

A Survey on Temporal Interaction Graph Representation Learning: Progress, Challenges, and Opportunities

Pengfei Jiao¹, Hongjiang Chen¹, Xuan Guo², Zhidong Zhao¹, Dongxiao He² and Di Jin^{2*}

¹School of Cyberspace, Hangzhou Dianzi University, China

²College of Intelligence and Computing, Tianjin University, China

{pjiao, hchen, zhaozd}@hdu.edu.cn, {guoxuan, hedongxiao, jindi}@tju.edu.cn

Abstract

Temporal interaction graphs (TIGs), defined by sequences of timestamped interaction events, have become ubiquitous in real-world applications due to their capability to model complex dynamic system behaviors. As a result, temporal interaction graph representation learning (TIGRL) has garnered significant attention in recent years. TIGRL aims to embed nodes in TIGs into low-dimensional representations that effectively preserve both structural and temporal information, thereby enhancing the performance of downstream tasks such as classification, prediction, and clustering within constantly evolving data environments. In this paper, we begin by introducing the foundational concepts of TIGs and emphasizing the critical role of temporal dependencies. We then propose a comprehensive taxonomy of state-of-the-art TIGRL methods, systematically categorizing them based on the types of information utilized during the learning process to address the unique challenges inherent to TIGs. To facilitate further research and practical applications, we curate the source of datasets and benchmarks, providing valuable resources for empirical investigations. Finally, we examine key open challenges and explore promising research directions in TIGRL, laying the groundwork for future advancements that have the potential to shape the evolution of this field.

1 Introduction

Temporal Interaction Graphs (TIGs), also known as dynamic networks, are a powerful data structure for modeling sequences of timestamped interaction events (i.e., edges) among entities (i.e., nodes). These graphs have been widely employed to capture the evolving nature of real-world systems such as e-commerce platforms, social networks, and recommendation engines. For example, as depicted in Fig. 1, a social network can be modeled as a TIG, where nodes represent users and edges denote interactions occurring at specific timestamps. Unlike static graphs, TIGs evolve over time and can capture complex temporal dynamics and dependencies, making them indispensable in data mining and machine learning

*Corresponding authors.

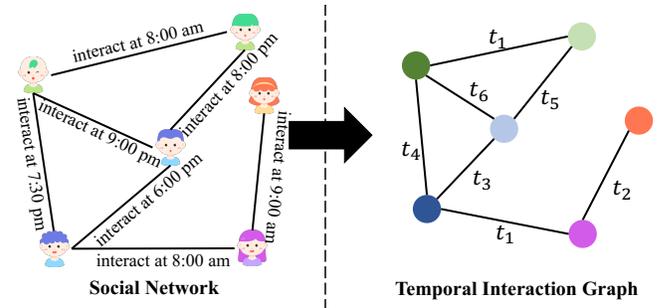


Figure 1: The overview of TIG.

applications—including recommendation systems, chemical synthesis modeling, and social network analysis.

As the adoption of TIGs grows, the need for learning effective embeddings to enhance downstream tasks (e.g., link prediction, node classification, and node clustering) has become increasingly crucial. Traditionally, to extract useful information from a TIG, temporal data is either discarded or transformed into a sequence of static graph snapshots, which may generate potential features but often overlook fine-grained regularities in the network’s evolution. To overcome this limitation, temporal interaction graph representation learning (TIGRL) has recently attracted considerable attention. TIGRL aims to learn mappings from the input space to a low-dimensional representation while preserving both temporal dynamics and structural evolution. However, methods developed for static graphs or snapshot-based temporal graphs are not directly applicable to TIGs because of their inherently dynamic nature and irregular temporal granularity. The challenges in learning representations for TIGs are multifaceted. First, TIGs exhibit dependencies in both spatial and temporal dimensions. Their topology varies considerably over time, and the evolution patterns are often intricate and non-linear; hence, TIGRL methods must simultaneously capture static structural relationships and temporal changes. Second, many real-world systems experience rapid and unpredictable dynamics, necessitating fine-grained representations that account for both short-term fluctuations and long-term trends. In contrast, snapshot-based approaches tend to preserve temporal order while oversimplifying temporal details, thereby rendering fine-grained modeling infeasible. Moreover, these methods must also satisfy real-time reasoning requirements with computational efficiency. Finally, the domain-specific nature of TIGs

introduces additional challenges: incorporating auxiliary information (e.g., node and edge attributes) often requires specialized domain knowledge, and the large-scale, continuously evolving characteristics of TIGs raise issues of scalability, storage, and parallel training.

Several surveys have investigated learning on TIGs; however, they have primarily focused on specific aspects. For example, some surveys offer a general overview of TIGs and classify models based on encoder–decoder architectures [Kazemi *et al.*, 2020], while others emphasize the distinctions between TIGs and static graphs by categorizing models according to temporal modeling techniques [Longa *et al.*, 2023; Zheng *et al.*, 2025]. Additionally, certain works focus on specialized downstream tasks such as temporal link prediction [Qin and Yeung, 2023] or temporal graph generation [Gupta and Bedathur, 2022]. These efforts, while valuable, have limitations: (1) outdated coverage of recent advancements, (2) insufficient categorization of TIG methodologies, (3) limited discussion of datasets and benchmarks, and (4) inadequate attention to emerging applications and challenges.

To address these gaps, this paper provides a comprehensive review of recent advancements in TIGRL. In particular, our contributions are as follows:

- **A Novel Taxonomy.** We introduce a structured taxonomy that categorizes existing works based on the types of information utilized during the learning process, providing a systematic framework for navigating the field. An overview of these approaches is presented in Fig. 2.
- **A Comprehensive Review.** Leveraging the proposed taxonomy, we conduct a thorough analysis of recent advancements in TIGRL, offering an in-depth summary of representative methods along with their respective strengths and limitations. Furthermore, we curate publicly available datasets and benchmarks to facilitate ongoing research and practical applications.
- **Future Directions.** We highlight persisting challenges in the field and propose potential research avenues, including emerging topics such as prompt-based technologies and the integration of foundational models.

2 Preliminary

In this section, we introduce the fundamental concepts and notation used throughout the paper.

Definitions. A **Temporal Interaction Graph (TIG)** is a dynamic graph that models sequences of interaction events, each occurring at a specific timestamp, while also accounting for the insertion and deletion of nodes and edges. Formally, a TIG observed over a time interval Γ is defined as a tuple $G_\Gamma = (V_\Gamma, E_\Gamma)$, where $V_\Gamma = \{v_1, v_2, \dots, v_N\}$ represents the set of nodes observed during Γ . $E_\Gamma = \{(v_i, v_j, w_{ij}, t_e) \mid v_i, v_j \in V_\Gamma, w_{ij} \in \mathbb{R}, t_e \in \Gamma\}$ denotes the set of edges, where each edge (v_i, v_j) is associated with a real-valued weight w_{ij} and a timestamp t_e . This definition captures the dynamic nature of the graph’s topology, where the intervals between successive timestamps can be irregular, allowing for a detailed yet efficient representation of temporal interactions.

