Beyond Graph Priors: A Co-Evolving Framework Under Uncertainty for Enterprise Resilience Assessment

Yanzhe Xie[1], Li Huang[1], Qiang Gao[1*], Xueqin Chen[2], Fan Zhou[3], Kunpeng Zhang[4]

[1]Southwestern University of Finance and Economics, Chengdu, China
[2]Sichuan University, Chengdu, China
[3]University of Electronic Science and Technology of China, Chengdu, China
[4]University of Maryland, College Park, USA
223081200056@smail.swufe.edu.cn, lihuang@swufe.edu.cn, qianggao@swufe.edu.cn, nedchen0728@gmail.com, fan.zhou@uestc.edu.cn, kpzhang@umd.edu

Abstract

Assessing enterprise resilience under uncertainty necessitates capturing both intrinsic attributes and evolving inter-enterprise dependencies. However, real-world enterprise systems pose substantial structural challenges: redundant or loosely correlated links can trigger spurious relational inferences, while missing or latent dependencies often hinder the propagation of informative signals. Moreover, most existing approaches adopt static graph priors or decouple structural refinement from semantic learning, lacking a co-evolutionary paradigm that allows structure and representation to inform one another. We propose **CFU**, a novel **C**o-evolving **F**ramework under **U**ncertainty, which reconceptualizes graph structure as a dynamic and learnable component evolving alongside node semantics. Specifically, CFU begins with a structure-aware contrastive pretraining phase to distill latent relational semantics without supervision. It then performs bidirectional structural refinement, filtering structurally redundant edges through semantic agreement scoring, and uncovering temporally contingent, task-relevant dependencies via similarity-guided inference. These operations are integrated through a dynamic fusion procedure that continuously aligns the evolving topology with the resilience objective. By embedding structural adaptation within the learning loop, CFU enables context-aware resilience assessment across incomplete, ambiguous, and structurally volatile enterprise environments. Ultimately, extensive experiments conducted on real-world datasets demonstrate its superior performance across diverse evaluation scenarios.

Code — https://github.com/xieyanzhe/CFU

Introduction

In an era of growing uncertainty and systemic interdependence, assessing the resilience of enterprises has become essential for maintaining economic continuity and macro-level stability (Li et al. 2024; Zhang et al. 2022). Enterprise resilience refers to a company's ability to sustain healthy operations without triggering adverse events such as risk warnings, trading suspensions, disciplinary disclosures, or delisting actions. Unlike consumer-level assessment tasks (Liang et al. 2021), enterprise-level evaluation faces the signals of instability that are often sparse, incomplete, and distributed across multiple dimensions. Crucially, resilience is rarely determined by enterprise-level indicators alone but often emerges from the structural dependencies that connect enterprises through supply chains, financial exposure, ownership networks, or regional co-movements (Bi et al. 2022).

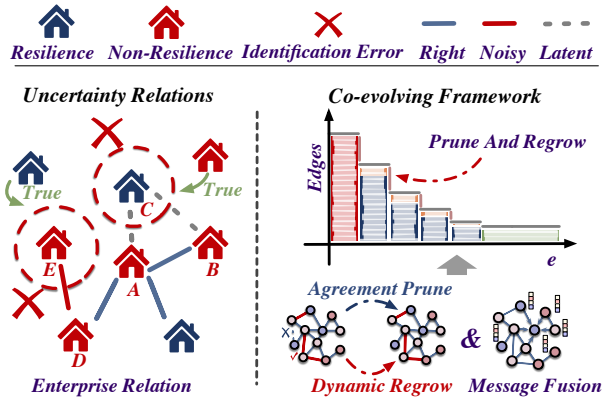

Figure 1: Insights of our solution design.

Early efforts in enterprise resilience assessment predominantly employed probabilistic scoring methods and machine learning algorithms based on tabular features of enterprise-level attributes, including logistic regression (LR) (LaValley 2008), gradient boosting machines (Chen and Guestrin 2016),tree-based methods (Sun et al. 2018), and other techniques (Danenas and Garsva 2015; Yeh, Chi, and Lin 2014; Kou et al. 2021). In parallel, deep neural network architectures based on multilayer perceptrons (MLP) (Taud and Mas 2017) have been employed to capture nonlinear patterns embedded within isolated feature vectors. These practices typically treated each enterprise as an independent observation, abstracting away any structural interdependence within the economic system. As shown in Figure 1, resilience is often influenced by internal attributes or connections within investment-linked networks, where associated enterprises may exhibit similar operational patterns (Bi et al. 2022). For instance, enterprise C is misclassified as resilient be-

*Corresponding Author (qianggao@swufe.edu.cn)

cause its latent relations with non-resilient enterprises A and B remain undiscovered, while enterprise E is misclassified due to false risk propagation induced by redundant edges. Hence, overlooking such linkages weakens the model's ability to reflect external dependencies, limiting the accuracy of resilience evaluation (Yang et al. 2021).

In contrast to instance-centric techniques, recent studies have adopted graph neural network (GNN)-based frameworks that incorporate inter-enterprise relations as structured priors. Static GNNs mark the initial progress in this direction, enriching enterprise node representations via localized message propagation via representative architectures, e.g., GCN (Kipf and Welling 2016), GAT (Veličković et al. 2017), DAGNN (Liu, Gao, and Ji 2020), and others (Cheng et al. 2019; Zheng et al. 2021). By utilizing structural signals derived from co-investment, governance overlap, or transactional ties, these methods have achieved notable gains in enterprise-level assessment (Wei et al. 2024). However, the assumption of stable and fully observed relationships limits their capacity to address real-world enterprise systems, where interactions are often incomplete, time-sensitive, and influenced by external factors (Zhang et al. 2024).

