contrast to graph-temporal models which only append textual attributes at input, GAI-RTPNet permits external disturbances to continuously contribute to edge weights, message-intensity regulation, and state updates of nodes at various stages. Such semantic shifts in the information space can thus be translated across the supply-network topology as learnable traces of risk propagation. All modules have access to the historical window node states, each module has its own prediction heads that maintain different supervision signals to local risk, edge-level propagation, and graph-level resilience.

3.1. Generative-AI Disturbance Sensing

Generative-AI Disturbance Encoder then uses distinct encoders in order to encode text semantics, temporal metadata, propagation metadata and source attributes, which results in event-level representations. The Event Relevance Attention subsequently utilizes the node state as the Query and the event representation as the Key and Value to predict the conditional relevance of the event and the node. The Source Credibility Gate changes the evidence of an event based on source credibility, diminishing the immediate effect of low-credibility data or data without multisource validation on node states. Generative Interaction Learning continues to learn higher-order interactions between Generative Probability, propagation intensity, novelty, and semantic consistency. Figure 2 shows the computational process of this module. This architecture divides the evaluation of both the relevance and the credibility of the information, and the generative features and whether they are effective disturbances into different learning steps, thus eliminating the control over the exogenous risk input by a solitary text score.

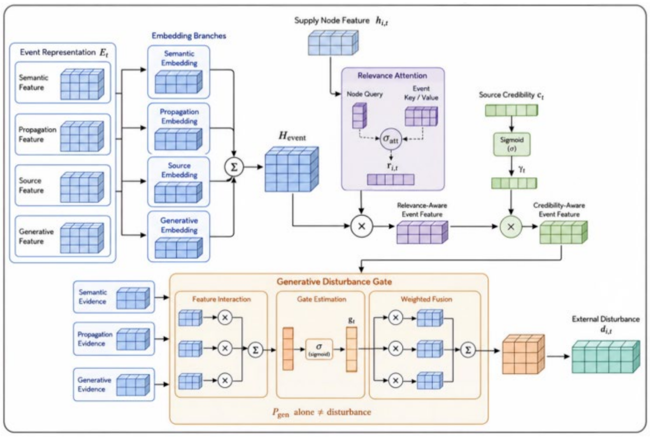

Figure 2. Generative-AI-Aware Disturbance Encoding Module

$$e_{k,t}=\mathrm{TextEnc}\left(x_{k,t}\right)+e_{k,t}^{\mathrm{time}}+e_{k,t}^{\mathrm{prop}} \tag{20}$$

$$\eta_{i,k,t}=\mathrm{softmax}_k\left(\frac{\left(W_q \mathrm{h}_{i,t}^{\mathrm{op}}\right)^{\top} W_k e_{k,t}}{\sqrt{d}}\right) \tag{21}$$

$$\gamma_{k,t}^{c}=\sigma\left(W_c\left[c_{k,t}, q_{k,t}, p_{k,t}\right]+b_c\right) \tag{22}$$

$$\gamma_{k,t}^{g}=\sigma\left(W_g\left[g_{k,t}, p_{k,t}, n_{k,t}, q_{k,t}\right]+b_g\right) \tag{23}$$

$$z_{i,t}^{D}=\mathrm{LN}\left(W_f \sum_k \eta_{i,k,t}\gamma_{k,t}^{c}\gamma_{k,t}^{g}\left(e_{k,t}\| u_{k,t}\right)+W_o \mathrm{h}_{i,t}^{\mathrm{op}}\right) \tag{24}$$

The disturbance-fusion effect in Eq. (24) establishes that Generative Probability is merely a conditional evidence in node-disturbance estimation. In case of high Generative Probability, Event Relevance Attention minimizes the role of text that has no high relevance to the node operations. On the contrary, an event with a lower Generative Probability may still produce a significant exogenous disturbance as long as its negative semantics, intensity of propagation, and level of node exposure are progressively increasing. There is thus a joint determination of the disturbance representation by various evidences, with source credibility and semantic consistency acting as a constraint on extreme propagation signals. In the course of training, the disturbance representation is jointly minimized using backpropagation with the downstream risk, propagation and resilience tasks. It allows the encoder to acquire event properties linked to the state modifications of the supply chain instead of trying to maximize each machine-generated-text detection accuracy independently.

3.2. Disturbance-Driven Dynamic Risk Propagation

The dynamic risk-propagation module of disturbance-driven dynamics describes the mechanism underlying the influence of external information shocks on the propagation effectiveness among nodes, once they enter a supply network. The initial relation graph specifies the possible channels of propagation. The dynamic weight of an edge at time t depends on both the source-node state and the target-node state, relationship type, current risk, and node perturbation simultaneously. Studies on supply chain risk inference based on graph-neural-networks have indicated that explicit relational structure representation can enhance the learning of complex risk dependencies [29]. As illustrated in Fig. 3, this is where GAI-RTPNet offers Disturbance Edge Modulation to the original topology. Edge weights are dynamically changing in response to external shocks and operating conditions without changing the candidate connection set, which allows differentiating structurally present relations with a weak current propagation effect from critical ones working in a high-risk amplification condition.

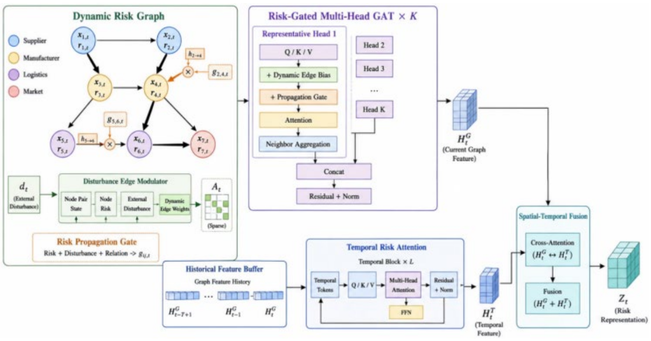

Figure 3. Disturbance-Driven Dynamic Risk Propagation Module

$$c_{i,j,t}=\left[\mathrm{h}_{i,t}^{\mathrm{op}}\|\mathrm{h}_{j,t}^{\mathrm{op}}\|e_{i,j}^{\mathrm{type}}\|R_{i,t}\|R_{j,t}\|z_{i,t}^{D}\|z_{j,t}^{D}\right] \tag{25}$$

$$l_{i,j,t}=\mathrm{LeakyReLU}\left(a^{\top}W_{a}c_{i,j,t}+w_{D}^{\top}\left(z_{i,t}^{D}-z_{j,t}^{D}\right)\right) \tag{26}$$

$$a_{i,j,t}=\frac{\exp\left(l_{i,j,t}\right)A_{i,j}^{\mathrm{prior}}}{\sum_{\mathrm{l}\in\mathrm{N}(i)}\exp\left(l_{i,\mathrm{l},t}\right)A_{i,\mathrm{l}}^{\mathrm{prior}}+\epsilon} \tag{27}$$

In order to minimize the similar treatment of various neighbors during conventional graph aggregation, the model also presents the Risk Propagation Gate. The inputs considered by this gate are neighbor risk, external disruption, type of relationship, and edge-level context and it uses the coefficient to control the size of every message that is sent to the target node. High-risk neighbors with a significant amount of disturbances and relation types that match with the propagation mechanism have large propagation weights, whereas messages of low-relevancy neighbors, with low credibility, or stable conditions are inhibited. The multihead graph attention then combines neighborhood information across various representation subspaces. They estimate the importance of relations independently, followed by residual connections and normalization on the fused representation to enhance training stability during complex network states. Such a design allows modeling the existence of a relationship and the fact that it serves as a high-risk propagation channel independently. To avoid notation ambiguity, e_i,j^type denotes the relation-type embedding, whereas ℓ_i,j,t in Eqs. (26)-(27) denotes the dynamic edge-compatibility logit; e_k,t is reserved for the event representation in Eq. (20). The disturbance-difference term in Eq. (26) represents an exposure gradient across a supply relation rather than an assumed positive effect of disagreement: its contribution is learned through w_D, so a large difference increases the edge logit only when that contrast is informative for propagation in the training data.

$$g_{i,j,t}^{R}=\sigma\left(W_{r}\left(R_{j,t},D_{j,t},FP_{j,t},DV_{j,t},e_{i,j}^{\mathrm{type}}\right)+b_{r}\right) \tag{28}$$

$$m_{j\rightarrow i,t}=a_{i,j,t}g_{i,j,t}^{R}W_{m}\left[\mathrm{h}_{j,t}^{\mathrm{op}}\|z_{j,t}^{D}\right] \tag{29}$$

$$\mathrm{h}_{i,t}^{G}=\mathrm{GRU}\left(\mathrm{h}_{i,t}^{\mathrm{op}},\sum_{j\in\mathrm{N}(i)}m_{j\rightarrow i,t}\right) \tag{30}$$