Remark. It is important to highlight a subclass of temporal graphs in which the graph structure remains static while node

attributes evolve over time. These graphs, often referred to as *spatio-temporal graphs*, are outside the scope of this paper. Additionally, in the literature, TIGs [Chen *et al.*, 2023] are sometimes referred to as continuous-time dynamic graphs or event-based temporal graphs. TIGs can be easily converted into discrete-time dynamic graphs [Jiao *et al.*, 2024], which we consider a special case. However, as demonstrated in [Souza *et al.*, 2022], TIGs are more general and expressive, and thus will not be discussed.

3 Proposed Taxonomy

The dynamic nature of TIGs, characterized by continuously evolving node behaviors and interaction patterns, creates intricate structural relationships and diverse temporal dependencies that complicate representation learning. To address these challenges, we introduce a taxonomy classifying existing TIGRL methods into three paradigms based on their core information utilization strategies: (1) **Structure-Oriented Approaches**, which prioritize evolving topological preservation through structural elements like node proximity, subgraph dynamics, and multi-scale motifs; (2) **Temporal-Oriented Approaches**, which model time-dependent interactions through chronological sequences, temporal event intervals, and trajectory evolution encoding; and (3) **Application-Oriented Approaches**, which adapt learning frameworks to domain-specific objectives such as recommendation systems or anomaly detection. We conduct a comprehensive methodological comparison in Table 1, including paradigm characteristics, model techniques, task applicability, and sources.

3.1 Structure-Oriented Approaches

Structure-oriented approaches in TIGRL focus on preserving the evolving graph topology and capturing dynamic structural patterns. These methods are designed to handle continuously changing neighborhood configurations. Broadly, current methodologies are divided into two paradigms: (1) temporal random walk-based approaches and (2) temporal neighborhood-based approaches.

Temporal Random Walk-Based

Temporal random walk-based approaches capture spatiotemporal patterns by generating node sequences constrained by chronological order. A temporal random walk W is defined as a sequence $W = [(v_0, t_0), (v_1, t_1), \dots, (v_k, t_k)]$, where timestamps follow certain constraints, ensuring that each transition (v_i, v_{i+1}, t_{i+1}) represents a valid temporal edge in the TIG. The node representation h_u^t is computed by aggregating multiple walk instances:

$$h_u^t = \frac{1}{M} \sum_{i=1}^M f_{\text{Enc}}(W_i), \quad (1)$$

where M is the number of walks, and $f_{\text{Enc}}(\cdot)$ is an encoding function implemented through sequential models like recurrent neural networks (RNNs) or transformers. Early approaches, such as CTDNE [Nguyen *et al.*, 2018], prioritize temporal coherence during walk generation, ensuring the preservation of the chronological order of interactions. CAWs [Wang *et al.*, 2021] refine this approach by incorporating causal anonymization, which enhances structural induction and captures temporal causality. NeurTWs [Jin *et al.*, 2022] apply neural message

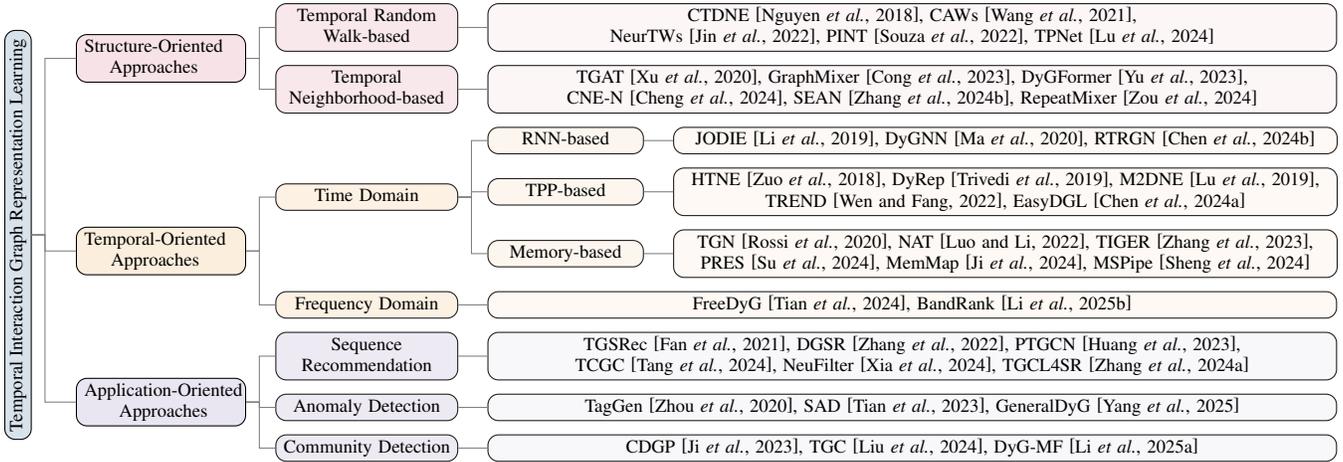


Figure 2: A taxonomy of Temporal Interaction Graph Representation Learning (TIGRL).

passing to anonymized walk sequences, while TPNet [Lu *et al.*, 2024] integrates temporal decay effects into the walk matrices, allowing for adaptive neighborhood weighting.

Despite their theoretical advantages in capturing temporal dependencies, these methods face practical limitations in computational efficiency due to the combinatorial explosion of valid temporal walks, particularly in dense graphs with long interaction histories. The inherent trade-off between walk length and temporal resolution further complicates parameter tuning, suggesting opportunities for future research in adaptive walk sampling strategies.

Temporal Neighborhood-Based

Temporal neighborhood-based approaches derive node representations by aggregating information from dynamically evolving neighborhoods within defined time windows. The fundamental formulation for a node u 's embedding h_u^t at time t integrates weighted neighborhood features and temporal-edge attributes:

$$h_u^t = f_{\text{enc}} \left(\sum_{i \in \mathcal{N}_u^t} \alpha_{ui} x_i^N, x_i^E \right), \quad (2)$$

where \mathcal{N}_u^t denotes the temporal neighborhood, α_{ui} adaptively weights neighbor i 's influence, x_i^N represents node-level features or relational metrics, and x_i^E encodes temporal interaction properties. The encoder $f_{\text{enc}}(\cdot)$ distills these into low-dimensional node embedding. TGAT [Xu *et al.*, 2020] fuses temporal decay with attention mechanisms, DyGFormer [Yu *et al.*, 2023] tokenizes historical interactions for transformer-based sequence modeling, and CNE-N [Cheng *et al.*, 2024] prioritizes co-neighbors for efficient dynamic link prediction. SEAN [Zhang *et al.*, 2024b] and RepeatMixer [Zou *et al.*, 2024] further refine α_{ui} via adaptive neighborhood scaling and periodic pattern detection.

