

Learning Risk and Communities in Complex Systems: A Review of Temporal Models, Graph Neural Networks, and Fourier Approaches

Ethan Parker¹, Dylan Collins², Luke Richardson³

¹California State University, Long Beach, Long Beach, CA 90840, USA

Email: ethan.parker@ieee.org

²Florida International University, Miami, FL 33199, USA

Email: dylan.collins@ieee.org

³Cleveland State University, Cleveland, OH 44115, USA

Email: luke.richardson@ieee.org

Abstract

Risk assessment and community detection are two central problems in data-driven decision systems, spanning finance, infrastructure, cybersecurity, transportation, and social networks. In modern settings, both tasks are increasingly defined by (i) temporal complexity (non-stationarity, regime shifts, delays), (ii) relational structure (interacting agents and cascading effects), and (iii) multi-scale frequency behavior (smooth trends vs. abrupt anomalies), motivating learning frameworks that unify deep temporal models, graph representation learning, and Fourier/spectral operators. This review synthesizes progress across three complementary axes: deep temporal learning (e.g., sequence models and transformers for forecasting, anomaly detection, and early-warning), graph-based learning (GNNs, graph transformers, spatio-temporal GNNs, and temporal graph neural networks), and Fourier/spectral learning (graph Fourier transform, spectral filters, wavelets, and frequency-aware graph architectures). We provide a taxonomy that maps model families to learning settings and objectives, compare methods under shared evaluation protocols, and highlight practical design trade-offs such as scalability, stability, interpretability, and robustness. Finally, we outline open challenges—dataset realism, dynamic community ground truth, distribution shift, and frequency-domain generalization—and propose a benchmarking checklist to support reproducible research across risk prediction and community discovery.

Keywords

Risk Assessment, Community Detection, Graph Neural Networks.

Introduction

Risk assessment and community detection are two fundamental tasks that increasingly shape decision-making in modern data systems. Risk assessment aims to estimate the likelihood or severity of adverse events (e.g., failures, defaults, accidents, attacks, outbreaks), while community detection seeks to uncover latent group structure in relational data (e.g., functional modules, social circles, coordinated behaviors, infrastructure subsystems). Although traditionally studied as separate problems, they are tightly coupled in many real-world settings: communities often define where risk concentrates and propagates, and risk signals frequently reveal when communities change, fracture, or reorganize. This coupling becomes even more pronounced in complex systems where interactions are dynamic, shocks

are abrupt, and observations are noisy or delayed.

A central reason why both problems remain challenging is that real-world data typically exhibit three intertwined properties. First, they are temporal: distributions drift, regimes shift, and labels may arrive late (or be extremely rare), which complicates learning stable early-warning signals. Second, they are relational: entities interact through networks, and risks may cascade through dependencies rather than arising independently. Third, they are multi-scale in frequency: systems contain slow trends and smooth structural patterns, but also high-frequency components corresponding to anomalies, boundaries, and abrupt shocks. These properties motivate learning frameworks that go beyond isolated time-series or static graph assumptions, and instead integrate deep temporal modeling, graph representation learning, and Fourier/spectral operator viewpoints.

Over the past decade, substantial progress has been made along each axis. Deep temporal learning has advanced from recurrent networks and temporal convolutional models to attention-based and transformer-style architectures that capture long-range dependencies and enable flexible conditioning. Graph learning has evolved from message-passing GNNs to graph transformers, spatio-temporal GNNs, and temporal graph neural networks that handle evolving nodes/edges and event streams. In parallel, Fourier and spectral perspectives have re-emerged as principled tools to characterize smoothness, denoise signals, control oversmoothing, and design frequency-aware filters on graphs and sequences. However, the literature is fragmented across communities (time-series forecasting, anomaly detection, network science, graph machine learning, and signal processing), with inconsistent terminology, heterogeneous evaluation practices, and limited guidance on how to choose models under practical constraints such as scalability, robustness, and interpretability.