Following graph propagation, Temporal Self-Attention is utilized on the historical graph representations of every node to discover acute shocks, risk accumulation, delayed transmission, and cross-period dependencies. Graph neural networks have been used to predict post-disaster enterprise-states and have found that state inference in networked supply chains can be enhanced by joint representations of structural dependencies and temporal information [30]. The spatial topology of GAI-RTPNet is not reconstructed with the help of temporal attention. It rather estimates the role of each past time step in predicting the prediction window of all the nodes on the dynamic-graph encoding sequence and adds them up to the current graph state. High attention weights are given to more recent time steps to reflect short-term disturbances, whereas long-term disturbances are accumulated via cross-period attention relations. The shared risk representation produced comprises local node state, dynamic neighborhood propagation, and historical evolution giving a homogeneous spatiotemporal foundation to the various prediction heads. Related temporal-graph models such as TGAT and DySAT also motivate explicit separation of structural and temporal dependencies [31,32]. In GAI-RTPNet, graph attention is evaluated only on observed candidate edges, with approximate cost O(T|E|d), while temporal self-attention is applied independently to each node over the history window, O(NT^2d). The dominant cost is therefore O(T|E|d + NT^2d), rather than all-pairs O(N^2T). With N=41, |E|=684, and T=14, however, the present experiment supports only a complexity analysis and not empirical scalability to thousand-node networks.

$$Q_{i}=H_{i}^{G}W_{Q},\quad K_{i}=H_{i}^{G}W_{K},\quad V_{i}=H_{i}^{G}W_{V},$$

$$Z_{i}=\mathrm{softmax}\left(\frac{Q_{i}K_{i}^{\top}}{\sqrt{d_{\mathrm{h}}}}+B\right)V_{i} \tag{31}$$

$$H_{i,t}=\mathrm{LN}\left[Z_{i,t}+\mathrm{FFN}\left(\mathrm{LN}\left(Z_{i,t}+\mathrm{h}_{i,t}^{G}\right)\right)\right] \tag{32}$$

3.3. Multi-Horizon Joint Risk-Resilience Prediction

There are three task branches that are built on a common spatiotemporal representation, i.e., the Node-Level Risk Head, Edge-Level Propagation Head, and Graph-Level Resilience Head. The node-risk head produces continuous risk estimates 1, 3, and 7 days and the propagation head determines future propagation risks of candidate edges, the resilience head does graph-level pooling of node representations and then predicts the system-resilience pathway during the prediction interval. There is an overall need to consider the combined effects of functional retention, recovery capability, and state change when evaluating system resilience [33]. The suggested model thus preserves the graph-level supervision instead of substituting system resilience with the average risk of each node. Figure 4 demonstrates the interactions between the three prediction heads and the shared representation. The multi-horizon formulation allows the model to recognize both short-term anomalies as well as delayed propagation in comparing the time of risk diffusion and system-resilience reduction.

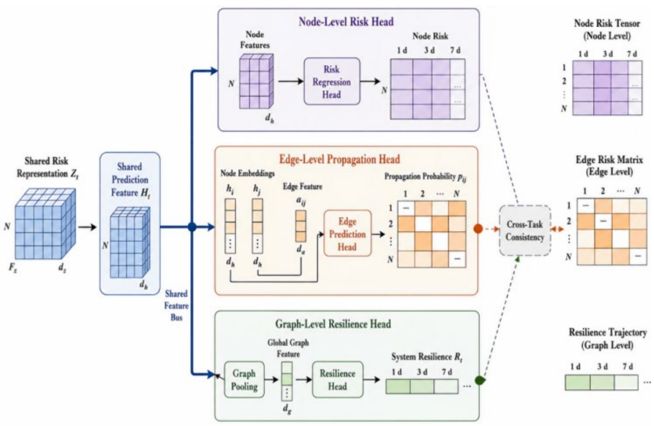

Figure 4. Multi-Task Risk–Propagation–Resilience Prediction Module

$$\widehat{R}_{i,t+\tau}=\sigma\left(W_R^{(\tau)}H_{i,t}+b_R^{(\tau)}\right),\quad \tau=1,\ldots,H \tag{33}$$

$$\widehat{P}_{i,j,t+\tau}=\sigma\left(\mathrm{MLP}_P\left(H_{i,t}\Vert H_{j,t}\Vert a_{i,j,t}\Vert D_{i,t}\Vert D_{j,t}\right)\right) \tag{34}$$

$$\widehat{\mathrm{Res}}_{t+\tau}=\sigma\left(W_{\mathrm{Res}}^{(\tau)}\mathrm{AttnPool}\left(\{H_{i,t}\}_{i=1}^{N}\right)+b_{\mathrm{Res}}^{(\tau)}\right) \tag{35}$$

$$\mathrm{L}_{\mathrm{multi}}=\lambda_1\mathrm{L}_{\mathrm{risk}}+\lambda_2\mathrm{L}_{\mathrm{prop}}+\lambda_3\mathrm{L}_{\mathrm{res}} \tag{36}$$

The three supervision signals create a mutual constraint by means of the common representation. Node-risk loss focuses on the size of local anomalies, propagation loss limits the direction of risk propagation and identifying critical edges, and the resilience loss demands that the graph level representation must maintain recovery and retention trends. Simultaneous training makes it less likely that a model trained on a single task will demonstrate minor local errors but be unable to represent system state transitions. In order to avoid discrepancies in task-loss scales influencing optimization, losses are averaged with predefined weights when training, and the model selection is based on an identical validation set. Prediction-wise, the three task heads have independent outputs. The growth of node risk and propagation probability and a reduction of resilience can then be used as a correlative signal of warning which cannot be compressed into one uninterpretable score.

3.4. Uncertainty-Aware Early Warning and Optimization

The early warnings depending merely on point estimates or classification probabilities may generate high-confidence errors in unseen events, missing inputs and distribution changes. The model together calculates heteroscedastic uncertainty in the continuous-risk problem and model uncertainty via stochastic forward propagation to separate data uncertainty due to observation noise and epistemic uncertainty due to lack of coverage of training data [34]. The Heteroscedastic Risk Head outputs the conditional variance and risk mean at the same time, whereas the MC Uncertainty Head produces predictive variance using a combination of multiple stochastic forward passes. Following the normalization, the two uncertainty parts are combined in Uncertainty Fusion to create confidence evidence. These are left independent of the risk prediction branch allowing the model to make the distinction between high risk predictions that have high confidence and high risk predictions that have insufficient evidence.

$$\left(\mu_{i,\tau}^{(m)},\log\sigma_{i,\tau}^{2,(m)}\right)=f_\theta^{(m)}\left(H_{i,t}\right) \tag{37}$$

$$\bar{R}_{i,\tau}=\frac{1}{M}\sum_m\mu_{i,\tau}^{(m)},\quad U_{i,\tau}=\frac{1}{M}\sum_m\sigma_{i,\tau}^{2,(m)}+\mathrm{Var}_m\left(\mu_{i,\tau}^{(m)}\right) \tag{38}$$

Warning severity combines expected risk, the positive component of the risk-growth rate, future resilience decline, and predictive uncertainty. Warning confidence is defined as $C_{i,t}=e^{-\kappa U_{i,t}}$ and decreases as fused uncertainty increases. Separating severity from confidence prevents intervention intensity from being conflated with model confidence: severity represents the potential impact, whereas confidence indicates the sufficiency of current evidence. When severity is high but confidence is insufficient, the system retains the risk notification and marks it as low-confidence. A higher warning level is issued only when both severity and confidence are high.

$$S_{i,t}=\omega_1\bar{R}_{i,t+H}+\omega_2\max\left(\Delta\bar{R}_{i,t},0\right)+\omega_3\left(1-\widehat{\mathrm{Res}}_{t+H}\right)+\omega_4U_{i,t} \tag{39}$$

$$W_{i,t}=\begin{cases}\text{Normal}, & S<\theta_1\\ \text{Attention}, & \theta_1\le S<\theta_2\\ \text{Warning}, & \theta_2\le S<\theta_3\\ \text{Critical}, & S\ge\theta_3\end{cases} \tag{40}$$

$$\mathrm{L}_{\mathrm{total}}=\mathrm{L}_{\mathrm{multi}}+\lambda_4\mathrm{L}_{\mathrm{unc}}+\lambda_5\mathrm{L}_{\mathrm{cal}}+\lambda_6\mathrm{L}_{\mathrm{reg}} \tag{41}$$

The four warning thresholds are not assigned subjectively. Instead, they are jointly determined on the validation set according to the False Alarm Rate, Missed Alarm Rate, and Mean Early Warning Time. The uncertainty loss $\mathrm{L}_{\mathrm{unc}}$ adopts a heteroscedastic negative log-likelihood, $\mathrm{L}_{\mathrm{cal}}$ constrains the statistical calibration of warning probabilities and prediction intervals, and L_{reg} regularizes dynamic-edge sparsity and temporal smoothness. Statistical calibration of uncertainty estimates improves the interpretability of regression-prediction confidence [35] and is consistent with prospective risk-management approaches to supply chain resilience and survivability [36]. The complete process is shown in Fig. 5. Recent supply-chain early-warning research likewise emphasizes the trade-off among detection performance, false alarms, and actionable lead time [37].