Despite excelling in temporal link prediction through neighborhood similarity modeling, these methods struggle with node-centric tasks like classification due to their relational bias and reliance on fixed time windows. This design trade-off enhances computational efficiency but limits long-term dependency modeling, balancing temporal granularity against representational complexity.

3.2 Temporal-Oriented Approaches

Beyond topology, the temporal dimension is critical in TIGRL. Temporal-oriented approaches specifically address the challenges associated with irregular time intervals and the non-linear evolution of graph structures. Unlike conventional static graph embedding techniques that often disregard temporal variance through uniform projections, these methods embrace the irregularity of temporal signals. To achieve this, temporal-oriented approaches are further categorized into: (1) Time-Domain Approaches and (2) Frequency-Domain Approaches.

Time Domain

Based on the technique employed to model temporal dependencies, we categorize these approaches into three primary classes: recurrent neural network-based (RNN-based), temporal point process-based (TPP-based), and memory-based. Each category offers distinct advantages in capturing temporal patterns across diverse application scenarios.

RNN-based. RNN-based approaches for temporal interaction graphs dynamically update node representations through sequential processing of interactions, synthesizing historical states and neighborhood context using recurrent architectures. The core mechanism can be expressed through a unified framework where a node's hidden state h_u^t evolves via:

$$h_u^t = \text{RNN} \left(h_u^{t-}, \sum_{i \in \mathcal{N}_u^t} \alpha_{ui} \cdot h_i^{t-} \right), \quad (3)$$

combining a node's memory h_u^{t-} with aggregated neighborhood information from temporal neighbors \mathcal{N}_u^t , where attention weights α_{ui} modulate influence distribution. JODIE [Li *et al.*, 2019] implements dual RNNs that bidirectionally update user and item embeddings, capturing mutual evolution patterns between interacting entities. RTRGN [Chen *et al.*, 2024b] enhances this foundation by integrating historical neighbor states with current interactions via a Time Revision module (directly corresponding to the neighborhood aggregation term) and refining attention weights α_{ui} through a bias-correcting mechanism during recursive aggregation.

Taxonomy	Methods	Learning Paradigms	Inference Setting	Technique	Task	Sources	Code
Structure-Oriented Approaches	CTDNE	Supervised Learning	Transductive	Temporal Walks	Node, Link	WWW 2018	link
	HNIP	Supervised Learning	Transductive	Temporal Walks	Node, Link, Graph	AAAI 2020	-
	TGAT	Supervised Learning	Inductive	GAT and GNNs	Node, Link	ICLR 2020	link
	CAWs	Supervised Learning	Inductive	Temporal Walks and GNNs	Link	ICLR 2021	link
	NeurTWs	Unsupervised Learning	Inductive	Temporal Walks and GNN	Node, Link	NeurIPS 2022	link
	PINT	Supervised Learning	Inductive	Temporal Walks and GNNs	Node, Link	NeurIPS 2022	link
	GraphMixer	Supervised Learning	Inductive	MLP and GNNs	Node, Link	ICLR 2023	link
	DyGFormer	Supervised Learning	Inductive	Transformer and GNNs	Node, Link	NeurIPS 2023	link
	TPNet	Supervised Learning	Inductive	Temporal Walks and GNNs	Link	NeurIPS 2024	link
	CNE-N	Supervised Learning	Inductive	Hashtable-based Memory and GNNs	Link	SIGKDD 2024	link
	SEAN	Supervised Learning	Inductive	RNN and GNNs	Node, Link	SIGKDD 2024	-
	RepeatMixer	Supervised Learning	Inductive	MLP and GNNs	Link	SIGKDD 2024	link
Temporal-Oriented Approaches	HTNE	Supervised Learning	Transductive	Hawkes process	Node, Link, Graph	SIGKDD 2018	-
	JODIE	Supervised Learning	Inductive	GNNs and RNNs	Node, Link	SIGKDD 2019	link
	DyREP	Supervised Learning	Inductive	TPPs and GAT	Node, Link	ICLR 2019	link
	M2DNE	Supervised Learning	Transductive	TPPs and GNNs	Node, Link, Graph	ICLR 2019	link
	DyGNN	Supervised Learning	Inductive	RNN and GNNs	Node, Link	SIGIR 2020	link
	TGN	Supervised Learning	Inductive	Memory and GNNs	Node, Link	ICML 2020	link
	NAT	Supervised Learning	Inductive	RNNs and GNNs	Node, Link	LOG 2022	link
	TREND	Supervised Learning	Inductive	TPPs and GNNs	Node, Link	WWW 2022	link
	TIGER	Supervised Learning	Inductive	Memory and GNNs	Link	WWW 2023	link
	FreeDyG	Supervised Learning	Inductive	Frequency MLP and GNNs	Node, Link	ICLR 2024	link
	PRES	Supervised Learning	Inductive	Memory and GNNs	Node, Link	ICLR 2024	link
	RTRGN	Supervised Learning	Inductive	RNNs and GNNs	Link	NeurIPS 2024	-
	EasyDGL	Supervised Learning	Inductive	TPPs and GNNs	Node, Link	TPAMI 2024	link
	MemMap	Supervised Learning	Inductive	Memory and GNNs	Node, Link	SIGKDD 2024	link
	MSPipe	Supervised Learning	Inductive	Memory and GNNs	Node, Link	SIGKDD 2024	link
	BandRank	Supervised Learning	Inductive	Band-pass Disentangle with GNNs	Node, Link	WWW 2025	-
Application-Oriented Approaches	TagGen	Supervised Learning	Inductive	Temporal Walks and GAT	Link, Detection	SIGKDD 2020	link
	APAN	Supervised Learning	Inductive	GNNs and RNNs	Node, Link	SIGMOD 2021	link
	TGSRec	Supervised Learning	Inductive	Transformer and GNNs	Recommendation	CIKM 2021	link
	DGSR	Supervised Learning	Inductive	Collaborative Signals with GNNs	Recommendation	TKDE 2022	link
	CDGP	Unsupervised Learning	Inductive	Memory and GNNs	Node, Detection	SIGKDD 2023	link
	SAD	Unsupervised Learning	Inductive	Memory and GNNs	Node, Detection	IJCAI 2023	link
	PTGCN	Supervised Learning	Inductive	GCN and GAT	Recommendation	TOIS 2023	link
	TCGC	Supervised Learning	Inductive	Collaborative Information with GNNs	Recommendation	TOIS 2024	-
	NeuFilter	Supervised Learning	Inductive	Kalman Filtering and GNNs	Recommendation	WSDM 2024	link
	TGCL4SR	Unsupervised Learning	Inductive	GCL and GCN	Recommendation	AAAI 2024	-
	TGC	Unsupervised Learning	Inductive	Graph Clustering with GNNs	Node, Graph, Detection	ICLR 2024	link
	GeneralDyG	Supervised Learning	Inductive	Transformer and GNNs	Node, Detection	AAAI 2025	link
	DyG-MF	Unsupervised Learning	Inductive	Matrix Factorization and GNNs	Node, Graph, Detection	WWW 2025	link

Table 1: A summary of TIGRL models, ordered by their release time. Acronyms in Task: Node refers to node-level tasks; Link refers to link-level tasks; Graph refers to graph-level tasks; Recommendation refers to Graph Sequence Recommendation; Detection refers to Anomaly Detection or Community Detection.