To capture structural evolution, dynamic graph frameworks (Huang et al. 2025; Han et al. 2025) have been proposed to incorporate temporal variations in both node attributes and connectivity patterns. Heterogeneous graphs have been explored for modeling multi-relational data (Zhu et al. 2020; Zhao et al. 2021; Wang et al. 2022), but these methods rely on complex paradigms and task-specific inputs, limiting generality; by treating structural dynamics as external and decoupling structure refinement from representation learning, they amplify noise and miss latent dependencies, degrading performance. Concerned with the limitations of existing studies under enterprise evolution uncertainty, we identify three key challenges in enterprise resilience assessment: ***C1:*** *Noisy relations undermine structural reliability, as outdated or loosely correlated edge links introduce misleading signals.* ***C2:*** *Missing connections hinder structural completeness due to limited data, privacy constraints, or unrecorded dependencies.* ***C3:*** *Fixed graph priors restrict adaptability, as most methods separate structure and representation learning, failing to capture the co-evolving nature of dynamic enterprise systems.* Correspondingly, we propose a novel **C**o-evolving **F**ramework under **U**ncertainty (**CFU**) for enterprise resilience assessment. As shown in Figure 1, CFU relaxes the fixed-graph prior by jointly evolving topology and node representations. It begins with structure-aware contrastive pretraining to capture latent relations, then refines the graph via latent agreement–based pruning and node-similarity–based augmentation. The resulting structures are fused into a unified, dynamically updated topology. Guided by a resilience objective, CFU enables mutual reinforcement between structure and representation, adapting robustly to noisy and evolving enterprise environments. Our contributions are summarized as follows:

- In this study, we introduce a unified co-evolving solution for enterprise resilience assessment, which jointly optimizes graph structure and node representations under uncertainty. This design breaks the reliance on static graph priors and enables adaptive modeling of evolved enterprise systems.
- We propose an uncertainty-aware structural refinement strategy that adaptively prunes edges with low consistency by leveraging similarity at the embedding level, which effectively diminishes structural redundancy and alleviates misleading relational cues.
- We introduce a learnable structure augmentation strategy to recover latent dependencies and a fusion procedure to integrate pruned and added edges into a unified, evolving topology. This jointly enhances structural completeness and supports robust message propagation.
- In the end, we conduct extensive experiments on real-world enterprise datasets and demonstrate that our approach consistently outperforms state-of-the-art baselines, particularly in scenarios characterized by incomplete data, evolved relations, and structural uncertainty.

Related Work

Attribute-driven Enterprise Assessment. Initial investigations of enterprise resilience typically used machine learning and statistical models that treat each enterprise as an isolated instance (Erdogan 2013; Chen, Chen, and Shi 2020). Logistic regression (LaValley 2008), gradient boosting (Chen and Guestrin 2016), and decision tree methods (Sun et al. 2018) focus on structured attribute data to estimate enterprise stability or default probability. Deep learning methods (Mai et al. 2019; Taud and Mas 2017) have been used to capture nonlinear feature interactions. While they perform well in clean, feature-rich environments, they overlook the structural dependencies among enterprises, which are essential for understanding resilience under uncertainty. As a result, these methods fail to account for how interconnected enterprises may influence each other through financial ties, ownership structures, or collaborative operations.

Learning with Inter-Enterprise Dependencies. More recent work introduces relational structures to capture interdependencies within enterprise ecosystems. Static graph neural networks, e.g., GCN (Kipf and Welling 2016), GAT (Veličković et al. 2017), GIN (Xu et al. 2018), and DAGNN (Liu, Gao, and Ji 2020), enhance node representations by propagating messages across predefined edges, leveraging connections, including co-investment, governance overlap, or supply-chain links (Cheng et al. 2019; Zheng et al. 2021). Although effective, these methods typically rely on fixed graph topologies constructed in advance, which can include noisy or outdated relations that mislead the learning process. To improve flexibility, dynamic graph methods (Huang et al. 2025; Han et al. 2025) allow temporal changes in node attributes and graph structure, enabling better adaptation to evolving systems. Heterogeneous graphs (Wang et al. 2022; Zhao et al. 2021) further integrate multiple types of interactions, enriching the semantic space. However, both directions often treat structural variation as external to the learning objective and require complex relational schema, which limit their scalability in real-world applications where data is sparse, noisy, or partially observed.

Prerequisites

Enterprise Relational Graph. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ denote an enterprise relational graph, where each node $v \in \mathcal{V}$ represents an enterprise, and each edge $(u, v) \in \mathcal{E}$ indicates a real-world interaction between enterprises. The node feature matrix is denoted as $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times d}$, where d is the dimensionality of enterprise attributes. Each edge is associated with a scalar weight $w_{uv} \in \mathbb{R}$, reflecting the strength or reliability of the connection between enterprise u and v. In practice, the observed enterprise graph is often noisy, incomplete, or semantically misaligned with the Enterprise Resilience Assessment (ERA) objective, thereby motivating the need for adaptive structural refinement.

Enterprise Resilience Assessment (ERA). We now formally define the ERA task. Given an enterprise relational graph $\mathcal{G}$ with node features $\mathbf{X}$ and edge weights $\mathcal{W}$ as defined above, the objective is to construct a resilience assessment function $\mathcal{F}_\theta$ that simultaneously produces resilience scores $\hat{\mathbf{Y}}$ and an optimized graph structure $\tilde{\mathcal{G}} = (\mathcal{V}, \tilde{\mathcal{E}}, \tilde{\mathcal{W}})$. Formally, $(\hat{\mathbf{Y}}, \tilde{\mathcal{G}}) = \mathcal{F}_\theta(\mathcal{G}, \mathbf{X})$. Supervision is provided by a partially labeled set $\mathcal{D} = \{(v, y_v)\}$, where $y_v \in \{0, 1\}$ denotes whether enterprise v is identified as non-resilience in the context of assessment.

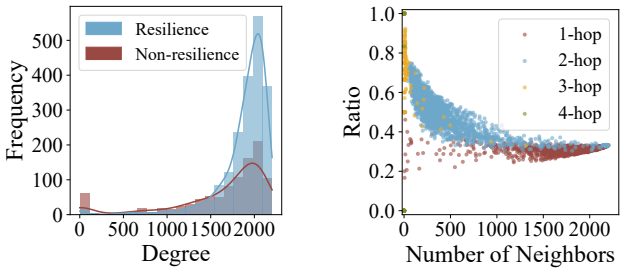

(a) (b)

Figure 2: (a) Degree distribution of enterprise labeled as resilience or non-resilience. (b) Ratio of non-resilience neighbors within multi-hop neighborhoods.