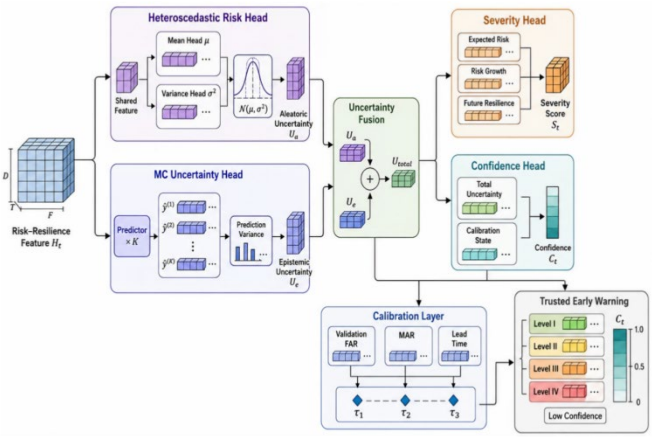

Figure 5. Dual-Uncertainty Trusted Early-Warning Module

4. Experiments and Results

4.1. Datasets and Experimental Settings

The experiments assume a validation approach that integrates both multisource public data and controlled disturbances. The public data sets include SupplyGraph, HC3, and the GDELT 2.0 Event Database/Global Knowledge Graph. They are applied to supply chain operational-state simulation, AI-generated-content probability estimation, and construction of public-opinion disturbances. SupplyGraph offers time series of FMCG supply planning and a product relation graph, and has recently been applied to graph learning and multitasking assessment in intelligent supply chain research [38]. HC3 is a Human-ChatGPT comparison corpus that has been utilized in recent AI-generated-text detection research to measure performance on detection and cross-dataset extrapolation [39]. GDELT 2.0 includes features like news events, organizations, themes, tone, and frequency of reporting and has been studied in terms of its coverage, redundancy, and use in terms of systematic analysis [40]. The three datasets are not directly concatenated by samples. Rather, they generate operating states, AI-generation probabilities and external disturbance inputs based on their individual functions. The HC3-based detector is used only to provide a probabilistic auxiliary feature. It is not used to infer authorship, identity, misconduct, or individual credibility; detector uncertainty and domain bias are treated as model uncertainty rather than as definitive labels.

SupplyGraph spans January 1, 2023 through August 9, 2023 with 221 days of data, 41 product nodes, 684 unique relations, and 62 edge types. It can also give time-series data on Production, Sales Order, Delivery to Distributors, and Factory Issue. Since the node identities are anonymized, GDELT news and SupplyGraph nodes are weakly aligned in terms of time and manufacturing/FMCG industry but not necessarily in a one-to-one manner of entity matching. In particular, daily public-opinion disturbances are filtered by manufacturing, FMCG, and supply chain topics and the node-level disturbance weights are defined by node operational states and structural exposure. This structure will save the statistical connection between the external information and supply chain observations without interpreting entities more than is granulated by the available public data.

The GDELT events are filtered by topics like manufacturing, FMCG, consumer goods, production, supply disruption, logistics, shortage, recall, demand, and distribution and the results are aggregated on a daily basis scale. Demand pressure, fulfillment gaps, production stability and capacity pressure, and changes in factory-side flows are built by Sales Order, Delivery to Distributor, Production, and Factory Issue, respectively, in SupplyGraph. The continuous variables are initially standardized throughout the training period and then grouped as inputs to the model based on a 14 d history. Prediction target horizons of 1, 3 and 7 d are regarded as the short-, medium- and long-term predictions respectively. Supervision of risk and resilience are based on the public time-series variables through the standard principles. The controlled disturbances change solely exogenous shocks or operating conditions in the experimental settings but do not replace the initial observations. The basic statistics and temporal resolution of the three datasets are summarized in Table 1, while their graph structure, main variables, task usage, and data splitting are reported in Table 2.

The preprocessing order is therefore: topic filtering and daily aggregation for GDELT; training-only standardization and 14-day windowing for SupplyGraph; stratified calibration of the HC3 auxiliary detector; and date-level weak alignment before disturbance fusion.

Table 1. Public dataset statistics and temporal resolution

Dataset	Domain	Period/Scale	Verified Structure	Resolution
GDELT 2.0 Event/GKG	Global news/events	2023-01-01–2023-08-09 study window	Continuous public event/news stream; filtered count obtained by query	15 min → daily
SupplyGraph	FMCG manufacturing supply planning	221 daily timestamps	41 product nodes; 684 unique edges; 62 edge types; 8 temporal matrices	Daily
HC3	Human–ChatGPT text	Public release corpus	24,322 question-level records in released corpus	Text sample

Table 2. Dataset graph structure, variables, task usage, and data splits

Dataset	Graph	Main Variables	Task Usage	Split
GDELT 2.0 Event/GKG	Organization/ theme/event associations	event time, organization, theme, tone, mentions, source	Public-opinion disturbance	Aligned by date
SupplyGraph	Product relation graph	production, sales order, delivery, factory issue	Risk/propagation/resilience	Chronological 70/15/15
HC3	None	text, source class	Auxiliary generative-probability detector	Stratified 70/15/15

Supplied Graph data are continuously segmented into 70% training, 15 percent validation and 15 percent test. To avoid future leakage of information, all the standardization parameters are estimated solely based on the training range. HC3 is divided into subsets of 70/15/15 per cent based on the source type of training and calibrating the auxiliary AI-generation-probability detector. The base models consist of XGBoost, LSTM, TCN, GAT-LSTM, DCRNN, TGCN, and PatchTST, and cover both fixed machine learning and purely time-based models, as well as graph-temporal models and more recent sequence-based modeling approaches. Identical data splits, prediction horizons, random seeds and evaluation rules are applied to all models. Data reported in the original publications when the data were exposed to various conditions are not directly used. Such coherent protocol guarantees that the differences in performance are mainly caused by the architecture of models and its use of information, not by the unevenness of data division or assessment standards. The common experimental environment and hyperparameter settings used for all models are listed in Table 3. These hyperparameter settings are specific to the present data scale and validation regime; changes in graph size, relation density, disturbance prevalence, or weak-label noise may alter the optimal configuration and therefore require re-tuning in new operating regimes.

MAE, RMSE and R^2 are used to assess continuous risk prediction. Precision, Recall, F1, AUC and Propagation Accuracy are used to evaluate propagation prediction, and Resilience MAE and Resilience RMSE are used to measure resilience prediction. Performance of early warning is evaluated in terms of Mean Early Warning Time, False Alarm Rate and Missed Alarm Rate. All models are trained multiple times on five different random seeds each time, where the validation set serves both to early-stop and to calibrate the threshold. Monte Carlo uncertainty is calculated based on 20 stochastic forward passes. Evaluation is done at the same time and report point-prediction accuracy, edge-level identification ability, graph-level resilience errors, as well as warning timeliness, thus not being dependent on a single metric when evaluating model efficacy. Graph learning may be able to discover underlying supply chain relationships. As was demonstrated earlier, GNNs have been found to restore missing links in the network structure and node features [41], which gives a theoretical foundation to treating propagation edges as a separate prediction object. For deployment, the three functional inputs correspond to an external news/event feed, internal operational/ERP and supply-network data, and an AI-generation detector; interface synchronization, access control, and data governance are therefore required. The FAR-MAR-lead-time thresholds are validation-set calibrated and should be recalibrated when operating costs, data distributions, or intervention policies change. Wall-clock inference latency was not benchmarked in the present study and remains an implementation limitation. The three datasets are publicly available from their respective source repositories; the model implementation and preprocessing scripts can be provided by the corresponding author as a reproducibility package. The reported FAR, MAR, and warning lead-time values therefore correspond to the present validation-calibrated operating point; because an exhaustive threshold-sensitivity sweep is not included, these metrics should not be interpreted as invariant to alternative threshold settings.