While these methods effectively preserve temporal dependencies through RNN mechanisms, they inherit challenges including computational demands from sequential processing and scalability limitations for large sparse graphs.

TPP-based. TPP modeling in TIGs focuses on continuous-time interactions, using conditional intensity functions to capture both historical influence decay and spontaneous events. The Hawkes process, a variant of TPP, formalizes this dynamic through the equation:

$$\lambda(t) = \mu(t) + \int_{-\infty}^t \kappa(t-s) dn(s), \quad (4)$$

where $\mu(t)$ represents baseline interactions, and $\kappa(t-s)$ encodes temporal dependencies between events. DyRep [Trivedi *et al.*, 2019] exemplifies a deep architectural approach by coupling dual-intensity functions with a time-aware RNN. Here, $\lambda(t)$ governs the dynamics of interactions, while temporal attention computes $\kappa(t-s)$ as $\exp(-\delta(t-s))$, emphasizing recent events. HTNE [Zuo *et al.*, 2018] extends this paradigm by modeling neighborhood sequences as Hawkes processes, with trainable exponential kernels that capture historical neighbor influences. TREND [Wen and Fang, 2022] advances the framework by decoupling intensity into event-level and node-level components, using coupled graph neural networks (GNNs) to

model both localized interactions and global trend propagation. EasyDGL [Chen *et al.*, 2024a] introduces curriculum learning, adaptively weighting $\mu(t)$ to prioritize recent events during training through TPP-based attention gates.

These methods share common strengths, such as handling asynchronous events and modeling temporal decay patterns. However, challenges remain in scaling to large graphs and addressing sparse interaction scenarios.

Memory-based. Memory-based approaches in temporal interaction graphs focus on updating node embeddings dynamically to capture temporal dependencies via event-driven memory updates. These methods follow a unified computational paradigm, consisting of three key components: message generation, memory update, and embedding computation. The process is mathematically represented as:

$$\begin{aligned} m_u^t &= \text{msg}(s_u^{t-}, s_v^{t-}, e_{uv}^t, t), \\ s_u^t &= \text{mem}(s_u^{t-}, m_u^t), \\ h_u^t &= \text{emb}(s_u^t, \mathcal{N}_u^t), \end{aligned} \quad (5)$$

where s_u^{t-}, s_v^{t-} are the memory states of nodes u and v before time t , m_u^t is the message generated from the event e_{uv}^t , \mathcal{N}_u^t represents the temporal neighbors of u up to time t , and h_u^t is the dynamic embedding of node u at time t . The functions

$\text{msg}(\cdot)$, $\text{mem}(\cdot)$, $\text{emb}(\cdot)$ are learnable modules that generate the message, update memory, and compute the embedding, respectively. TGN [Rossi *et al.*, 2020] introduces RNN-based memory updates triggered by streaming events. NAT [Luo and Li, 2022] improves computational efficiency with GPU-friendly dictionary structures for neighborhood representation. TIGER [Zhang *et al.*, 2023] addresses long-term dependency loss by incorporating dual memory modules, while PRES [Su *et al.*, 2024] uses Gaussian mixture models for memory correction through time-series gradient analysis. MSPipe [Sheng *et al.*, 2024] advances parallel training by introducing controlled memory staleness, helping mitigate the temporal dependency bottleneck.

These methods excel at preserving temporal patterns through continuous memory updates and enabling real-time inference. However, they also share common challenges, including high memory storage requirements for long interaction histories, sensitivity to the event processing order, and potential error accumulation in memory states. A key design trade-off in this domain remains balancing memory fidelity with computational efficiency.

Frequency Domain

The frequency domain decomposes temporal signals into constituent frequency components through mathematical transformations like Fourier analysis, providing a powerful paradigm to capture periodic patterns and long-range dependencies that are challenging to model in time-domain approaches. Building on this foundation, recent advances in TIGRL have developed innovative frequency-enhanced architectures. FreeDyG [Tian *et al.*, 2024] employs a frequency-enhanced MLP-Mixer layer that first transforms node embeddings into the frequency space via Fast Fourier Transform (FFT), then applies learnable complex-valued filters to amplify task-relevant frequencies before reconstructing time-domain features through inverse FFT. This dual-domain architecture effectively captures periodic interaction patterns and temporal shift phenomena in dynamic graphs. BandRank [Li *et al.*, 2025b] addresses frequency aliasing through parallel frequency-specific MLPs that extract multi-scale temporal features ranging from long-term evolutionary trends to short-term fluctuations, complemented by a harmonic ranking loss that mitigates gradient vanishing in listwise optimization.

While these methods demonstrate superior performance in capturing cyclical patterns compared to conventional time-domain approaches, they introduce computational overhead from frequency transformations and require careful handling of phase information during inverse transforms. The frequency-domain paradigm shows particular promise for applications with inherent periodicity like social network dynamics and traffic prediction, though its effectiveness on irregular temporal patterns warrants further investigation.

3.3 Application-Oriented Approaches

TIGRL has also been closely integrated with some specific applications, e.g., nodes, links, graphs that do not vary enough to fit a specific application. Application-oriented approaches need to be carefully considered in how to incorporate auxiliary information—such as domain knowledge or contextual attributes—to enhance representation learning for targeted

tasks. Application-oriented approaches are further categorized into: (1) Sequence recommendation, (2) Anomaly Detection, and (3) Community Detection.

Sequence Recommendation

The integration of TIGs into sequence recommendation systems has modeling of dynamic user-item interactions and contextual attributes. Unlike static graphs, temporal graphs encode time-stamped interactions (e.g., purchases) and auxiliary relationships (e.g., social ties), enabling fine-grained analysis of behavioral shifts, interest evolution, and activity bursts. By transforming recommendation tasks into temporal link prediction problems, methods like TIGRL effectively reduce dimensionality while preserving heterogeneous features and temporal dependencies. Recent advances address these challenges through diverse approaches: TGSRec [Fan *et al.*, 2021] employs Transformer-based continuous-time representations to handle sparse, irregular data, while DGSR [Zhang *et al.*, 2022] integrates temporal attention with graph convolution for dynamic preference modeling. PTGCN [Huang *et al.*, 2023] enhances adaptability through positional encoding and time-aware features, whereas TCGC [Tang *et al.*, 2024] jointly models interaction evolution via collaborative filtering and graph co-evolution. The TGCL4SR [Zhang *et al.*, 2024a] framework further advances temporal pattern extraction through contrastive learning, demonstrating robust noise resistance and multi-scale behavioral analysis. Collectively, these methods address critical challenges including data sparsity, temporal dynamics, and noise in recommendation systems.