Exploratory Analysis. As shown in Figure 2, the enterprise interactions from the 2023 Shanghai Stock Exchange (SSE) exhibit a highly dense graph structure. Notably, even enterprises with low degrees can have a high proportion of non-resilient neighbors within their two-hop or three-hop neighborhoods. This indicates that resilience-related signals may propagate via indirect or weakly correlated paths, potentially distorting local structures and complicating accurate assessment.

Methodology

We propose a co-evolving framework, namely CFU, for the ERA task that jointly refines graph structure and node representations. As shown in Figure 3, CFU performs contrastive pretraining to capture latent semantics, followed by semantic pruning to remove redundant connections and similarity-driven inference to uncover uncertain relational patterns. These components are integrated into an evolving graph, enabling structure and representation to mutually adapt throughout training, all of which are described below.

Contrastive Representation Pretraining (CRP)

To establish a robust and structure-aware initialization, we first adopt a contrastive pretraining approach that operates independently of labeled supervision. This component is designed to extract intrinsic relational semantics by maximizing the agreement between embeddings of locally connected enterprises. Formally, given an enterprise graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with a node feature matrix $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times d}$, we employ an encoder $\mathcal{J}_\phi$ to compute the initial enterprise representations:

$$\mathbf{H}^{(1)} = \sigma(\hat{\mathbf{A}}\mathbf{X}\mathbf{W}_\phi^{(1)}), \quad \mathbf{H} = \hat{\mathbf{A}}\mathbf{H}^{(1)}\mathbf{W}_\phi^{(2)}, \tag{1}$$

where $\hat{\mathbf{A}} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ is the row-normalized adjacency matrix with self-loops explicitly incorporated, $\mathbf{W}_\phi^{(1)}, \mathbf{W}_\phi^{(2)}$ are learnable weights, and $\sigma(\cdot)$ denotes a nonlinear activation function, instantiated as ReLU. The output matrix $\mathbf{H} \in \mathbb{R}^{|\mathcal{V}| \times h}$ comprises h-dimensional latent embeddings $\mathbf{h}_u$ for each enterprise node $u \in \mathcal{V}$.

To produce contrastive signals, we randomly shuffle the observed edge set $\mathcal{E}$ and select a subset of fixed size $\mathcal{E}_s \subseteq \mathcal{E}$ to form contrastive pairs. For each edge $(u, v) \in \mathcal{E}_s$, node v serves as the positive instance for anchor node u, while all other nodes $v' \in \mathcal{V} \setminus \{v\}$ are treated as implicit negatives. The objective encourages alignment between connected enterprises while suppressing misleading similarities:

$$\mathcal{L}_c = -\sum_{(u,v)\in\mathcal{E}_s} \log \frac{\exp(\mathbf{h}_u^\top \mathbf{h}_v/\tau)}{\sum_{v'\in\mathcal{V}} \exp(\mathbf{h}_u^\top \mathbf{h}_{v'}/\tau)}, \tag{2}$$

here, τ is a temperature hyperparameter, and $\mathbf{h}_u, \mathbf{h}_v$ are enterprise embeddings generated by the encoder $\mathcal{J}_\phi$. After convergence, $\mathcal{J}_\phi$ is frozen and reused in subsequent pruning and augmentation stages as a stable, task-agnostic initializer.

Agreement-Guided Edge Pruning (AEP)

To eliminate structural redundancy and suppress semantically inconsistent connections, we estimate semantic consistency for each observed edge based on latent representations from the frozen encoder $\mathcal{J}_\phi$. Specifically, for each edge $(u, v) \in \mathcal{E}$ with embeddings $\mathbf{h}_u, \mathbf{h}_v \in \mathbb{R}^h$, we compute a normalized agreement score S_{uv}^p, denoted as:

$$S_{uv}^p = \frac{\mathbf{h}_u^\top \mathbf{h}_v - \min\limits_{(u,v)\in\mathcal{E}} (\mathbf{h}_u^\top \mathbf{h}_v)}{\max\limits_{(u,v)\in\mathcal{E}} (\mathbf{h}_u^\top \mathbf{h}_v) - \min\limits_{(u,v)\in\mathcal{E}} (\mathbf{h}_u^\top \mathbf{h}_v) + \epsilon}, \tag{3}$$

where ϵ is a small constant for numerical stability. This min-max normalization facilitates fair comparison across edges and training stages. To gradually filter unreliable connections while preserving essential structure in early training, we introduce a dynamic threshold τ_e that increases linearly with training step e:

$$\tau_e = \tau_i + \left(\frac{e}{T}\right)(\tau_f - \tau_i), \tag{4}$$

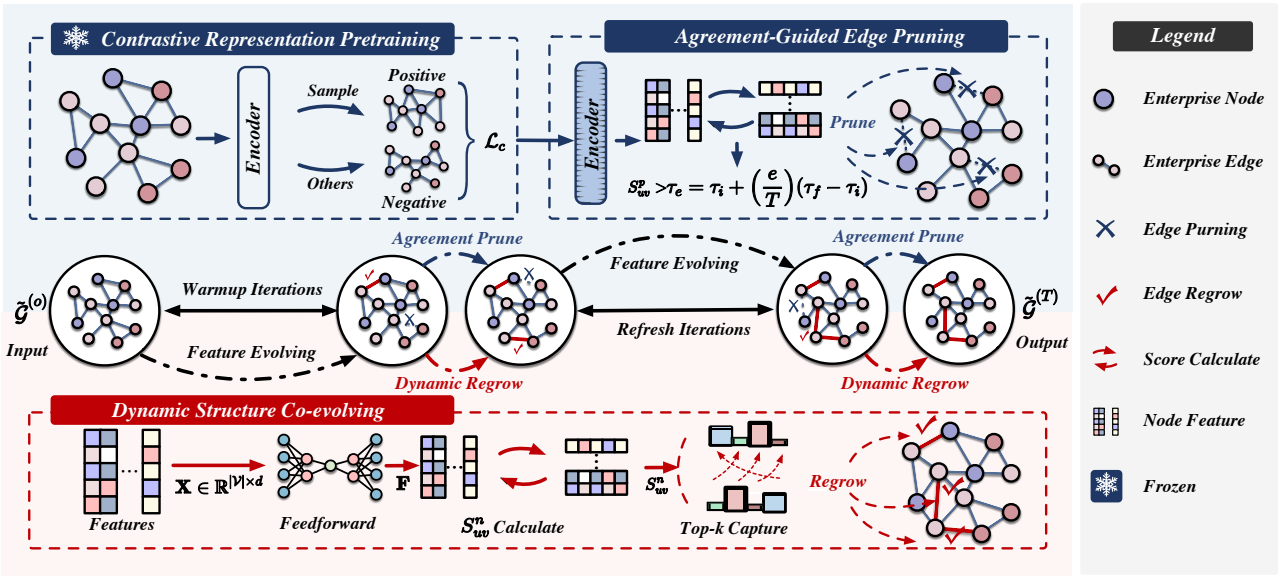

Figure 3: The Co-evolving Framework under Uncertainty (CFU) for addressing the ERA task.

where τ_i and τ_f denote the initial and final thresholds, and T is the predefined saturation iterations. At step e, we retain edges with agreement scores above τ_e:

$$\mathcal{E}_p^{(e)} = \{(u, v) \in \mathcal{E} \mid S_{uv}^p > \tau_e\}, \tag{5}$$

resulting in a pruned subgraph $\mathcal{G}_p = (\mathcal{V}, \mathcal{E}_p)$ that serves as a structurally refined prior with reduced redundancy. This pruning is performed at fixed training iterations, enabling efficient topology adaptation without introducing additional parameters, while maintaining alignment with learned semantic signals. In practice, pruning is activated after a warm-up phase of e_w iterations to allow stable feature formation. Subsequent updates occur every e_r iterations, ensuring that structural refinement progresses in synchrony with representation learning while avoiding excessive fluctuation.

Dynamic Structure Co-evolving (DSC)

To capture latent and evolving dependencies under structural uncertainty, we introduce a learnable structure inference component DSC that dynamically regrows implicit relations among enterprises via semantic cues. All node features are assembled into matrix $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times d}$. These are projected through a two-layer feedforward network:

$$\mathbf{F} = \mathbf{W}_\theta^{(2)} \cdot \sigma(\mathbf{W}_\theta^{(1)} \mathbf{X} + \mathbf{b}_\theta^{(1)}) + \mathbf{b}_\theta^{(2)}, \tag{6}$$

where $\mathbf{W}_\theta^{(1)} \in \mathbb{R}^{d \times h}$ and $\mathbf{W}_\theta^{(2)} \in \mathbb{R}^{h \times d}$ denote learnable weights, and $\sigma(\cdot)$ represents a nonlinear activation function ReLU. The resulting matrix $\mathbf{F} \in \mathbb{R}^{|\mathcal{V}| \times d}$ encodes enriched semantic embeddings that comprehensively preserve higher-order feature interactions, thereby facilitating the reliable discovery of latent relational patterns across enterprises. We then compute a pairwise similarity matrix $S^n \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ to robustly quantify potential inter-enterprise affinities:

$$S_{uv}^n = \frac{\mathbf{F}_u^\top \mathbf{F}_v}{|\mathbf{F}_u|_2 \cdot |\mathbf{F}_v|_2}. \tag{7}$$

Let $\mathbf{S}_u^n = [S_{uv_1}^n, S_{uv_2}^n, \ldots, S_{uv_{|\mathcal{V}|}}^n] \in \mathbb{R}^{|\mathcal{V}|}$ denote the similarity profile of node u, measuring its cosine affinity comprehensively to other all nodes. We then induce directed connections from u to its top-k most semantically aligned peers, thereby forming a sparse, data-driven relational structure via this similarity vector. Finally, we construct a directed edge from u to its top-k most similar enterprises:

$$\mathcal{E}_d = \{(u, v) \mid v \in \text{Top-k}(S_u^n)\}, \tag{8}$$

where yielding a sparse graph $\mathcal{G}_d = (\mathcal{V}, \mathcal{E}_d)$ that captures implicit relational cues not observed in the raw topology. We then fuse the structure-pruned graph $\mathcal{G}_p$ with the learnable dynamic graph $\mathcal{G}_d$ to obtain a co-evolved structure: hereby inducing a sparse auxiliary graph $\mathcal{G}_d = (\mathcal{V}, \mathcal{E}_d)$ that encodes latent, potentially task-relevant dependencies beyond the observed topology. To integrate both structural priors and dynamically inferred cues, we fuse this learned structure with the previously pruned graph $\mathcal{G}_p$ to obtain a complement relational backbone:

$$\tilde{\mathcal{E}} = \mathcal{E}_p \cup \mathcal{E}_d, \quad \tilde{\mathcal{G}} = \mathcal{G}_p \cup \mathcal{G}_d = (\mathcal{V}, \tilde{\mathcal{E}}). \tag{9}$$

This co-evolved structure achieves dual objectives by suppressing irrelevant or spurious connections through edge pruning and simultaneously uncovering latent, task-relevant dependencies via similarity-driven graph construction. The resulting topology, denoted as $\tilde{\mathcal{G}}$, continuously aligns with the evolving semantic representation space, thereby enabling more reliable and informative message propagation throughout the graph.

Resilience-Oriented Prediction and Optimization

Building upon the dynamically refined structure, we further align representation learning with the resilience objective through supervised prediction and structure-aware optimization. This stage enables downstream signals to directly guide the evolution of topological composition.