Table 3. Experimental environment and hyperparameter settings

Item	Setting
Implementation	Python 3.11; PyTorch 2.x; PyTorch Geometric 2.x
Computing device	CUDA-compatible GPU, ≥16 GB VRAM
Optimizer	AdamW
Initial learning rate	1.0×10-3
Batch size	32
Maximum epochs	200
Hidden dimension	128
Dynamic graph layers	2
Graph attention heads	4
Temporal attention heads	8
Historical window T	14 daily steps
Prediction horizons H	1, 3, 7 daily steps
Dropout	0.20
Weight decay	1.0×10-4
Early stopping patience	20 epochs
Monte-Carlo forward passes	20
Random seeds	5 independent seeds
Threshold selection	Validation FAR–MAR–lead-time joint calibration

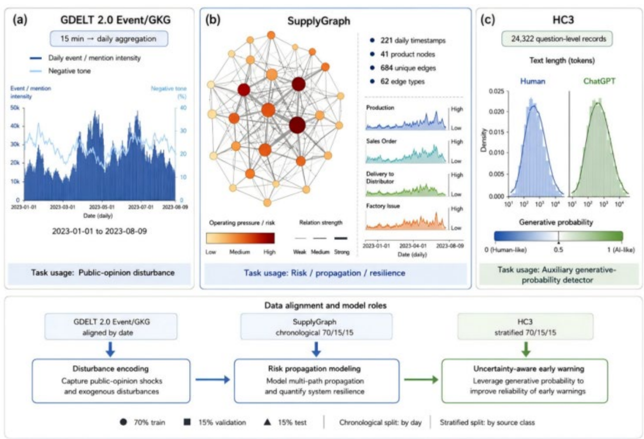

Figure 6. Overview of the public datasets and their functions in GAI-RTPNet

GDELT 2.0 Event/GKG, SupplyGraph, and HC3 represent the multisource public-data experiment structure see Fig. 6. In figure 6(a), the variation in GDELT event/reporting intensity and negative semantic polarity are shown in the study window. The initial high-frequency events after daily aggregation become public-opinion disturbance sequences that are synchronized with the time and supply chain operation conditions. Figure 6(b) gives the structure of the SupplyGraph network and its main operational parameters. The 41 anonymized product nodes and 684 unique relations represent the candidate topology of risk propagation, and the continuous trajectories of Production, Sales Order, Delivery to Distributor, and Factory Issue are employed to generate the node operating states. Figure 6(c) illustrates the distributional properties of Human and ChatGPT texts within HC3. This dataset is purely used to train and calibrate the ancillary AI-generation-probability detector; text provenance is not indirectly understood as supply chain risk. The three datasets eventually enter GAI-RTPNet based on various organizational principles: GDELT offers event-level exogenous evidence, SupplyGraph offers node-edge temporal states, and HC3 offers AI-generated-content probabilities. All these together make up a closed data loop of information disturbance, network proliferation, and reliable early warning.

4.2. Overall Prediction and Early-Warning Performance

Tables 4 and 5 summarize the performance of all models in terms of node-risk prediction, risk-propagation identification, system-resilience prediction, and early warning. XGBoost mainly learns static nonlinear maps, LSTM and TCN learns continuous time dependencies and GAT-LSTM, DCRNN and TGCN additionally learn the structure of graphs and PatchTST performs better at multistep prediction by learning patch-based sequence representations. Besides these features, GAI-RTPNet consists of the ability to consider exogenous public-opinion perturbations, dynamic updating of edges, joint control of both risks and resilience, and calibration of uncertainty. This assessment thus encompasses not just node-risk mistakes but also propagation-edge detection, graph-level resilience mistakes, and warning lead time. GAI-RTPNet also outperforms all other categories in all the four tasks with no significant increase in false or missed alarms.

Table 4. Overall node-risk prediction and propagation-classification performance

Model	Risk MAE↓	Risk RMSE↓	R²↑	Precision↑	Recall↑	F1↑
XGBoost	0.094	0.132	0.781	0.802	0.768	0.784
LSTM	0.081	0.114	0.838	0.831	0.804	0.817
TCN	0.076	0.108	0.853	0.842	0.821	0.831
GAT-LSTM	0.068	0.097	0.881	0.861	0.844	0.852
DCRNN	0.065	0.093	0.889	0.869	0.852	0.860
TGCN	0.062	0.089	0.898	0.875	0.861	0.868
PatchTST	0.059	0.086	0.906	0.883	0.869	0.876
GAI-RTPNet	0.048	0.071	0.936	0.912	0.901	0.906

Table 5. Overall resilience-prediction and early-warning performance

Model	AUC↑	Prop. Acc.↑	Res. MAE↓	Res. RMSE↓	EWT (d)↑	FAR↓	MAR↓
XGBoost	0.861	0.713	0.087	0.119	1.2	0.163	0.208
LSTM	0.889	0.742	0.074	0.103	1.6	0.145	0.184
TCN	0.902	0.758	0.069	0.096	1.8	0.134	0.171

Model	AUC↑	Prop. Acc.↑	Res. MAE↓	Res. RMSE↓	EWT (d)↑	FAR↓	MAR↓
GAT-LSTM	0.921	0.792	0.061	0.087	2.3	0.118	0.149
DCRNN	0.928	0.807	0.058	0.083	2.5	0.111	0.141
TGCN	0.934	0.819	0.055	0.079	2.7	0.104	0.133
PatchTST	0.941	0.811	0.052	0.076	2.9	0.101	0.126
GAI-RTPNet	0.963	0.867	0.041	0.061	4.1	0.072	0.083

The Risk MAE and RMSE of GAI-RTPNet are 0.048 and 0.071, respectively, which are 18.6 percent and 17.4 percent lower than the respective PatchTST values of 0.059 and 0.086. Its R^2 rises to 0.936. Such findings suggest that when given the same historical windows and prediction horizon, external disturbances and dynamic graph relations offer more information towards continual risk estimation not accounted by purely temporal modeling. The benefit is even more straightforward in the propagation task. The accuracy in Propagation is 0.867, higher than TGCN 0.819, DCRNN 0.807, and GAT-LSTM 0.792, and F1 and AUC are 0.906 and 0.963, respectively. Dynamic edge weights and propagation gates may thus also help identify structurally existing candidate relations based on their effective risk-propagation strength. The use of artificial intelligence and big-data analysis is considered to be a significant data asset in enhancing supply chain resiliency. The current experiments also show that learning external information, network structure, and operational time series together can convert such a data functionality into actual improvement in predictions of risk and propagation.

The GAI-RTPNet attains a smaller Resilience MAE 0.041 and RMSE 0.061 compared to any baseline. The outcome is predictable by the multitask design where graph-level states are limited by the combination of node risk and propagation relations. In the early-warning problem, the average warning lead time is 4.1 d and the values of FAR and MAR are 0.072 and 0.083 respectively. EWT is raised by 1.2 d when compared with PatchTST whereas FAR and MAR decrease by 0.101 and 0.126 respectively to 0.072 and 0.083. Earlier warning activation does not lead to a similar increase in false alarm. The model can respond with an Attention or Warning alert when the node risks have not yet reached peak level and propagation-edge weights keep increasing and system resiliency starts to reduce by using multisource evidence. To make supply chain resilience choices, it is typically necessary to preallocate preparatory resources to anticipate a disruption and postallocate recovery resources to recover after the disruption [42,43]. A greater warning lead time with regulated false alarms has obvious decision-making value.

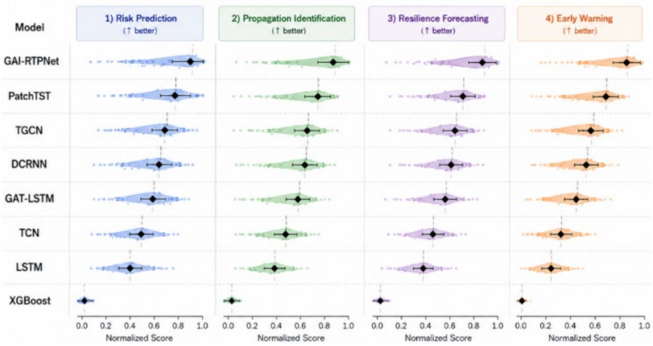

Figure 7. Task-Wise Normalized Performance Profiles of Competing Models

The normalized scores show comparisons between the baseline models and GAI-RTPNet in various task domains Figure 7. In order to remove metric scale differences and optimization direction, all metrics are initially oriented by whether higher or lower values are more desirable. Then, min-max normalization is done per metric, which maps four-task domains of Risk Prediction, Propagation Identification, Resilience Forecasting, and Early Warning into the range [0,1]. The highest aggregated normalized score is achieved by GAI-RTPNet in all four domains as it implies that the gains in performance are not limited to a particular task. Both PatchTST and TGCN form the third tier overall. DCRNN and GAT-LSTM continue to be competitive in propagation identification and resilience prediction, but TCN, LSTM and XGBoost decrease significantly more on tasks related to network structures. It can be observed that this visualization is similar to the raw metrics in Tables 4 and 5 and suggests that the benefit provided by GAI-RTPNet comes through the systematic combination of exogenous information with relational propagation and temporal dependencies and not through the optimization of a specific evaluation metric locally.