Anomaly Detection

Anomaly Detection identifies deviations from normal behavioral patterns in TIGs by leveraging their temporal dynamics and structural evolution. Unlike static methods, TIG-based approaches exploit time-dependent node/edge variations to detect subtle anomalies like fraudulent transactions in financial networks or malicious activities in social platforms. A common framework integrates temporal graph encoders with anomaly scoring modules, where node representations are dynamically updated to reflect evolving interactions, and deviations are measured through reconstruction errors or outlier scoring. SAD [Tian *et al.*, 2023] advances semi-supervised detection via contrastive learning on temporal subgraphs, reducing dependency on labeled anomalies. GeneralDyG [Yang *et al.*, 2025] enhances generalizability through temporal ego-graph sampling and hybrid GNN-Transformer architectures, addressing feature diversity across domains. While these methods excel in capturing temporal irregularities, they often face challenges in balancing detection sensitivity with computational efficiency, particularly in streaming scenarios. Key limitations include limited robustness to adversarial temporal perturbations and reliance on predefined temporal windows.

Community Detection

Community detection in TIGs analyzes topological patterns and temporal dependencies to uncover evolving network structures, capturing dynamic behaviors like community splitting, merging, and membership migration through continuous interaction updates. Unlike static methods, TIG-based approaches enable incremental updates that balance efficiency with temporal coherence while adapting to structural shifts. Recent

technical advances employ diverse strategies: CDGP [Ji *et al.*, 2023] dynamically integrates temporal-structural signals via graph propagation for evolution tracking, while TGC [Liu *et al.*, 2024] adjusts deep clustering techniques to suit the interaction sequence-based batch-processing pattern of TIG. DyG-MF [Li *et al.*, 2025a] reformulates dynamic matrix factorization using temporal latent factors to detect gradual community changes. Despite progress, these methods struggle with abrupt structural shifts due to temporal pattern over smoothing in long sequences, computational bottlenecks from continuous updates, and sensitivity to interaction sparsity. Current solutions partially address these through hybrid architectures like residual compression and multi-task learning frameworks, yet fundamental trade-offs persist between temporal resolution and scalability.

4 Frontier Technology for TIG

This section presents frontier technologies for TIG, including generative learning, prompt learning, and foundational models, which are pivotal advancements shaping the future of TIGRL.

4.1 Generative Learning

Generative learning for TIGs focuses on modeling data distributions to synthesize instances preserving both structural and temporal dynamics inherent in evolving entity interactions. Traditional frameworks like VAEs, GANs, and diffusion models have been adapted to TIGs by integrating temporal mechanisms such as TPP and RNN, enabling the generation of chronologically valid interaction sequences with accurate topological patterns. Recent advancements employ hybrid architectures to address the interplay between temporal coherence and structural fidelity. For example, TG-GAN [Zhang *et al.*, 2021] utilizes a continuous-time GAN with recurrent generators to produce timestamped edge sequences while maintaining chronological constraints, whereas TIGGER [Gupta *et al.*, 2022] combines TPPs with autoregressive decoders to capture interaction bursts and long-term dependencies. MTM [Liu and Sariyüce, 2023] introduces higher-order structural motifs (e.g., triadic closures) to guide generation, preserving local-global dynamics through transformer-based temporal embeddings. These models successfully replicate critical metrics like degree distributions, temporal motif frequencies, and inter-event intervals in real-world datasets, demonstrating their capacity to generate realistic TIGs despite challenges in modeling irregularly sampled or continuous-time interactions.

4.2 Prompt Learning

Prompt learning, initially rooted in natural language processing for adapting pre-trained models via task-specific prompts without modifying core parameters, has been extended to TIGs to address dynamic and time-sensitive relational patterns. Unlike static graphs, TIGs require prompts to encapsulate structural dependencies and temporal dynamics, such as evolving user-item interactions in e-commerce or social network communications over time. Recent advancements integrate temporal encoding mechanisms—like dynamic context-aware templates and node-time conditional prompts—to align pre-training objectives with time-dependent tasks. For example, DyGPrompt [Yu *et al.*, 2024] introduces prompts conditioned on node attributes and timestamps, dynamically adjusting

temporal neighborhood focus during message propagation, while TIGPrompt [Chen *et al.*, 2025] employs task-specific tokens to modulate attention to temporal granularity (e.g., short-term vs. long-term dependencies) through joint learning with positional encodings. By freezing backbone networks and tuning lightweight prompts, these methods achieve efficiency and scalability for large-scale TIGs with continuous temporal streams, balancing adaptability with computational constraints. This paradigm shift underscores the potential of temporal prompt engineering in bridging static model architectures with dynamic real-world interaction patterns, offering insights for both graph mining and broader AI communities.

4.3 Foundation Model

Foundation models have transformed machine learning by learning transferable representations through large-scale pre-training, with emerging applications in TIGs that model evolving relationships in domains like social networks and financial transactions. These models unify temporal dependencies, structural evolution, and contextual patterns into a universal framework through self-supervised learning on massive dynamic graph datasets. While traditional approaches rely on task-specific architectures, recent TIG-oriented foundation models integrate temporal-aware attention mechanisms, continuous-time encoding, and multi-scale aggregation to handle heterogeneous scenarios ranging from recommendation systems to epidemic prediction. A key innovation involves adapting large language models (LLMs) for temporal graph reasoning: methods like LLM4DyG [Zhang *et al.*, 2024c] encode temporal interactions and node attributes into text prompts, enabling LLMs to predict future states via text-generation paradigms. Hybrid architectures also combine pre-trained temporal graph encoders with task-specific adapters, using techniques like time-decayed attention to prioritize recent interactions while maintaining long-term dependencies. However, most current solutions repurpose existing frameworks like transformers rather than building architectures with native support for TIG inductive biases, creating tension between computational efficiency and the inherent complexity of temporal-structural patterns in TIGs.

5 Datasets and Benchmarks

This section introduces commonly used datasets and benchmarks in TIGRL, which are crucial for validating and comparing algorithms to drive research progress in this field.