Resilience Prediction To infer the resilience status of each enterprise, we deploy an interchangeable message propagation $\mathcal{Z}_\theta$ (Hamilton, Ying, and Leskovec 2017) over the fused graph $\tilde{\mathcal{G}}$, which consolidates both structurally retained and dynamically discovered enterprise relations. This graph-guided propagation enables the integration of localized and latent signals into richer node representations. Given the input feature matrix $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times d}$, the representation refinement is formulated as:

$$\mathbf{Z} = \mathcal{Z}_\theta(\tilde{\mathcal{G}}, \mathbf{X}), \quad \mathbf{Z} \in \mathbb{R}^{|\mathcal{V}| \times h}. \tag{10}$$

The results are passed via a feedforward prediction head:

$$\hat{\mathbf{Y}} = \mathbf{W}_\theta^{(4)} \cdot \sigma(\mathbf{W}_\theta^{(3)} \mathbf{Z} + \mathbf{b}_\theta^{(3)}) + b_\theta^{(4)}, \tag{11}$$

where $\mathbf{W}_\theta^{(3)} \in \mathbb{R}^{h \times h}$ and $\mathbf{b}_\theta^{(3)} \in \mathbb{R}^h$ define the intermediate transformation layer in detail, and $\mathbf{W}_\theta^{(4)} \in \mathbb{R}^{h \times C}$ and $\mathbf{b}_\theta^{(4)} \in \mathbb{R}^C$ denote the final classification parameters. The activation function $\sigma(\cdot)$ is instantiated as LeakyReLU. Supervision is applied directly to the labeled subset $\mathcal{D} \subset \mathcal{V}$ via a standard cross-entropy objective:

$$\mathcal{L}_p = -\sum_{v \in \mathcal{D}} \sum_{c=1}^{C} y_{v_c} \log \hat{y}_{v_c}, \tag{12}$$

where $y_{v_c} \in \{0, 1\}$ denotes the ground-truth one-hot indicator for node v and category c. This formulation ensures that the model is directly optimized for the resilience classification objective in a label-efficient manner.

Structure Evolution and Optimization Throughout training, the fused topology $\tilde{\mathcal{G}}$ undergoes continual iterative refinement, acting as a flexible and adaptive structural prior that evolves in close tandem with the task-driven objective. At each training iteration e, the graph $\tilde{\mathcal{G}}^{(e)} = (\mathcal{V}, \tilde{\mathcal{E}}^{(e)})$ is carefully reconstructed through a combination of latent relation inference and pruning-based denoising, guided jointly by the frozen encoder $\mathcal{J}_\phi$ and the total function $\mathcal{F}_\theta$.

$$\tilde{\mathcal{G}}^{(0)} \to \tilde{\mathcal{G}}^{(1)} \to \cdots \to \tilde{\mathcal{G}}^{(T)}, \tag{13}$$

where the evolving topology thus remains responsive to task-specific supervision while maintaining semantic stability from pretrained embeddings. This iterative process ensures that the structure co-adapts with representation learning, enhancing both resilience, interpretability, and predictive performance.

Experiments

Datasets. In this study, our experiments use enterprise networks from Shanghai (SH), Shenzhen (SZ), and Beijing (BJ) stock exchanges over 2022–2023. Due to data sparsity in BJ, we construct a fused dataset (BSS) by aggregating BJ, SH, and SZ to ensure sufficient scale and structural richness. Nodes represent enterprises and edges encode ties such as investment and governance links, with resilience labels (R/Non-R) assigned from financial disclosures and adverse events. All datasets are split into training, validation, and test sets (6:2:2).

Baselines. To validate the effectiveness of our proposed CFU, we compare it with three categories of baselines: (1) tabular classifiers (LR (LaValley 2008), XGBoost (Chen and Guestrin 2016), MLP (Taud and Mas 2017)) that ignore structural context, (2) static GNNs (GCN (Kipf and Welling 2016), GraphSAGE (Hamilton, Ying, and Leskovec 2017), GAT (Veličković et al. 2017), DAGNN (Liu, Gao, and Ji 2020)) that rely on fixed topologies but suffer from noisy and incomplete edges, and (3) dynamic methods (PinSAGE (Ying et al. 2018), Praga (Huang et al. 2025), UnGSL (Han et al. 2025)) designed to adapt message passing to evolving graph structures.

Evaluation Metrics. To evaluate the identification of enterprise resilience in this study, we use four standard metrics: F1 score, ROC-AUC, PR-AUC, and Kolmogorov–Smirnov (KS). These respectively reflect classification balance, discriminative power, ranking quality for non-resilient cases, and distributional separability, providing a comprehensive and multidimensional assessment of model performance under varying class distributions and noise conditions.

Experimental Configurations. All experiments are conducted on a single NVIDIA RTX 4090 GPU. Our CFU is trained using the Adam optimizer with a learning rate of 0.001 and a weight decay of 0.0005. We set the hidden dimensionality d to 64, apply a dropout rate of 0.2, and train for 2,000 iterations. The similarity graph retains top-$k = 10$ neighbors per node. Structural refinement is triggered every $e_r = 5$ iterations, following a warm-up phase of $e_w = 10$ iterations to stabilize the representation space.

Overall Performance. As shown in Table 1, methods without graph structures (e.g., LR, XGBoost, MLP) achieve reasonable performance in several datasets. This indicates that enterprise-level financial and ownership features contain meaningful signals for resilience assessment. However, their effectiveness is constrained by the lack of relational context, especially in datasets with interdependencies (e.g., 2022-BSS). Static GNNs (e.g., GCN, GraphSAGE, GAT) incorporate neighborhood information but often suffer from performance degradation when the underlying graph contains high redundancy. For instance, GCN consistently underperforms, which can be attributed to its uniform neighbor aggregation mechanism that propagates redundancy or weakly relevant features, amplifying pseudo-correlations in dense graphs. Dynamic GNNs (e.g., PinSAGE, Praga, UnGSL) attempt to construct task-adaptive graphs from features, yet they frequently overlook structural priors and exhibit instability across datasets, particularly when feature noise is non-negligible. In contrast, our proposed CFU achieves consistently superior results, benefiting from a co-evolving framework that simultaneously removes redundant structural links and uncovers latent, dynamic, uncertain dependencies. This dual refinement enables more reliable message passing and better aligns the evolving graph with semantic and predictive objectives, which proves especially advantageous in large and structurally diverse enterprise networks.