4.3. Disturbance Propagation, Robustness, and Generalization Analysis

Disturbance experiments distinguish between observational data and controlled experimental situations. Low, Moderate, and Severe public opinion shocks utilize the real daily GDELT sequence as the baseline and only change the level of exogenous shocks that occur at the node-disturbance-representation layer. Scenarios with upstream availability failure, demand surge, logistics delay, and multi-node cascading disruption start with real SupplyGraph trajectories and then have controlled changes in the relevant operational variables. The predictive ability is therefore tested on real time structures using public data and the model response to different disturbances and their severity are tested using controlled injections. Synthetic scenarios are not viewed as historical events, but as scenarios under which mechanisms can be validated and tested on their robustness. An area where the combination of AI and blockchain has been considered to manage external information on a trusted basis, especially in regard to the truthfulness and traceability of supply chain data [44]. The current model enforces these constraints of trust by gating sources and using early warning based on uncertainty. The scenario-wise prediction and propagation results are reported in Table 6, and the corresponding resilience and early-warning results are reported in Table 7.

Table 6. Prediction and propagation performance under different disturbance and disruption scenarios

Scenario	Level	Risk MAE↓	F1↑	AUC↑	Prop. Acc.↑	Peak Risk
GAI public-opinion shock	Low	0.044	0.919	0.969	0.881	0.43
GAI public-opinion shock	Moderate	0.049	0.907	0.961	0.869	0.61
GAI public-opinion shock	Severe	0.057	0.891	0.949	0.851	0.78
Upstream availability failure	Moderate	0.052	0.901	0.956	0.858	0.69
Demand surge	Moderate	0.050	0.904	0.959	0.861	0.66
Logistics delay	Moderate	0.055	0.894	0.952	0.854	0.72
Multi-node cascading disruption	Severe	0.063	0.876	0.938	0.832	0.84
5% Gaussian noise	—	0.052	0.899	0.956	0.856	0.64
10% Gaussian noise	—	0.058	0.885	0.946	0.842	0.66
10% random missing	—	0.054	0.893	0.951	0.850	0.65
20% random missing	—	0.061	0.878	0.940	0.835	0.68
Cross-scenario testing	Unseen	0.066	0.869	0.932	0.821	0.75

Table 7. Resilience and early-warning performance under different disturbance and disruption scenarios

Scenario	Level	Affected Nodes	Res. Drop↓	EWT (d)↑	FAR↓	MAR↓
GAI public-opinion shock	Low	5	0.07	4.7	0.058	0.071
GAI public-opinion shock	Moderate	10	0.14	4.2	0.069	0.083
GAI public-opinion shock	Severe	17	0.23	3.6	0.083	0.102
Upstream availability failure	Moderate	12	0.18	4.0	0.074	0.094
Demand surge	Moderate	11	0.16	4.3	0.071	0.090
Logistics delay	Moderate	14	0.20	3.8	0.079	0.099
Multi-node cascading disruption	Severe	23	0.31	3.1	0.096	0.124
5% Gaussian noise	—	11	0.15	4.0	0.076	0.092
10% Gaussian noise	—	12	0.17	3.7	0.086	0.108
10% random missing	—	12	0.16	3.9	0.081	0.101
20% random missing	—	13	0.19	3.4	0.092	0.117
Cross-scenario testing	Unseen	16	0.24	3.0	0.101	0.132

Tables 6 and 7 indicate that the increase in the intensity of the public-opinion disturbance Low-Moderate-Severe causes the increase in the Peak Risk 0.43-0.61-0.78, the increase in the number of Affected Nodes 5-10-17 and the increase in the Resilience Drop 0.07-0.14-0.23. These findings show constant risk escalation and weakening resiliency. At the same time, F1 drops by 0.919 to 0.907 and 0.891, and Propagation Accuracy declines by 0.881 to 0.869 and 0.851. The high-intensity shocks do not only expand the risk range but they also make it harder to detect the propagation channels. EWT reduces to 4.2 and 3.6 d as a result of which serious disturbances lead to a faster growth of risks and reduce the possible reaction time. This monotonic trend is in line with the dynamic-edge and propagation-gate design. With the strengthening of the exogenous shock, the larger number of candidate relations are given higher effective propagation weights, pushing risk out through the local nodes to the rest of the network.

Operational-disturbance scenarios also confirm the sensitiveness of the model to various mechanisms. The most difficult controlled scenario is multi-node cascading disruption with a Risk MAE of 0.063, a Propagation Accuracy of 0.832, a Peak Risk of 0.84, 23 Affected Nodes, and a Resilience Drop of 0.31. However, the average warning lead time is 3.1d. Experiments on noise and missingness assess the reliability of inputs. F1 scores are 0.899 and 0.885 with 5% and 10% Gaussian noise respectively, and 0.893 and 0.878 with 10% and 20% random missingness. Cross-scenario testing models a combination of disturbances not seen in training with a Risk MAE of 0.066, a Propagation Accuracy of 0.821, and an EWT of 3.0 d. Even with distribution shift, the model still has practicable propagation-identification and early-warning properties. Instantaneous visibility and end-to-end connectivity are taken to be the key elements of fast response in the digital supply chain [45]. These findings suggest that dynamic graphs and multitask constraints may maintain some level of state identifiability even with incomplete information.

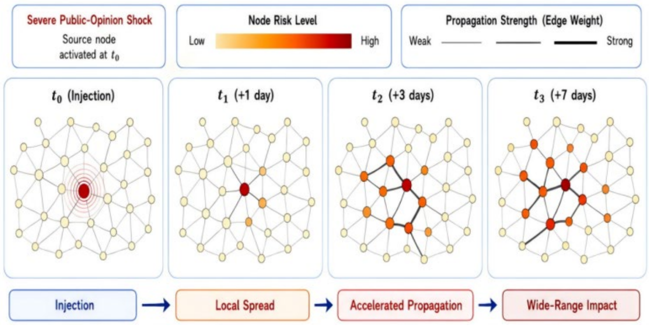

Figure 8. Temporal Risk Propagation under a Controlled Severe Public-Opinion Shock

Figure 8 illustrates the time-based spread of supply-network risks during a severe controlled public-opinion shock. Light to dark node colors represent the increasing risk level and the edge width is used to indicate the dynamic propagation strength caused by the disturbances. When the disturbance is introduced into the source node at t0, risk is localized among other nodes that are locally connected. At t1 and t2, the edge weights as well as the propagation gates of some of the most important connections increase simultaneously leading to an increase of high-risk nodes which are no longer isolated events but rather form local clusters. By t3 this high-risk area continues to grow and even becomes a more continuous propagating cluster. This development supports the numbers in Tables 6 and 7: when in the Severe case, the Peak Risk equals 0.78, the number of Affected Nodes is 17, and the number of Resilience Drop is 0.23. The dynamic growth of the propagation range is not induced by a static topology but happens together with the rise in dynamic edge weights, which can be used to demonstrate the direct link between exogenous public-opinion shocks and useful propagation relations.