Datasets	Domain	$ V $	$ E $	$ T $	Sizes	Sources
Wikipedia	Social	9,227	157,474	152,757	534MB	link
Reddit	Social	10,984	672,447	669,065	2.2GB	link
Enron	Social	184	125,235	22,632	35MB	link
UCI	Social	899	33,720	58,911	668KB	link
ML25M	Social	221,588	25,000,095	866,346	647MB	link
GDELT	Social	16,682	191,290,882	170,522	82.3GB	link
MOOC	Interaction	7,144	411,749	345,600	40MB	link
LastFM	Interaction	1,980	1,293,103	1,283,614	37MB	link
Social Evolution	Proximity	74	2,099,520	565,932	148MB	link
Contact	Proximity	692	2,426,279	8,065	160MB	link
Flights	Transport	13,169	1,927,145	122	32MB	link
Bitcoin-OTC	Economics	5,881	35,592	35,592	988KB	link
Bitcoin-Alpha	Economics	3,782	24,186	24,186	492KB	link
DGraphFin	Economics	4,889,537	4,300,999	1,433	649MB	link

Table 2: Dataset statistics

5.1 Datasets

High-quality datasets form the cornerstone of TIG research by enabling rigorous algorithm validation and comparative analysis across diverse scenarios. As summarized in Table 2, publicly available TIG datasets exhibit significant variations in scale, ranging from compact networks like UCI (thousands of nodes/edges) to massive real-world systems such as ML25M (millions of interactions). This scalability spectrum supports both granular pattern discovery and large-scale industrial applications. The datasets span multiple domains including social dynamics, transportation flows, and economic transactions, demonstrating TIG’s broad applicability in modeling temporal relational systems. Notably, these resources primarily capture structural evolution patterns while lacking rich node/link attributes—a limitation stemming from practical data collection challenges. This constraint paradoxically encourages research into generalizable models that infer latent features from temporal topology changes, potentially enhancing robustness across application scenarios.

Benchmark	Code	Year
PyTorch Geometric Temporal	https://github.com/benedekrozemberczki/pytorch_geometric_temporal	2021
TGL	https://github.com/amazon-science/tgl	2022
Dynamic Graph Library	https://github.com/yule-BUAA/DyGLib	2023
Temporal Graph Benchmark	https://github.com/shenyangHuang/TGB	2024
BenchTemp	https://github.com/qianghuangwhu/benchtemp	2024
Dynamic Graph Benchmark	https://github.com/gravins/dynamic_graph_benchmark	2024
BenchTGNN	https://github.com/Yang-yuxin/BenchTGNN	2024

Table 3: Links of benchmarks

5.2 Benchmarks

Open-source benchmarks are essential for advancing temporal interaction graph research by providing standardized evaluation protocols that facilitate reproducible research and systematic performance comparisons. Table 3 catalogs several key resources: PyTorch Geometric Temporal offers lightweight temporal GNN implementations optimized for small-scale prototypes, while TGL overcomes scalability challenges with distributed training architectures for billion-edge graphs. The Temporal Graph Benchmark establishes multi-domain evaluation standards across more than ten datasets with rigorous temporal integrity checks, BenchTGNN reveals critical design patterns in temporal message passing and neighborhood sampling through ablation studies, and the Dynamic Graph Benchmark focuses on industrial-grade metrics to monitor concept drift robustness in streaming scenarios; collectively, these benchmarks address diverse aspects of temporal graph learning and provide a comprehensive toolkit for the field.

6 Conclusion and Future Directions

In this survey, we provide a comprehensive overview of recent advancements in TIGRL. We introduce a novel taxonomy that categorizes existing research from multiple perspectives. Our analysis reveals that the majority of current efforts focus on addressing node correlation, capturing temporal patterns, enhancing robustness to noise, and scaling to large-scale networks. Despite these significant advancements, TIGRL faces several challenges and presents a variety of promising directions for future research:

More Complex TIGs. The development of TIG learning is constrained by critical gaps in dataset construction and evaluation practices. While recent advances have produced

large-scale dynamic graph datasets, their structural and complexity lag far behind static counterparts, as evidenced by EdgeBank’s [Poursafaei *et al.*, 2022] competitive performance only using simple historical interaction patterns. This highlights the inadequacy of existing benchmarks in distinguishing sophisticated models, necessitating datasets with multi-scale hierarchies, realistic evolution dynamics (e.g., edge deletions, attribute shifts), and specialized graph types like temporal signed networks and heterogeneous graphs for applications such as financial fraud detection. Current repositories often neglect crucial real-world temporal operations and multidimensional features, limiting the comprehensive evaluation of model robustness. The research community needs to prioritize creating annotated TIG benchmarks that integrate heterogeneous interaction modalities, temporal granularity, and domain-specific topological patterns. Such efforts will enable rigorous validation of spatiotemporal dependency modeling and advance the practical deployment of TIG models across industries requiring dynamic relational reasoning.

Explainability and Interpretability. Explainability and interpretability are critical for enhancing trust and transparency in TIG models, particularly in dynamic, time-sensitive domains such as finance, healthcare, and social networks. These properties aim to clarify how models infer temporal patterns, attribute importance to interactions, and rationalize predictions in evolving graph structures. Recent studies have proposed tailored methods to address TIG-specific challenges. For instance, the T-GNNExplainer [Xia *et al.*, 2023] decomposes temporal reasoning into exploration (identifying influential temporal paths) and navigation (highlighting key interactions), enabling stepwise explanations for dynamic graph predictions. Another approach, TempME [Chen and Ying, 2023], leverages motif discovery to uncover recurring temporal subgraph patterns that drive model decisions, bridging low-level interactions and high-level predictions. These studies show that although the interpretability of TIGRL is a relatively new and less discussed area, there are already some studies attempting to address this issue.

Unified Evaluation Standards. While benchmarks like TGL and DyGLib have advanced TIG evaluation, current standards suffer from critical limitations in protocol design and task adaptability. The oversimplified evaluation for link prediction tasks, though achieving high performance on specific datasets, inadequately reflects real-world complexity where temporal interactions require multi-class classification, regression, or anomaly detection capabilities. This discrepancy stems from insufficient standardization in handling non-classification tasks, particularly the absence of unified normalization for continuous interaction intensities and inconsistent negative sampling strategies for anomaly detection. To enable fair cross-model comparisons, future frameworks must implement multi-dimensional metrics incorporating temporal dynamics and interaction intensity. Crucially, this requires establishing unified preprocessing pipelines, dynamic evaluation protocols simulating real-world streaming scenarios, and domain-specific benchmarks mirroring practical deployment conditions.

Acknowledgments

This work was supported in part by the Zhejiang Provincial Natural Science Foundation of China under Grant LDT23F01012F01, in part by the National Natural Science Foundation of China under Grants 62372146, 62422210, and 92370111.