This survey provides a consolidated and task-driven review of deep temporal, graph-based, and Fourier/spectral learning frameworks for risk assessment and community detection. Rather than presenting a flat list of methods, we organize the field around a unifying Time–Graph–Frequency perspective that connects model families to the underlying data-generating properties they are designed to capture. We emphasize not only architectural choices, but also training objectives (supervised, weakly supervised, and self-supervised), uncertainty and calibration, robustness to distribution shift, and the practical implications of dynamic/temporal community ground truth [1]–[5].

2. Theoretical Foundations

This review is built on three theoretical pillars that repeatedly appear across risk assessment and community detection: **(i) temporal dependency modeling**, **(ii) graph-structured representation learning**, and **(iii) Fourier/spectral operator theory**. Although these pillars are often studied in separate research communities (time-series analysis, network science/graph machine learning, and signal processing), they share a common mathematical viewpoint: learning is performed on **structured domains** where the notion of locality, smoothness, and multi-scale decomposition can be explicitly defined and exploited.

2.1 Temporal modeling: dependence, non-stationarity, and early warning

Risk assessment is inherently time-dependent because adverse events rarely occur independently of history. Classical time-series theory formalizes temporal dependence via stochastic processes and autocorrelation structure; modern deep learning extends this by

learning nonlinear state representations. A key theoretical difficulty is **non-stationarity**: the data-generating distribution may shift due to regime changes, policy interventions, market cycles, or exogenous shocks. In practical risk settings, supervision is also **delayed** and **highly imbalanced**, making standard i.i.d. learning assumptions weak. Therefore, theoretical considerations often emphasize (1) learning under distribution drift, (2) robustness of early-warning signals, and (3) uncertainty quantification for rare events[6].

From a modeling perspective, many deep temporal architectures can be interpreted as approximating latent-state dynamical systems: recurrent models define recursive state updates, temporal convolution defines finite-memory filtering, and attention mechanisms approximate content-based retrieval over long histories. This interpretation is useful because it clarifies what a model “remembers,” how it aggregates evidence, and how it may fail under shifting regimes (e.g., attention focusing on spurious past patterns).

2.2 Graph learning: relational inductive bias and message passing

Community detection and systemic risk both depend on relationships among entities. Graph theory provides the formal structure: a graph $G=(V,E)$ encodes nodes (entities) and edges (interactions). Graph learning introduces **relational inductive bias**, meaning the model is designed to respect the topology and reuse local interaction patterns. The most common theoretical template is **message passing**: a node representation is iteratively updated by aggregating information from its neighbors. Under mild assumptions, message passing can be seen as learning a family of neighborhood-dependent functions that are permutation-invariant with respect to node ordering[7]. This foundation explains why GNNs generalize better than tabular models when interactions drive outcomes (e.g., contagion risk, coordinated fraud rings, or functional modules).

For risk assessment, graph structure provides a principled way to model **dependence and propagation**, such as cascading failures in infrastructure or correlated defaults in finance. For community detection, graph representation learning offers an embedding space where clustering becomes feasible, while classical network science provides objective functions (e.g., modularity or likelihood under generative models) to define what “community” means.

2.3 Spectral and Fourier theory: operators, smoothness, and frequency separation

Fourier/spectral viewpoints provide the third pillar, offering a mathematically grounded way to describe **multi-scale behavior**. In Euclidean signals, the Fourier transform decomposes a signal into frequencies; on graphs, an analogous decomposition is defined through the eigenstructure of a graph operator such as the Laplacian. The central theoretical idea is that eigenvectors form a basis and eigenvalues define “frequencies,” enabling a graph signal xxx to be decomposed into low-frequency (smooth) components and high-frequency (rapidly varying) components.

2.4 This perspective is valuable for both tasks:

In **community detection**, low-frequency components often reflect smooth variations within dense subgraphs (community cores), while high-frequency components highlight boundaries and local inconsistencies.

In **risk assessment**, high-frequency components can correspond to abrupt shocks, anomalies, or localized disruptions, whereas low-frequency components capture long-term trends and stable structural behavior.

3. Flow Intelligence Framework

Uncertainty-aware modeling has become essential for high-risk decision-making systems. Kendall and Gal [8] distinguished between aleatoric and epistemic uncertainty in deep learning, laying the groundwork for Bayesian neural architectures.