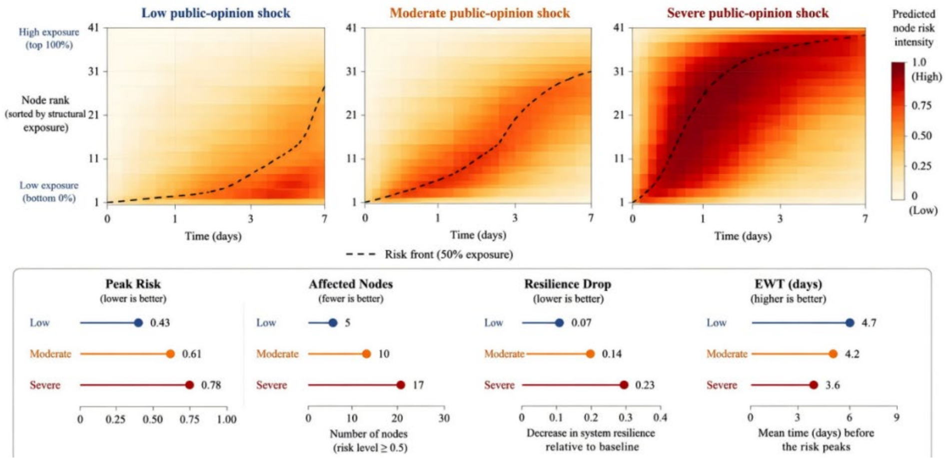

Figure 9. Spatiotemporal Risk Propagation under Controlled Public-Opinion Disturbances

The comparison of risk development between Low, Moderate, and Severe controlled public-opinion disturbances is shown in Figure 9. With the increase in the intensity of the disturbance, high-risk areas expand faster based on the structural exposure of the nodes, and the risk area expands over time to cover more nodes. These aggregated metrics indicate that Peak Risk rises 0.43-0.78, Affected Nodes rise 5-17 and Resilience Drop rises 0.07-0.23, whereas EWT declines 4.7-3.6 days. The figure illustrates that the severity level is rising and therefore amplifies the local risk and network diffusion and narrows the window to react proactively by comparing the same situation in terms of risk intensity, range of propagation, loss of resilience and amount of time available to respond. Multisource conditions and disturbance experiments in virtual environment and scenario modeling in digital supply chains can be organized [46]. The controlled-scenario analysis in this paper will follow this logic, but all performance measures will be obtained via common data partitioning and model-inference methods employed in this article.

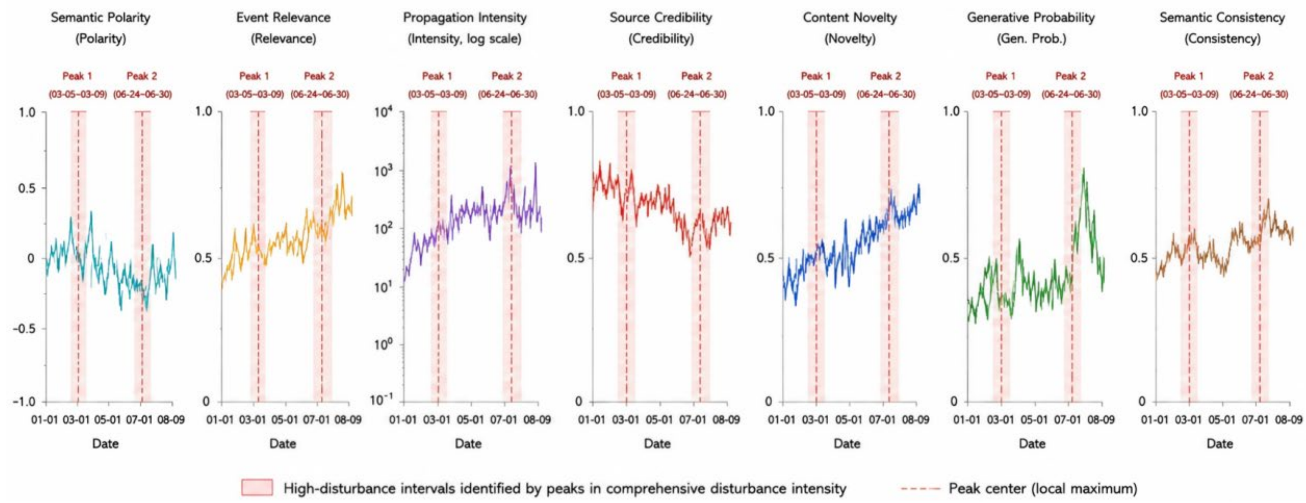

Figure 10. Temporal Dynamics of Generative-AI Disturbance Features over the Study Period

Figure 10 shows the time changes of seven-dimensional generative-AI-driven public-opinion disturbance characteristics during the course of the study. In the vicinity of the chosen intervals of high disturbances, the intensity of propagation and the relevance of the events, the probability of AI generation, credibility of sources and other semantic variables demonstrate significant variability. This created disturbance representation can thus be used to effectively represent the changing dynamics that come about due to external information shocks.

4.4. Analysis Ablation and Interpretability Analysis

The ablation experiments each eliminate the GAI Disturbance Encoder, Dynamic Edge Attention, Temporal Transformer, RiskResilience Joint Learning and Uncertainty Calibration modules in order to determine the marginal contributions of these models to the various tasks. The evaluation focuses on Risk MAE, RMSE, F1, AUC, Propagation Accuracy, Resilience MAE, EWT, FAR, and MAR. Thus, it is able to reflect the shift in point-prediction performance, propagation identification, resilience estimation, and reliability of warnings. In summary, the modules have obvious task-related effects. The dynamic-edge mechanism has the most substantial impact on the identification of propagation, the temporal modeling has the largest effect on the medium and long term risk and warning lead time, the joint multitask learning has more restrictive effect on the system resilience, and uncertainty calibration mostly reduces false alarms. The complete quantitative ablation results are listed in Table 8. In the variant 'w/o GAI Disturbance Encoder,' all GDELT/HC3-derived disturbance inputs are removed and the model retains SupplyGraph operational states and topology only. This variant therefore isolates the aggregate contribution of the external-information pathway. It does not isolate Generative Probability g alone; accordingly, the present evidence supports the combined disturbance representation rather than an independent causal effect of g.

Table 8. Ablation study of GAI-RTPNet

Variant	Risk MAE↓	RMSE↓	F1↑	AUC↑	Prop. Acc.↑	Res. MAE↓	EWT (d)↑	FAR↓	MAR↓
Full GAI-RTPNet	0.048	0.071	0.906	0.963	0.867	0.041	4.1	0.072	0.083
w/o GAI Disturbance Encoder	0.056	0.081	0.887	0.950	0.842	0.049	3.2	0.081	0.105
w/o Dynamic Edge Attention	0.061	0.087	0.879	0.946	0.801	0.053	3.4	0.085	0.112
w/o Temporal Transformer	0.064	0.091	0.872	0.941	0.824	0.056	2.8	0.089	0.119
w/o Risk–Resilience Joint Learning	0.057	0.083	0.891	0.952	0.846	0.064	3.6	0.078	0.103
w/o Uncertainty Calibration	0.049	0.073	0.901	0.957	0.862	0.043	3.7	0.119	0.091

Once Dynamic Edge Attention is eliminated, Propagation Accuracy drops from 0.867 to 0.801, which is an absolute change of 0.066 and the greatest degradation of propagation behavior of all ablation variants. The fixed or weakly dynamic relationships also adversely affect the estimation of risks at nodes, as indicated by a rise in Risk MAE to 0.061.

The elimination of the Temporal Transformer also leads to the increase of Risk MAE to 0.064 and the decrease of EWT to 2.8 d, which is one of the most significant degradations of risk error and warning lead time. The cross-temporal dependencies are hence imperative in defining the delay in transmission and medium to long term trends. Upon deletion of the Risk-Resilience Joint Learning, Resilience MAE rises 0.041 to 0.064 whereas F1 is equal to 0.891. The supervision of Graph-level resilience thus mainly restricts the development of system functionality rather than simply replicating the edge-level classification messages.

The impact of Uncertainty Calibration on point-prediction errors is relatively low. Once removed, Risk MAE varies by 0.048 to 0.049 and Propagation Accuracy moves by 0.867 to 0.862. But, FAR increases from 0.072 to 0.119 (65.3%), while MAR increases from 0.083 to 0.091 (9.6%). The calibration module is thus primarily enhancing the match between risk outputs and confidence estimates so as to reduce overactive warnings in cases of high uncertainty. Additional generative-content detectors may also be susceptible to poor generalization of models and data domains. More recent works have thus highlighted that the detection processes need to take into account the performance, generalization and linguistic variance simultaneously [47]. That is why GAI-RTPNet cannot directly use the AI-generation probability as the warning outcome. Rather, the end warning is based on the combination of event relevance, source credibility, dynamic propagation evidence and the two kinds of uncertainty.

The model stores dynamism of edge weights, propagation gates, event relevance, warning severity and uncertainty across any time step. These intermediary variables help identify critical supply/manufacturing nodes, high-risk propagation edges, and high-risk periods. Such variables also enable explaining the origins of risk in an event timeline, the path of risk propagation, its amplification time, and why the warning is triggered. The anonymized SupplyGraph data node level explanations only involve the product-node identifiers and structural positions and no actual recovery or inference of the true firm identities are made.