References

- [Chen and Ying, 2023] Jialin Chen and Rex Ying. Tempme: Towards the explainability of temporal graph neural networks via motif discovery. *NeurIPS*, 36:29005–29028, 2023.
- [Chen *et al.*, 2023] Hongjiang Chen, Pengfei Jiao, Huijun Tang, and Huaming Wu. Temporal graph representation learning with adaptive augmentation contrastive. In *ECML PKDD*, pages 683–699, 2023.
- [Chen *et al.*, 2024a] Chao Chen, Haoyu Geng, Nianzu Yang, Xiaokang Yang, and Junchi Yan. Easydgl: Encode, train and interpret for continuous-time dynamic graph learning. *TPAMI*, 2024.
- [Chen *et al.*, 2024b] Yizhou Chen, Anxiang Zeng, Qingtao Yu, Kerui Zhang, Cao Yuanpeng, Kangle Wu, Guangda Huzhang, Han Yu, and Zhiming Zhou. Recurrent temporal revision graph networks. *NeurIPS*, 36, 2024.
- [Chen *et al.*, 2025] Xi Chen, Siwei Zhang, Yun Xiong, Xixi Wu, Jiawei Zhang, Xiangguo Sun, Yao Zhang, Feng Zhao, and Yulin Kang. Prompt learning on temporal interaction graphs. *ICLR*, 2025.
- [Cheng *et al.*, 2024] Ke Cheng, Peng Linzhi, Junchen Ye, Leilei Sun, and Bowen Du. Co-neighbor encoding schema: A light-cost structure encoding method for dynamic link prediction. In *SIGKDD*, pages 421–432, 2024.
- [Cong *et al.*, 2023] Weilin Cong, Si Zhang, Jian Kang, Baichuan Yuan, Hao Wu, Xin Zhou, Hanghang Tong, and Mehrdad Mahdavi. Do we really need complicated model architectures for temporal networks? In *ICLR*, 2023.
- [Fan *et al.*, 2021] Ziwei Fan, Zhiwei Liu, Jiawei Zhang, Yun Xiong, Lei Zheng, and Philip S Yu. Continuous-time sequential recommendation with temporal graph collaborative transformer. In *CIKM*, pages 433–442, 2021.
- [Gupta and Bedathur, 2022] Shubham Gupta and Srikanta Bedathur. A survey on temporal graph representation learning and generative modeling. *arXiv preprint arXiv:2208.12126*, 2022.
- [Gupta *et al.*, 2022] Shubham Gupta, Sahil Manchanda, Srikanta Bedathur, and Sayan Ranu. Tigger: Scalable generative modelling for temporal interaction graphs. In *AAAI*, volume 36, pages 6819–6828, 2022.
- [Huang *et al.*, 2023] Liwei Huang, Yutao Ma, Yanbo Liu, Bohong Danny Du, Shuliang Wang, and Deyi Li. Position-enhanced and time-aware graph convolutional network for sequential recommendations. *TOIS*, 41(1):1–32, 2023.
- [Ji *et al.*, 2023] Shuo Ji, Xiaodong Lu, Mingzhe Liu, Leilei Sun, Chuanren Liu, Bowen Du, and Hui Xiong. Community-based dynamic graph learning for popularity prediction. In *SIGKDD*, pages 930–940, 2023.
- [Ji *et al.*, 2024] Shuo Ji, Mingzhe Liu, Leilei Sun, Chuanren Liu, and Tongyu Zhu. Memmap: An adaptive and latent memory structure for dynamic graph learning. In *SIGKDD*, pages 1257–1268, 2024.
- [Jiao *et al.*, 2024] Pengfei Jiao, Hongjiang Chen, Huijun Tang, Qing Bao, Long Zhang, Zhidong Zhao, and Huaming Wu. Contrastive representation learning on dynamic networks. *Neural Networks*, 174:106240, 2024.
- [Jin *et al.*, 2022] Ming Jin, Yuan-Fang Li, and Shirui Pan. Neural temporal walks: Motif-aware representation learning on continuous-time dynamic graphs. *NeurIPS*, 35:19874–19886, 2022.
- [Kazemi *et al.*, 2020] Seyed Mehran Kazemi, Rishab Goel, Kshitij Jain, Ivan Kobyzev, Akshay Sethi, Peter Forsyth, and Pascal Poupart. Representation learning for dynamic graphs: A survey. *JMLR*, 21(1):2648–2720, 2020.
- [Li *et al.*, 2019] Jia Li, Zhichao Han, Hong Cheng, Jiao Su, Pengyun Wang, Jianfeng Zhang, and Lujia Pan. Predicting path failure in time-evolving graphs. In *Proceedings of KDD 2019*, pages 1279–1289, 2019.
- [Li *et al.*, 2025a] Dongyuan Li, Satoshi Kosugi, Ying Zhang, Manabu Okumura, Feng Xia, and Renhe Jiang. Revisiting dynamic graph clustering via matrix factorization. In *WWW*, 2025.
- [Li *et al.*, 2025b] Yingxuan Li, Yuanyuan Xu, Xuemin Lin, Wenjie Zhang, and Ying Zhang. Ranking on dynamic graphs: An effective and robust band-pass disentangled approach. In *WWW*, 2025.
- [Liu and Sariyüce, 2023] Penghang Liu and Ahmet Erdem Sariyüce. Using motif transitions for temporal graph generation. In *SIGKDD*, pages 1501–1511, 2023.
- [Liu *et al.*, 2024] Meng Liu, Yue Liu, KE LIANG, Wenxuan Tu, Siwei Wang, sihang zhou, and Xinwang Liu. Deep temporal graph clustering. In *ICLR*, 2024.
- [Longa *et al.*, 2023] Antonio Longa, Veronica Lachi, Gabriele Santin, Monica Bianchini, Bruno Lepri, Pietro Lio, franco scarselli, and Andrea Passerini. Graph neural networks for temporal graphs: State of the art, open challenges, and opportunities. *TMLR*, 2023.
- [Lu *et al.*, 2019] Yuanfu Lu, Xiao Wang, Chuan Shi, Philip S Yu, and Yanfang Ye. Temporal network embedding with micro-and macro-dynamics. In *CIKM*, pages 469–478, 2019.
- [Lu *et al.*, 2024] Xiaodong Lu, Leilei Sun, Tongyu Zhu, and Weifeng Lv. Improving temporal link prediction via temporal walk matrix projection. *NeurIPS*, 37, 2024.
- [Luo and Li, 2022] Yuhong Luo and Pan Li. Neighborhood-aware scalable temporal network representation learning. In *LOG*, pages 1–1, 2022.
- [Ma *et al.*, 2020] Yao Ma, Ziyi Guo, Zhaocun Ren, Jiliang Tang, and Dawei Yin. Streaming graph neural networks. In *SIGIR*, pages 719–728, 2020.