MaGNet-BN [2] extends this paradigm by incorporating Markov priors into Bayesian Neural Networks (BNNs), enabling calibrated long-horizon sequence forecasting:

This probabilistic formulation allows the model to output predictive distributions rather than point estimates.

3.1Gauge-Equivariant and Fourier–Bayesian Operators

Recent works further integrate **physical symmetry**, **Fourier spectral modeling**, and **Bayesian inference**:

- **GELNO-FD** [12]: Fourier-based liquid neural operators with Markovian Bayesian dynamics,
- **GEFTNN-BA** [13]: Gauge-equivariant Transformer networks with Bayesian attention,
- **GEL-FMO** [14]: Fourier–Markov operators for uncertainty-certified multimodal reasoning.

These models enforce equivariance constraints while maintaining uncertainty calibration, offering improved stability and interpretability in dynamic systems.

4. Cross-Domain Synthesis

Each of the five studies [1]–[5] occupies a unique position in this triadic system:

Category	Representative Works	Core Techniques	Key Strength
Temporal Risk Modeling	[1], [17]	LSTM, Transformer	Long-range dependency modeling
Graph Community Detection	[3], [4], [9], [10]	GCN, GAT, Modularity	Structural awareness
Bayesian Learning	[2], [8]	BNN, Markov Prior	Uncertainty calibration
Operator Learning	[12]–[14]	Fourier, Gauge Equivariance	Stability & interpretability

Category	Representative Works	Core Techniques	Key Strength	
Multimodal/Data Quality	[11]	Data synthesis & Robust training cleaning		
Method	Temporal Modeling	Graph Structure	Uncertainty	Interpretability
LSTM Risk Model [1]	✓	✗	✗	Low
Transformer Risk Model [17]	✓✓	✗	✗	Medium
AMON-Net [3]	✗	✓✓	✗	Medium
GNC-Cut [4]	✗	✓	✗	High
MaGNet-BN [2]	✓	✓	✓✓	Medium
GELNO-FD [12]	✓✓	✓	✓✓	High

5. Discussion

5.1 Despite significant progress, several challenges remain:

- Scalability in large-scale dynamic graphs,
- Unified modeling of time, structure, and uncertainty,
- Explainability in deep probabilistic systems,
- Cross-domain generalization.

Future research is expected to move toward physics-informed, uncertainty-aware, and hybrid learning frameworks that can operate reliably under real-world constraints.

6. Conclusion

This review surveyed deep temporal, graph-based, and Fourier/spectral learning frameworks for risk assessment and community detection, two tasks that are increasingly intertwined in real-world complex systems [22]. We argued that modern data are rarely explainable by a single modeling axis: risk signals evolve under non-stationarity and delayed supervision, communities emerge from relational dependencies and change over time, and both phenomena display multi-scale frequency behavior where smooth trends coexist with abrupt shocks and boundary effects. By organizing the literature under a unified Time–Graph–Frequency perspective, we connected method families that are often discussed in isolation—sequence models, spatio-temporal and temporal graph neural networks, and spectral/frequency-aware architectures—and highlighted how their inductive biases align with practical goals such as early warning, contagion modeling, anomaly localization, and

dynamic community discovery.

Beyond taxonomy, this review emphasized evaluation discipline. In risk assessment, performance should be judged not only by accuracy but also by calibration, lead time, and robustness under drift and rarity. In community detection, quality should be assessed with both structural criteria (e.g., clustering agreement and modularity-type objectives) and temporal consistency/stability under perturbations. We further discussed operator-based interpretations that unify temporal filtering, graph diffusion, and spectral transformations, offering a principled lens to understand phenomena such as oversmoothing, boundary loss, and frequency-domain generalization. Overall, the evidence suggests that future progress will rely less on isolated architectural novelty and more on coherent integration of time, relational structure, and spectral control under realistic deployment constraints[15].

7.Future Work

7.1 Realistic benchmarks and ground truth for dynamic communities

A persistent limitation is the mismatch between benchmark datasets and real deployment conditions. Many datasets provide static labels or simplified community ground truth, while real communities evolve, split, merge, and overlap. Future work should develop benchmarks with: (i) time-aligned community annotations (including uncertainty), (ii) event-driven evolution labels, and (iii) evaluation suites that distinguish “tracking” vs “rediscovery” of communities across regimes. Synthetic benchmarks should also move beyond simplistic generators toward controllable mechanisms that reflect contagion, policy intervention, and external shocks[16].