Dataset	Metric	Classic Methods			Static GNNs				Dynamic GNNs			Ours
		LR	XGBOOST	MLP	GCN	GraphSAGE	GAT	DAGNN	PinSAGE	Praga	UnGSL	CFU
2022-SH	F1	0.3448	0.3902	0.4834	0.2844	0.4419	0.4108	0.3200	0.5042	0.4118	0.5079	**0.5444**
	ROC-AUC	0.6213	0.5931	0.6021	0.6057	0.6310	0.6042	0.6071	0.6310	0.5053	0.6213	**0.6624**
	PR-AUC	0.4834	0.4417	0.4509	0.4558	0.5054	0.4738	0.4519	0.4914	0.3423	0.4507	**0.5312**
	KS	0.2331	0.1691	0.1924	0.1963	0.2145	0.1817	0.1856	0.2508	0.0317	0.2198	**0.2667**
2022-SZ	F1	0.4889	0.3833	0.4904	0.3028	0.4925	0.4694	0.2488	0.4106	0.4059	0.2335	**0.5591**
	ROC-AUC	0.7171	0.6489	0.7023	0.6599	0.7197	0.5896	0.6922	0.6784	0.5247	0.5531	**0.7461**
	PR-AUC	0.6173	0.5049	0.6063	0.5264	0.6259	0.4772	0.5534	0.5535	0.3193	0.3930	**0.6629**
	KS	0.3481	0.2536	0.3397	0.2557	0.3841	0.2011	0.2944	0.2687	0.0495	0.1248	**0.3999**
2022-BSS	F1	0.4492	0.4304	0.4258	0.2038	0.4420	0.3651	0.4458	0.4064	0.3657	0.1848	**0.4569**
	ROC-AUC	0.7062	0.6318	0.6863	0.6598	0.6922	0.5467	0.6716	0.6322	0.5295	0.6230	**0.7252**
	PR-AUC	0.6096	0.4618	0.5911	0.5203	0.6086	0.3943	0.5682	0.5072	0.3465	0.4702	**0.6270**
	KS	0.3138	0.1977	0.2886	0.2559	0.3096	0.1873	0.2817	0.2757	0.0813	0.2192	**0.3435**
2023-SH	F1	0.3143	0.3871	0.4315	0.2487	0.3158	0.3372	0.2176	0.3060	0.4190	0.3304	**0.4961**
	ROC-AUC	0.6492	0.5655	0.6680	0.6631	0.6701	0.5629	0.6381	0.6516	0.4997	0.6217	**0.7000**
	PR-AUC	0.5459	0.4302	0.5351	0.5244	0.5563	0.3962	0.4996	0.5030	0.3436	0.4726	**0.6057**
	KS	0.2618	0.1258	0.2445	0.2477	0.2649	0.1443	0.2327	0.2714	0.0147	0.1561	**0.2965**
2023-SZ	F1	0.4206	0.2843	0.4533	0.3485	0.4567	0.1453	0.1829	0.4126	0.3613	0.3986	**0.5314**
	ROC-AUC	0.6648	0.5954	0.7300	0.6292	0.7103	0.5377	0.6970	0.6342	0.5069	0.5416	**0.7759**
	PR-AUC	0.4925	0.4135	0.5802	0.4256	0.5795	0.3986	0.5109	0.4073	0.2803	0.3599	**0.6352**
	KS	0.3640	0.2248	0.3637	0.2349	0.3342	0.1817	0.2801	0.2520	0.0254	0.1042	**0.4457**
2023-BSS	F1	0.4374	0.3256	0.4324	0.0950	0.3919	0.4356	0.3912	0.4573	0.3669	0.3876	**0.4864**
	ROC-AUC	0.7310	0.6325	0.7315	0.5758	0.7161	0.5529	0.5064	0.7239	0.4964	0.5201	**0.7564**
	PR-AUC	0.5892	0.4640	0.5853	0.3996	0.5625	0.3822	0.3100	0.5656	0.2971	0.3753	**0.6054**
	KS	0.3591	0.2309	0.3536	0.1455	0.3494	0.1063	0.0572	0.3644	0.0159	0.1105	**0.3840**

Table 1: Performance comparison across datasets. Best results in bold, second-best underlined.

Method			2022-BSS				2023-BSS			
CFU-CRP	CFU-AEP	CFU-DSC	F1	ROC-AUC	PR-AUC	KS	F1	ROC-AUC	PR-AUC	KS
	✓	✓	0.4378	0.6900	0.5941	0.3018	0.4412	0.7373	0.5808	0.3692
✓		✓	0.4513	0.7019	0.6107	0.3237	0.4758	0.7068	0.5629	0.3280
✓	✓		0.4331	0.6651	0.5718	0.2739	0.4458	0.7289	0.5893	0.3609
✓	✓	✓	**0.4569**	**0.7252**	**0.6270**	**0.3435**	**0.4864**	**0.7564**	**0.6054**	**0.3840**

Table 2: Ablation study on the BSS datasets.

Ablation Study. Table 2 presents the ablation results on the BSS datasets, highlighting both the individual importance and the interdependence of each core component in CFU. When CFU-CRP is removed, noticeable declines in metric PR-AUC and separability metric KS suggest weakened ability to distinguish subtle resilience differences. This degradation also impairs the effectiveness of the CFU-AEP, which relies on stable semantics to filter redundant links. Excluding AEP further undermines graph discriminability, as excessive relational redundancy introduces noise into the propagation process, particularly harming KS and F1 scores. The absence of dynamic structure-aware propagation (CFU-DSC) results in the most comprehensive performance drop, reflecting its role in refining representations and reinforcing edge quality during augmentation. Overall, CRP, AEP, and DSC form a tightly coupled triad. Their synergy is indispensable for mitigating structural uncertainty, preserving graph integrity, and ultimately promoting robust and reliable resilience identification in volatile enterprise networks.

Impact of Scheduling in AEP. Figure 4 shows how performance varies with the warm-up threshold e_w and pruning refresh interval e_r on 2022-BSS, both key to the Agreement-Guided Edge Pruning (AEP). Too-small e_w causes premature pruning before representations stabilize, while overly large e_w delays structural correction. Similarly, small e_r introduces volatility, whereas large e_r misses timely redundancy reduction. Optimal performance is achieved near $e_w = 10$ and $e_r = 5$, confirming that AEP works best when synchronized with evolving node semantics.