- [Nguyen *et al.*, 2018] Giang Hoang Nguyen, John Boaz Lee, Ryan A Rossi, Nesreen K Ahmed, Eunye Koh, and Sungchul Kim. Continuous-time dynamic network embeddings. In *WWW*, pages 969–976, 2018.
- [Poursafaei *et al.*, 2022] Farimah Poursafaei, Shenyang Huang, Kellin Pelrine, and Reihaneh Rabbany. Towards better evaluation for dynamic link prediction. *NeurIPS*, 35:32928–32941, 2022.
- [Qin and Yeung, 2023] Meng Qin and Dit-Yan Yeung. Temporal link prediction: A unified framework, taxonomy, and review. *CSUR*, 56(4):1–40, 2023.
- [Rossi *et al.*, 2020] Emanuele Rossi, Ben Chamberlain, Fabrizio Frasca, Davide Eynard, Federico Monti, and Michael Bronstein. Temporal graph networks for deep learning on dynamic graphs. In *ICML 2020 Workshop on Graph Representation Learning*, 2020.
- [Sheng *et al.*, 2024] Guangming Sheng, Junwei Su, Chao Huang, and Chuan Wu. Mspipe: Efficient temporal gnn training via staleness-aware pipeline. In *SIGKDD*, pages 2651–2662, 2024.
- [Souza *et al.*, 2022] Amauri Souza, Diego Mesquita, Samuel Kaski, and Vikas Garg. Provably expressive temporal graph networks. *NeurIPS*, 35:32257–32269, 2022.
- [Su *et al.*, 2024] Junwei Su, Difan Zou, and Chuan Wu. Pres: Toward scalable memory-based dynamic graph neural networks. In *ICLR*, 2024.
- [Tang *et al.*, 2024] Haoran Tang, Shiqing Wu, Xueyao Sun, Jun Zeng, Guandong Xu, and Qing Li. Tcgc: Temporal collaboration-aware graph co-evolution learning for dynamic recommendation. *TOIS*, 43(1):1–27, 2024.
- [Tian *et al.*, 2023] Sheng Tian, Jihai Dong, Jintang Li, Wenlong Zhao, Xiaolong Xu, Baokun Wang, Bowen Song, Changhua Meng, Tianyi Zhang, and Liang Chen. Sad: semi-supervised anomaly detection on dynamic graphs. In *IJCAI*, pages 2306–2314, 2023.
- [Tian *et al.*, 2024] Yuxing Tian, Yiyang Qi, and Fan Guo. Freedyg: Frequency enhanced continuous-time dynamic graph model for link prediction. In *ICLR*, 2024.
- [Trivedi *et al.*, 2019] Rakshit Trivedi, Mehrdad Farajtabar, Prasenjeet Biswal, and Hongyuan Zha. Dyrep: Learning representations over dynamic graphs. In *ICLR*, 2019.
- [Wang *et al.*, 2021] Yanbang Wang, Yen-Yu Chang, Yunyu Liu, Jure Leskovec, and Pan Li. Inductive representation learning in temporal networks via causal anonymous walks. In *ICLR*, 2021.
- [Wen and Fang, 2022] Zhihao Wen and Yuan Fang. Trend: Temporal event and node dynamics for graph representation learning. In *WWW*, pages 1159–1169, 2022.
- [Xia *et al.*, 2023] Wenwen Xia, Mincai Lai, Caihua Shan, Yao Zhang, Xinnan Dai, Xiang Li, and Dongsheng Li. Explaining temporal graph models through an explorer-navigator framework. In *ICLR*, 2023.
- [Xia *et al.*, 2024] Jiafeng Xia, Dongsheng Li, Hansu Gu, Tun Lu, Peng Zhang, Li Shang, and Ning Gu. Neural kalman filtering for robust temporal recommendation. In *WSDM 2024*, pages 836–845, 2024.
- [Xu *et al.*, 2020] Da Xu, Chuanwei Ruan, Evren Korpeoglu, Sushant Kumar, and Kannan Achan. Inductive representation learning on temporal graphs. In *ICLR*, 2020.
- [Yang *et al.*, 2025] Xiao Yang, Xuejiao Zhao, and Zhiqi Shen. A generalizable anomaly detection method in dynamic graphs. *AAAI*, 2025.
- [Yu *et al.*, 2023] Le Yu, Leilei Sun, Bowen Du, and Weifeng Lv. Towards better dynamic graph learning: New architecture and unified library. *NeurIPS*, 2023.
- [Yu *et al.*, 2024] Xingtong Yu, Zhenghao Liu, Yuan Fang, and Xinming Zhang. Dygprompt: Learning feature and time prompts on dynamic graphs. *arXiv preprint arXiv:2405.13937*, 2024.
- [Zhang *et al.*, 2021] Liming Zhang, Liang Zhao, Shan Qin, Dieter Pfoser, and Chen Ling. Tg-gan: Continuous-time temporal graph deep generative models with time-validity constraints. In *WWW*, pages 2104–2116, 2021.
- [Zhang *et al.*, 2022] Mengqi Zhang, Shu Wu, Xueli Yu, Qiang Liu, and Liang Wang. Dynamic graph neural networks for sequential recommendation. *TKDE*, 35(5):4741–4753, 2022.
- [Zhang *et al.*, 2023] Yao Zhang, Yun Xiong, Yongxiang Liao, Yiheng Sun, Yucheng Jin, Xuehao Zheng, and Yangyong Zhu. Tiger: Temporal interaction graph embedding with restarts. In *WWW*, pages 478–488, 2023.
- [Zhang *et al.*, 2024a] Shengzhe Zhang, Liyi Chen, Chao Wang, Shuangli Li, and Hui Xiong. Temporal graph contrastive learning for sequential recommendation. In *AAAI*, volume 38, pages 9359–9367, 2024.
- [Zhang *et al.*, 2024b] Siwei Zhang, Xi Chen, Yun Xiong, Xixi Wu, Yao Zhang, Yongrui Fu, Yinglong Zhao, and Jiawei Zhang. Towards adaptive neighborhood for advancing temporal interaction graph modeling. In *SIGKDD*, pages 4290–4301, 2024.
- [Zhang *et al.*, 2024c] Zeyang Zhang, Xin Wang, Ziwei Zhang, Haoyang Li, Yijian Qin, and Wenwu Zhu. Llm4dyg: can large language models solve spatial-temporal problems on dynamic graphs? In *SIGKDD*, pages 4350–4361, 2024.
- [Zheng *et al.*, 2025] Yanping Zheng, Lu Yi, and Zhewei Wei. A survey of dynamic graph neural networks. *Frontiers of Computer Science*, 19(6):1–18, 2025.
- [Zhou *et al.*, 2020] Dawei Zhou, Lecheng Zheng, Jiawei Han, and Jingrui He. A data-driven graph generative model for temporal interaction networks. In *SIGKDD*, 2020.
- [Zou *et al.*, 2024] Tao Zou, Yuhao Mao, Junchen Ye, and Bowen Du. Repeat-aware neighbor sampling for dynamic graph learning. In *SIGKDD*, pages 4722–4733, 2024.
- [Zuo *et al.*, 2018] Yuan Zuo, Guannan Liu, Hao Lin, Jia Guo, Xiaoqian Hu, and Junjie Wu. Embedding temporal network via neighborhood formation. In *SIGKDD*, pages 2857–2866, 2018.