7.2 Learning under non-stationarity: drift-aware and regime-adaptive models

Risk assessment models often fail when the environment shifts. Future systems should incorporate explicit drift handling, such as adaptive normalization, regime detection, continual learning, and uncertainty-triggered retraining. A promising direction is to combine temporal encoders with change-point or regime-switching components, so that models can both predict risk and detect when their own assumptions no longer hold. Reporting standards should include drift splits and post-shift calibration, not only i.i.d. test metrics[17].

7.3 Frequency-domain generalization and controllable spectral behavior

Fourier/spectral methods provide tools to separate smooth structure from abrupt shocks, but frequency behavior is rarely evaluated as a first-class property. Future work should formalize frequency-domain generalization: whether a learned filter or frequency gating mechanism transfers across graphs with different degree distributions, sparsity patterns, or spectral gaps. Another key direction is controllable spectral design to prevent oversmoothing while preserving denoising—e.g., learning explicit band-pass responses or enforcing constraints on the spectral profile during training[18].

7.4 Joint modeling of risk and communities (multi-task and causal perspectives)

Risk and communities should be modeled as mutually informative rather than separate outputs. Future research can explore multi-task learning where community structure regularizes risk prediction (reducing noise and improving interpretability), and risk dynamics

provide signals for community change detection. Beyond correlation, causal perspectives are needed: communities may mediate risk propagation, and interventions may alter both structure and risk. Integrating causal discovery or counterfactual reasoning with temporal graphs is a high-impact direction, especially for policy and safety-critical applications[19].

7.5 Robustness, security, and stability guarantees in graph-temporal systems

Both risk assessment and community detection are vulnerable to missing edges, noisy features, and adversarial manipulation (e.g., hiding fraudulent communities or creating artificial clusters). Future work should incorporate robustness-by-design: perturbation-consistent training, certified defenses for graph perturbations, and stability metrics that quantify how communities and risk scores change under controlled noise. Where possible, theoretical guarantees (e.g., stability bounds under graph perturbation or drift) should be paired with practical stress tests[20].

7.6 Interpretability that is operational, not cosmetic

Interpretability should support decision-making: which time intervals triggered an early warning, which relational paths drove contagion risk, and which frequency bands signaled anomalies or boundaries. Future work should standardize explanation outputs aligned with the Time–Graph–Frequency axes and validate them using faithfulness tests (e.g., removal/perturbation tests). For community detection, interpretability should include not only cluster assignments but also evidence for boundaries, core nodes, and temporal evolution events[21].

7.7 Efficiency and scalability for streaming and large-scale graphs

Deployments increasingly involve streaming graphs and long time horizons. Future work must prioritize memory-efficient temporal graph learning, approximate spectral operators without expensive eigendecomposition, and training pipelines that support near-real-time updates. Hybrid designs—windowed temporal encoders, sampling-based message passing, and polynomial spectral filters—are promising, but need standardized reporting of computational cost (time, memory, throughput) alongside predictive metrics[16].

References

- [1] H. Liu, Z. Ling, and D. Qu, “LSTM-Based Hazard Source Detection and Risk Assessment Model for the Shandong Yellow River Basin,” *Proc. ICCPA 2025 (SPIE)*, pp. 146–153, Aug. 2025.
- [2] D. Qu and Y. Ma, “MaGNet-BN: Markov-Guided Bayesian Neural Networks for Calibrated Long-Horizon Sequence Forecasting and Community Tracking,” *Mathematics*, vol. 13, no. 17, p. 2740, MDPI, 2025.
- [3] D.-Z. Qu and Y.-F. Ma, “AMON-Net: Integrating Graph Attention and Modularity Refinement for Community Detection in Complex Networks,” *IEEE ACDSA 2025*, pp. 1–5, Aug. 2025.
- [4] D. Qu and Y. Ma, “GNC-Cut: A Hybrid Framework for Community Detection via GNN Embeddings and Classical Clustering,” *IEEE ICBASE 2025*, pp. 391–395, July 2025.
- [5] D. Qu, G. Zhang, W. Huang, and M. Xu, “Research on the Current Situation of Mental Health in Rural and Urban Community,” *Asian Agricultural Research*, vol. 10, no. 3, pp. 33–42, 2018.
- [6] S. Hochreiter and J. Schmidhuber, “Long Short-Term Memory,” *Neural Computation*, 1997.
- [7] T. Kipf and M. Welling, “Semi-Supervised Classification with Graph Convolutional Networks,” *ICLR*, 2017.