Edge Dynamics under Structural Evolution. Figure 5 illustrates how edge counts evolve with training iterations, comparing the full CFU with its DSC-disabled (w/o DSC) variant. While they both initially reduce edge density, CFU quickly stabilizes, maintaining a structurally compact yet informative graph. Without DSC, the CFU continues pruning aggressively, leading to excessive sparsity. This highlights the role of DSC in preserving informative links by align-

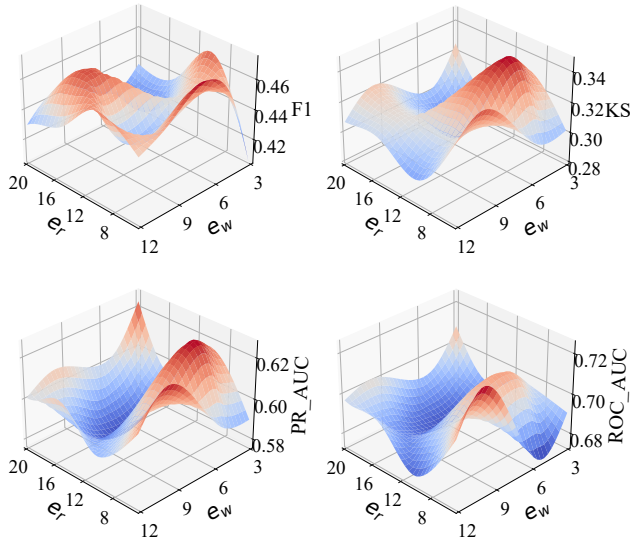

Figure 4: Performance landscape over different warm-up thresholds e_w and pruning intervals e_r on the 2022-BSS.

ing structure updates with task-guided semantics. Crucially, it allows the model to capture uncertain and evolving inter-enterprise relations that static pruning alone would overlook, ensuring the graph remains responsive to dynamic dependencies throughout training.

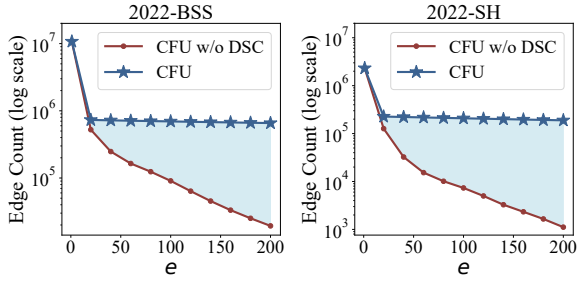

Figure 5: The count of edges on the 2022-BSS and 2022-SH datasets changes along with the iterations e.

Impact of k on DSC. Figure 6 examines the sensitivity of performance to the top-k parameter in our Dynamic Structure Co-evolving (DSC) module. We observe that moderate values of k yield the best overall results across both years, striking a balance between completeness and selectivity. When k is too small, the inferred structure lacks sufficient relational coverage, limiting the propagation of informative signals. In contrast, large k values introduce excess edges that may reintroduce structural redundancy, diminishing the discriminative power of the fused topology. The results confirm that a well-calibrated k is essential for capturing latent dependencies while maintaining a sparse and interpretable relational graph.

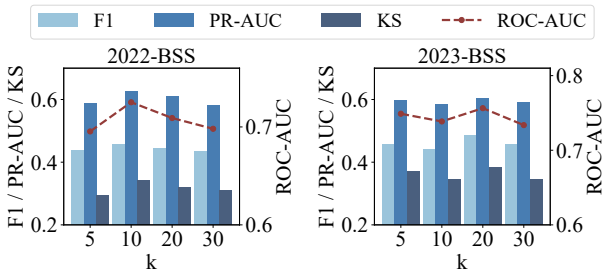

Figure 6: Effect of the top-k parameter in DSC on performance across two BSS datasets.

Effectiveness of CFU as a Plug-in Framework. Table 3 summarizes the relative performance gain when CFU is integrated into three commonly used GNN backbones (e.g., GCN, GAT, and GraphSAGE) on the 2022-BSS dataset. The improvement is observed when CFU is applied to GAT, particularly in PR-AUC (+44.4%) and KS (+35.6%), indicating CFU's strong capacity to suppress relational redundancy and enhance class separability. For GCN, which suffers from over-smoothing in dense graphs, CFU increases F1 by over 64%, highlighting its ability to refine noisy propagation paths. Even on the relatively robust GraphSAGE baseline, CFU yields steady gains in all metrics. These results verify that CFU serves as an effective standalone method and acts as a versatile plug-in component under varied structural assumptions.

Method	F1	ROC-AUC	PR-AUC	KS
GCN	0.2038	0.6598	0.5203	0.2559
GCN-CFU	0.3349	0.6730	0.5628	0.2561
↑ Improvement	+64.3%	+2.0%	+8.2%	+0.1%
GAT	0.3651	0.5467	0.3943	0.1873
GAT-CFU	0.3671	0.6592	0.5692	0.2542
↑ Improvement	+0.5%	+20.6%	+44.4%	+35.6%
GraphSAGE	0.4420	0.6922	0.6086	0.3096
GraphSAGE-CFU	0.4569	0.7252	0.6270	0.3435
↑ Improvement	+3.4%	+4.8%	+3.0%	+11.0%

Table 3: Performance improvement (%) by integrating CFU into different GNN backbones on 2022-BSS.

Conclusion

We introduced CFU, which effectively suppresses structural redundancy and uncovers latent dependencies, consistently outperforming strong baselines on large-scale, real-world enterprise datasets spanning multiple markets and varying temporal conditions. Beyond resilience prediction, our co-evolving design offers a general and unified paradigm for modeling dynamic and incomplete graphs, thereby opening promising new avenues for broader applications in financial systems, risk management, supply chains, critical infrastructure monitoring, and other complex, interdependent networks worldwide.

Acknowledgments

This work was supported by the Natural Science Foundation of Sichuan Province (No. 2025ZNSFSC1495), the National Natural Science Foundation of China (No. 62572097), the Sichuan Science and Technology Program (No. 2023ZYD0145), and the Chengdu Science and Technology Program (No. 2023-JB00-00016-GX).