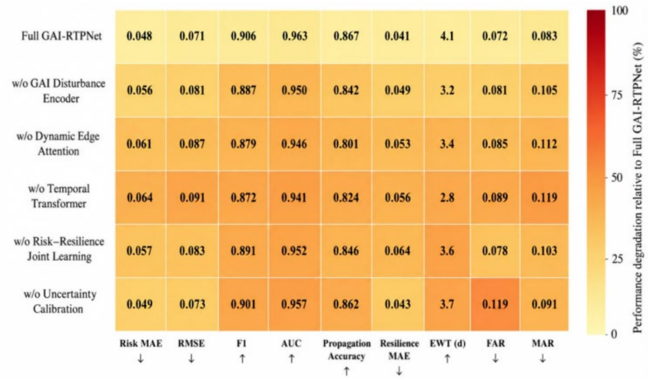

Figure 11. Ablation Study Heatmap of GAI-RTPNet

The results of the ablation of the GAI-RTPNet modules are presented in Figure 11 as a heatmap. All cells have kept the same evaluation measure, but the color denotes the level of performance decrease compared to the full model. Complete GAI-RTPNet has values of 0.048, 0.071, 0.906, 0.963, 0.867, 0.041, 4.1 d, 0.072, and 0.083 on Risk MAE, RMSE, F1, AUC, Propagation Accuracy, Resilience MAE, EWT, FAR, and MAR, respectively. Propagation Accuracy is lowered to 0.801 with elimination of Dynamic Edge Attention, Risk MAE is increased to 0.064 with elimination of the Temporal Transformer and EWT is lowered to 2.8 d with elimination of the Temporal Transformer, Resilience MAE is increased to 0.064 with elimination of Risk-Resilience Joint Learning, and FAR is raised to 0.119 with elimination of Uncertainty Calibration. The main benefit of every core module is thus aligned to its design purpose, and there is not one module that captures all the performance improvements.

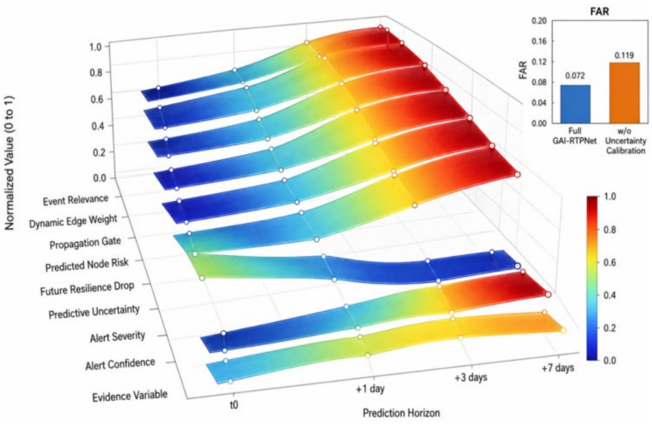

Figure 12. Three-Dimensional Evidence Landscape for Interpretable Warning Formation in GAI-RTPNet

As shown in Figure 12, this is how warning evidence develops over time during prediction horizons in a typical Severe public-opinion-shock situation. The horizontal dimensions are represented by t0, +1 d, +3 d, and +7 d, and the vertical dimensions represent normalized measurements such as Event Relevance, Dynamic Edge Weight, Propagation Gate, Predicted Node Risk, Future Resilience Drop, Predictive Uncertainty, Alert Severity and Alert Confidence. With increasing predictive evidence, the event relevance, dynamic edge weight, and propagation gates tend to increase and the predicted node risk and future resilience loss increase. External disturbances to information are thus gradually converted into network-based evidence of risk propagation. At the same time, predictive uncertainty decreases by a larger amount, which provides stronger evidence in favor of warning severity and confidence. The comparison of calibrations also indicates that the full GAI-RTPNet has an FAR of 0.072, which increases to 0.119 when Uncertainty Calibration is excluded, agreeing with the findings of the ablation experiments in Table 8. Evidence landscape creates a continuous chain of explanations of the external-information relevance, relation activation, risk propagation, system-resilience loss and prediction trustworthiness and thus increases the traceability of warnings.

5. Conclusions

Since it is difficult to quantify external information disturbances in manufacturing supply chains, there is a dynamic change in risk-propagation relationships, and the credibility of early warnings is low during the high rate of generation and spread of generative-AI content, this paper introduces the Generative-AI-Aware Risk Transmission and Resilience Prediction Network, GAI-RTPNet. The model converts the disturbances in public opinion which are independent textual characteristics into exogenous factors influencing the evolution of the supply network state. Incorporating dynamic-graph propagation, multi-horizon predictions and dual-uncertainty-aware mechanisms, it offers a common framework of identifying risks, inferring propagation, assessing resilience and reliable early warnings. The most important conclusions are as follows.

(i) GAI-RTPNet is a good representation of how generative-AI-based public opinion unrest can influence the spread of risk across supply chains. Generative-AI Disturbance Encoder incorporates semantic polarity, event relevancy, propagation strength, source credibility, content novelty, AI-generation probability, and semantic consistency in order to build exogenous disturbances at the level of nodes. Risk-propagation paths are better characterized, and their temporal changes are better depicted through external information shocks participating directly in the evolution of supply relations and node states due to dynamic edge-weight updating and risk-propagation gating.

(ii) The experimental findings confirm the efficacy of GAI-RTPNet in multitask prediction and early warning. The model has a Risk MAE of 0.048, a Propagation Accuracy of 0.867 and a Resilience MAE of 0.041 and a mean warning lead time of 4.1 days. It has FAR and MAR of 0.072 and 0.083, respectively. Stability of the model is achieved even during extreme multi-node cascading failures, random noises, missing data, and cross-scenario experiments. The ablation findings also illustrate that Dynamic Edge Attention, the Temporal Transformer, RiskResilience Joint Learning, and Uncertainty Calibration play different roles in the areas of propagation identification, temporal prediction, resilience modeling, and warning reliability respectively.

The current research has multiple constraints. Since SupplyGraph works with anonymized product nodes, GDELT events and supply-network nodes are poorly aligned mainly on the industry and time dimensions. Causal attribution at the firm level is thus not possible yet. The network topology, demand elasticity, recovery mechanisms, etc., differ by industry, and cross-industry generalization of the model should be further validated. Incorporation of firm-level multisource data, causal risk propagation, supply chain digital twins and online continual-learning mechanisms can be considered in future research. Reliable warnings might also be associated with more direct response strategies that include procurement reconfiguration, inventory adjustment, logistics rerouting and production recovery, which would enhance the practicality of the model in a real-world manufacturing supply chain. Because validation is limited to anonymized product-level SupplyGraph data, firm-level external validity and transfer across industries with different topology, demand elasticity, and recovery processes remain unresolved and require independent validation. These findings should also be interpreted in light of three inherent data-driven limitations. Weak event-to-node alignment can introduce assignment noise because entity-level ground truth is unavailable; the temporal-order propagation labels are weakly supervised and may contain false-positive or false-negative edges when co-movement is not genuine transmission; and model hyperparameters and warning thresholds are calibrated to the current data regime, so their optimal values may shift with graph scale, disturbance prevalence, operating costs, data distributions, or intervention policies. These sources of uncertainty are mitigated by topology constraints, multitask supervision, uncertainty calibration, and validation-set tuning, but they cannot be fully eliminated in the present public-data setting.

Acknowledgements

This work was support by Science Research Project of Hebei Education Department (Grant No. BJ2025339); 2023 Young Fund Project for Humanities and Social Sciences Research of the Ministry of Education of China (Grant No. 23YJCZH131).

References

[1] Jackson I, Ivanov D, Dolgui A, et al. Generative artificial intelligence in supply chain and operations management: A capability-based framework for analysis and implementation[J]. International Journal of Production Research, 2024, 62(17): 6120-6145.

[2] Boone T, Fahimnia B, Ganeshan R, et al. Generative AI: Opportunities, challenges, and research directions for supply chain resilience[J]. Transportation Research Part E: Logistics and Transportation Review, 2025, 199: 104135.

[3] Smyth C, Dennehy D, Fosso Wamba S, et al. Artificial intelligence and prescriptive analytics for supply chain resilience: A systematic literature review and research agenda[J]. International Journal of Production Research, 2024, 62(23): 8537-8561.

[4] Ivanov D. Transformation of supply chain resilience research through the COVID-19 pandemic[J]. International Journal of Production Research, 2024, 62(23): 8217-8238.