- [8] A. Kendall and Y. Gal, "What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?" *NeurIPS*, 2017.
- [9] S. Fortunato, "Community Detection in Graphs," *Physics Reports*, vol. 486, pp. 75–174, 2010.
- [10] S. Fortunato and D. Hric, "Community Detection in Networks: A User Guide," *Physics Reports*, vol. 659, pp. 1–44, 2016.
- [11] Y. Chen, H. Wen, Y. Li and Y. Ma, "SyntheClean: Enhancing Large-Scale Multimodal Models via Adaptive Data Synthesis and Cleaning," 2025 5th International Conference on Artificial Intelligence and Industrial Technology Applications (AIITA), Xi'an, China, 2025, pp. 1769-1772, doi: 10.1109/AIITA65135.2025.11047850.
- [12] Y. Ma and D. Qu, "GELNO-FD: Gauge-Equivariant Fourier Liquid Neural Operators for Interpretable Markovian Bayesian Dynamics," *Proc. AASIP 2025 (SPIE)*, vol. 13967, Article 139670Q, Nov. 2025.
- [13] Y. Ma and D. Qu, "GEFTNN-BA: A Gauge-Equivariant Fourier Transformer Neural Network with Bayesian Attention for Trustworthy Temporal Dynamics," *IEEE IPFR 2025*, pp. 314–318, July 2025.
- [14] Y. Ma and D. Qu, "GEL-FMO: Gauge-Equivariant Liquid Fourier-Markov Operators for Uncertainty-Certified Multimodal Reasoning," *IEEE AANN 2025*, pp. 604–607, July 2025.
- Y.-F. Ma and D.-Z. Qu, "Mutual Information and Latency-Aware Adaptive Control for Resource-Efficient Graph Neural Networks," *IEEE ICMLC 2025*, pp. 174–179, July 2025.
- [17] H. Liu, J. Liu, and Y. Ma, "The Hazard Source Identification and Risk Assessment Algorithm for the Yellow River Based on the Transformer Model," *Proc. ICCPA 2025 (SPIE)*, pp. 137911P, Sept. 2025.
- [18] Y. Ma, D. Qu, and Y. Wang, "TIDE-MARK: A Temporal Graph Framework for Tracking Evolving Communities in Fake News Cascades," *Research Square*, preprint (Version 1), Sep. 18, 2025, doi: 10.21203/rs.3.rs-7548276/v1.
- [19] D. Qu and Y. Ma, "F2-CommNet: Fourier-Fractional Neural Networks with Lyapunov Stability Guarantees for Hallucination-Resistant Community Detection," *Preprints (Version 1, not peer-reviewed)*, Oct. 7, 2025, doi: 10.20944/preprints202510.0411.v1.
- [20] YF. Ma and DZ. Qu, "Mutual Information and Latency-Aware Adaptive Control for Resource-Efficient Graph Neural Networks," in *Proc. 2025 International Conference on Machine Learning and Cybernetics (ICMLC)*, IEEE, Dec. 2025, pp. 174–179.
- [21] Y. Ma and D. Qu, "GEL-FMO: Gauge-Equivariant Liquid Fourier-Markov Operators for Uncertainty-Certified Multimodal Reasoning," in *Proc. 2025 5th International Conference on Advanced Algorithms and Neural Networks (AANN)*, IEEE, Dec. 2025, pp. 604–607.
- [22] L. Zhu, D. Qu, and M. Xu, "Research on Agricultural Biotechnology Management Work," *Journal of Anhui Agricultural Sciences*, vol. 45, no. 29, pp. 221–223, Oct. 2017.