References

Bi, W.; Xu, B.; Sun, X.; Wang, Z.; Shen, H.; and Cheng, X. 2022. Company-as-tribe: Company financial risk assessment on tribe-style graph with hierarchical graph neural networks. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2712–2720.

Chen, T.; and Guestrin, C. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794.

Chen, Z.; Chen, W.; and Shi, Y. 2020. Ensemble learning with label proportions for bankruptcy prediction. *Expert Systems with Applications*, 146: 113155.

Cheng, D.; Tu, Y.; Ma, Z.-W.; Niu, Z.; and Zhang, L. 2019. Risk Assessment for Networked-guarantee Loans Using High-order Graph Attention Representation. In *IJCAI*, 5822–5828.

Danenas, P.; and Garsva, G. 2015. Selection of support vector machines based classifiers for credit risk domain. *Expert systems with applications*, 42(6): 3194–3204.

Erdogan, B. E. 2013. Prediction of bankruptcy using support vector machines: an application to bank bankruptcy. *Journal of Statistical Computation and Simulation*, 83(8): 1543–1555.

Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30.

Han, S.; Zhou, Z.; Chen, J.; Hao, Z.; Zhou, S.; Wang, G.; Feng, Y.; Chen, C.; and Wang, C. 2025. Uncertainty-Aware Graph Structure Learning. In *Proceedings of the ACM on Web Conference 2025*, WWW '25, 4863–4874. Association for Computing Machinery.

Huang, X.; Ma, Z.; Meng, D.; Liu, Y.; Ruan, S.; Sun, Q.; Zheng, X.; and Qiao, Z. 2025. PRAGA: prototype-aware graph adaptive aggregation for spatial multi-modal omics analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 326–333.

Kipf, T. N.; and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.

Kou, G.; Xu, Y.; Peng, Y.; Shen, F.; Chen, Y.; Chang, K.; and Kou, S. 2021. Bankruptcy prediction for SMEs using transactional data and two-stage multiobjective feature selection. *Decision Support Systems*, 140: 113429.

LaValley, M. P. 2008. Logistic regression. *Circulation*, 117(18): 2395–2399.

Li, P.; Zhao, Q.; Liu, Y.; Zhong, C.; Wang, J.; and Lyu, Z. 2024. Survey and prospect for applying knowledge graph in enterprise risk management. *Computers, Materials and Continua*, 78(3): 3825–3865.

Liang, T.; Zeng, G.; Zhong, Q.; Chi, J.; Feng, J.; Ao, X.; and Tang, J. 2021. Credit risk and limits forecasting in e-commerce consumer lending service via multi-view-aware mixture-of-experts nets. In *Proceedings of the 14th ACM international conference on web search and data mining*, 229–237.

Liu, M.; Gao, H.; and Ji, S. 2020. Towards deeper graph neural networks. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 338–348.

Mai, F.; Tian, S.; Lee, C.; and Ma, L. 2019. Deep learning models for bankruptcy prediction using textual disclosures. *European journal of operational research*, 274(2): 743–758.

Sun, J.; Lang, J.; Fujita, H.; and Li, H. 2018. Imbalanced enterprise credit evaluation with DTE-SBD: Decision tree ensemble based on SMOTE and bagging with differentiated sampling rates. *Information Sciences*, 425: 76–91.

Taud, H.; and Mas, J.-F. 2017. Multilayer perceptron (MLP). In *Geomatic approaches for modeling land change scenarios*, 451–455. Springer.

Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; and Bengio, Y. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903*.

Wang, C.; Zhou, S.; Yu, K.; Chen, D.; Li, B.; Feng, Y.; and Chen, C. 2022. Collaborative knowledge distillation for heterogeneous information network embedding. In *Proceedings of the ACM web conference 2022*, 1631–1639.

Wei, S.; Lv, J.; Guo, Y.; Yang, Q.; Chen, X.; Zhao, Y.; Li, Q.; Zhuang, F.; and Kou, G. 2024. Combining intra-risk and contagion risk for enterprise bankruptcy prediction using graph neural networks. *Information Sciences*, 659: 120081.

Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2018. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*.

Yang, S.; Zhang, Z.; Zhou, J.; Wang, Y.; Sun, W.; Zhong, X.; Fang, Y.; Yu, Q.; and Qi, Y. 2021. Financial risk analysis for SMEs with graph-based supply chain mining. In *Proceedings of the twenty-ninth international conference on international joint conferences on artificial intelligence*, 4661–4667.

Yeh, C.-C.; Chi, D.-J.; and Lin, Y.-R. 2014. Going-concern prediction using hybrid random forests and rough set approach. *Information Sciences*, 254: 98–110.

Ying, R.; He, R.; Chen, K.; Eksombatchai, P.; Hamilton, W. L.; and Leskovec, J. 2018. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 974–983.

Zhang, C.; Chen, J.; Shu, T.; and Tan, J. 2022. Enterprise event risk detection based on supply chain contagion. In *2022 IEEE 9th International Conference on Data Science and Advanced Analytics (DSAA)*, 1–10. IEEE.

Zhang, Z.; Ji, Y.; Shen, J.; Chen, Y.; Zhang, X.; and Yang, G. 2024. Collaborative Metapath Enhanced Corporate Default Risk Assessment on Heterogeneous Graph. In *Proceedings of the ACM Web Conference 2024*, 446–456.

Zhao, J.; Wang, X.; Shi, C.; Hu, B.; Song, G.; and Ye, Y. 2021. Heterogeneous graph structure learning for graph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 4697–4705.

Zheng, Y.; Lee, V. C.; Wu, Z.; and Pan, S. 2021. Heterogeneous graph attention network for small and medium-sized enterprises bankruptcy prediction. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 140–151. Springer.

Zhu, Y.-N.; Luo, X.; Li, Y.-F.; Bu, B.; Zhou, K.; Zhang, W.; and Lu, M. 2020. Heterogeneous mini-graph neural network and its application to fraud invitation detection. In *2020 IEEE International Conference on Data Mining (ICDM)*, 891–899. IEEE.