[5] Li X, Krivtsov V, Pan C, et al. End-to-end supply chain resilience management using deep learning, survival analysis, and explainable artificial intelligence[J]. International Journal of Production Research, 2025, 63(3): 1174-1202.
[6] Modgil S, Singh R K, Hannibal C. Artificial intelligence for supply chain resilience: Learning from COVID-19[J]. The International Journal of Logistics Management, 2022, 33(4): 1246-1268.
[7] Belhadi A, Mani V, Kamble S S, et al. Artificial intelligence-driven innovation for enhancing supply chain resilience and performance under the effect of supply chain dynamism: An empirical investigation[J]. Annals of Operations Research, 2024, 333(2-3): 627-652.
[8] Dubey R, Bryde D J, Dwivedi Y K, et al. Impact of artificial intelligence-driven big data analytics culture on agility and resilience in humanitarian supply chain: A practice-based view[J]. International Journal of Production Economics, 2022, 250: 108618.
[9] Deiva Ganesh A, Kalpana P. Future of artificial intelligence and its influence on supply chain risk management—A systematic review[J]. Computers & Industrial Engineering, 2022, 169: 108206.
[10] Deiva Ganesh A, Kalpana P. Supply chain risk identification: A real-time data-mining approach[J]. Industrial Management & Data Systems, 2022, 122(5): 1333-1354.
[11] Ashraf M, Eltawil A, Ali I. Disruption detection for a cognitive digital supply chain twin using hybrid deep learning[J]. Operational Research, 2024, 24: 23.
[12] Ivanov D. Intelligent digital twin (iDT) for supply chain stress-testing, resilience, and viability[J]. International Journal of Production Economics, 2023, 263: 108938.
[13] Jackson I, Saenz M J, Ivanov D. From natural language to simulations: Applying AI to automate simulation modelling of logistics systems[J]. International Journal of Production Research, 2024, 62(4): 1434-1457.
[14] Akhtar P, Ghouri A M, Khan H U R, et al. Detecting fake news and disinformation using artificial intelligence and machine learning to avoid supply chain disruptions[J]. Annals of Operations Research, 2023, 327(2): 633-657.
[15] Chatterjee S, Chaudhuri R, Vrontis D. Role of fake news and misinformation in supply chain disruption: Impact of technology competency as moderator[J]. Annals of Operations Research, 2023, 327(2): 659-682.
[16] Petratos P N, Faccia A. Fake news, misinformation, disinformation and supply chain risks and disruptions: Risk management and resilience using blockchain[J]. Annals of Operations Research, 2023, 327(2): 735-762.
[17] Park S, Nan X. Generative AI and misinformation: A scoping review of the role of generative AI in the generation, detection, mitigation, and impact of misinformation[J]. AI & Society, 2026, 41: 1501-1515.
[18] Li L, Liu Y, Jin Y, et al. Generative AI-enabled supply chain management: The critical role of coordination and dynamism[J]. International Journal of Production Economics, 2024, 277: 109388.
[19] Gelastopoulos G, Keramydas C. A systematic review of text mining analytics for supply chain risk management using online data[J]. Supply Chain Analytics, 2025, 12: 100167.
[20] Statsenko L, Jayasinghe R S, Soosay C. Supply network resilience capabilities: A social-ecological perspective[J]. Supply Chain Management: An International Journal, 2024, 29(1): 1-26.
[21] Jauhar S K, Harinath S, Krishnaswamy V, et al. Explainable artificial intelligence to improve the resilience of perishable product supply chains by leveraging customer characteristics[J]. Annals of Operations Research, 2025, 354: 103-142.
[22] Habibi F, Chakrabortty R K, Abbasi A. Evaluating supply chain network resilience considering disruption propagation[J]. Computers & Industrial Engineering, 2023, 183: 109531.
[23] Habibi F, Chakrabortty R K, Abbasi A, et al. Investigating disruption propagation and resilience of supply chain networks: Interplay of tiers and connections[J]. International Journal of Production Research, 2025, 63(17): 6229-6251.
[24] Zhao P, Li Z, Han X, et al. Supply chain network resilience by considering disruption propagation: Topological and operational perspectives[J]. IEEE Systems Journal, 2022, 16(4): 5305-5316.
[25] Chen D, Sun D, Yin Y, et al. The resilience of logistics network against node failures[J]. International Journal of Production Economics, 2022, 244: 108373.
[26] Cao Y, Zhu X, Yan H. Data-driven Wasserstein distributionally robust mitigation and recovery against random supply chain disruption[J]. Transportation Research Part E: Logistics and Transportation Review, 2022, 163: 102751.
[27] Scheibe K P, Blackhurst J. Supply chain disruption propagation: A systemic risk and normal accident theory perspective[J]. International Journal of Production Research, 2018, 56(1-2): 43-59.
[28] Sokolov B, Ivanov D, Dolgui A, Pavlov A. Structural quantification of the ripple effect in the supply chain[J]. International Journal of Production Research, 2016, 54(1): 152-169.
[29] Kosasih E E, Margaroli F, Gelli S, et al. Towards knowledge graph reasoning for supply chain risk management using graph neural networks[J]. International Journal of Production Research, 2024, 62(15): 5596-5612.
[30] Yang S, Ogawa Y, Ikeuchi K, et al. Post-hazard supply chain disruption: Predicting firm-level sales using graph neural network[J]. International Journal of Disaster Risk Reduction, 2024, 110: 104664.
[31] Xu D, Ruan C, Korpeoglu E, Kumar S, Achan K. Inductive representation learning on temporal graphs[C]//International Conference on Learning Representations. 2020.
[32] Sankar A, Wu Y, Gou L, Zhang W, Yang H. DySAT: Deep neural representation learning on dynamic graphs via self-attention networks[C]//Proceedings of the 13th ACM International Conference on Web Search and Data Mining. 2020: 519-527.
[33] Cheng Y, Elsayed E A, Huang Z. Systems resilience assessments: A review, framework and metrics[J]. International Journal of Production Research, 2022, 60(2): 595-622.
[34] Tyralis H, Papacharalampous G. A review of predictive uncertainty estimation with machine learning[J]. Artificial Intelligence Review, 2024, 57: 94.
[35] Sluijterman L, Cator E, Heskes T. How to evaluate uncertainty estimates in machine learning for regression?[J]. Neural Networks, 2024, 173: 106203.
[36] Ivanov D, Dolgui A, Blackhurst J V, et al. Toward supply chain viability theory: From lessons learned through COVID-19 pandemic to viable ecosystems[J]. International Journal of Production Research, 2023, 61(8): 2402-2415.
[37] Wang C, Guo Z, Shi F. Tipping Point or False Alarm? An interpretable machine learning framework for early warning

of supply chain disruptions under multi-source uncertainty[J]. Applied Sciences, 2026, 16(13): 6457.
[38] Khaleghi M, Sheykhivand S, Khaleghi N, et al. An intelligent multi-task supply chain model based on bio-inspired networks[J]. Biomimetics, 2026, 11(2): 123.
[39] Ardeshirifar R. Comparing hand-crafted and deep learning approaches for detecting AI-generated text: Performance, generalization, and linguistic insights[J]. AI and Ethics, 2025, 5: 4197-4209.
[40] Hong D, Fu Z, Zhang X, et al. Research on the development and application of the GDELT Event Database[J]. Data, 2025, 10(10): 158.
[41] Kosasih E E, Brintrup A. A machine learning approach for predicting hidden links in supply chain with graph neural networks[J]. International Journal of Production Research, 2022, 60(17): 5380-5393.
[42] Zamani E D, Smyth C, Gupta S, et al. Artificial intelligence and big data analytics for supply chain resilience: A systematic literature review[J]. Annals of Operations Research, 2023, 327(2): 605-632.
[43] Aldrighetti R, Battini D, Ivanov D. Efficient resilience portfolio design in the supply chain with consideration of preparedness and recovery investments[J]. Omega, 2023, 117: 102841.
[44] Charles V, Emrouznejad A, Gherman T. A critical analysis of the integration of blockchain and artificial intelligence for supply chain[J]. Annals of Operations Research, 2023, 327(1): 7-47.
[45] Dolgui A, Ivanov D. 5G in digital supply chain and operations management: Fostering flexibility, end-to-end connectivity and real-time visibility through internet-of-everything[J]. International Journal of Production Research, 2022, 60(2): 442-451.
[46] Dolgui A, Ivanov D. Metaverse supply chain and operations management[J]. International Journal of Production Research, 2023, 61(23): 8179-8191.
[47] Alharthi R, Ojo S, Nathaniel T I, et al. Responsible detection and mitigation of AI-generated text using hybrid neural networks and feature fusion: Toward trustworthy content management in the era of large language models[J]. International Journal of Computational Intelligence Systems, 2025, 18: 274.