# Bridging the Divide: End-to-End Sequence-Graph Learning

#### **Anonymous Author(s)**

Affiliation Address email

## **Abstract**

Many real—world datasets are both *sequential* and *relational*: each node carries an event sequence while edges encode interactions. Existing methods in sequence modeling and graph modeling often neglect one modality or the other. We argue that sequences and graphs are not separate problems but complementary facets of the same dataset, and should be learned jointly. We introduce BRIDGE, a unified end-to-end architecture that couples a sequence encoder with a GNN under a single objective, allowing gradients to flow across both modules and learning task-aligned representations. To enable fine-grained token-level message passing among neighbors, we add TOKENXATTN, a token-level cross-attention layer that passes messages between events in neighboring sequences. Across two settings, friendship prediction (Brightkite) and fraud detection (Amazon), BRIDGE consistently outperforms static GNNs, temporal graph methods, and sequence-only baselines on ranking and classification metrics.

## 1 Introduction

2

3

5

6

7

8

9

10

11 12

13

14

Modern machine learning increasingly faces settings where data are both sequential and relational. 15 We study the general case in which an entity may have multiple relation with other entities, where 17 each entity generates a time-stamped event sequence. We adopt user-event modeling as our running example to motivate design choices and evaluation, since social platforms epitomize this sequence-on-18 graph regime; the approach, however, applies broadly beyond this domain. On social platforms, each 19 user generates a sequence of personal events (e.g., logins, posts) while simultaneously maintaining 20 friendships with others. In e-commerce, a user may purchase items or write reviews as event 21 sequences, while being connected to others through co-purchasing or co-review relations. In Figure 1, 22 we illustrate the e-commerce setting: each user generates a personal sequence of events (e.g., logins, 23 page visits, purchases, reviews), while also engaging in relational events that connect them to other 25 users (e.g., sending or receiving gifts, commenting on reviews). Such data are inherently multimodal, combining temporal dynamics at the sequence level with structural dependencies across the network, 26 making it essential to model both modalities jointly for accurate prediction. 27

There have been efforts to incorporate temporal and sequential information into relational models, 28 i.e., graphs. For example, the temporal heterogeneous graph formulations, such as CTDG [19] and 29 its derived models [23, 28, 37], encode relational events between multiple entities with a timestamp 30 assigned to each relation. However, this formulation only works for events describing interactions with 31 other entities, such as transitive actions. User events, on the other hand, describe intransitive actions 32 (e.g., 'sign in', 'login', and 'subscribe') or status changes (e.g., 'payment successful', 'payment 33 34 declined', 'subscription status update') that involve only a single user. Forcing such events into a graph requires self-loops or introducing event nodes, resulting in a misleading graph where, for 35 instance, everyone who 'logs in' looks related through a shared login node. 36

When dealing with datasets that have both sequential and relational components, existing approaches fall short of fully capturing this interplay. On one hand, sequence models such as recurrent networks Submitted to 39th Conference on Neural Information Processing Systems (NeurIPS 2025). Do not distribute.

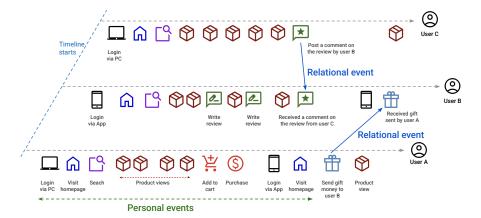


Figure 1: Illustration of data combining sequential and relational structure. Each entity is associated with a sequence of events, while edges capture interactions between entities. For example, in e-commerce, a user generates a sequence of purchases or reviews while also forming relations with other users through co-reviews or social interactions.

and Transformers excel at extracting temporal patterns within a sequence but ignore how sequences are connected. On the other hand, graph neural networks (GNNs) and temporal graph models focus on relational structure and evolving interactions, but typically compress each sequence into a single feature vector, discarding fine-grained temporal information. Spatial-temporal methods attempt to combine both, but often rely on the unrealistic assumption of synchronized time steps. As a result, current methods either capture sequential detail without relational context or relational structure without event-level granularity.

To bridge this gap, we propose a new perspective: treating each sequence not as a static node embedding but as a dynamic set of event-level representations that can communicate across the graph.

Our approach introduces two complementary components.

- At the *model level*, we design BRIDGE, an end-to-end architecture that integrates a sequential encoder with a GNN, to learn temporal and relational signals jointly.
- At the layer level, we develop TOKENXATTN, a token-wise cross-attention mechanism that
  allows individual events in one sequence to attend to events in neighboring sequences.

The two components together enable richer information exchange and preserve the temporal information of sequences while respecting the graph topology.

Across link prediction and fraud detection experiments, our approach consistently surpasses strong GNN and temporal-graph baselines by treating sequences and graphs as a single system. This joint view unlocks richer temporal-relational reasoning and points toward a new class of hybrid models for complex, multi-modal data.

#### 2 Related Works

49

50

51

52

59

66

67

68

69

70

71

72

This section reviews approaches that combine sequences and graphs. We leave additional related works in Appendix A. A growing body of work seeks to leverage sequence and graph information, motivated by applications where entities are connected but also evolve through temporal event sequences. This line of research spans two major directions: (i) spatial-temporal graph methods, which integrate time-series with graph structure, and (ii) temporal graph methods, which model the sequence as edges with timestamps.

**Spatial-Temporal Graph Methods** Spatial-temporal graph networks were originally designed for time-series prediction problems such as traffic forecasting. Models like STGCN [41], STGformer [32], MTGCN [36], and MST-GAT [7] encode spatial relations (e.g., road networks) and temporal patterns. In recommendation systems, hybrid approaches have also emerged. Shui et al. [25] models the relational graph among items while simultaneously capturing user consumption sequences. Similarly, Zhang et al. [42] treats users and items as different node types within a heterogeneous graph, enabling message passing between items (sequential) and users (relational). A common limitation of these methods is the assumption that all time series are synchronized and share identical time steps. By

contrast, our setting imposes no such constraint: sequences may be asynchronous, vary in length, and even be non-overlapping.

**Temporal Graph Methods** Temporal graph models such as TGN [23], TGAT [37], and DyRep [28], 76 attach timestamps to edges to model interactions without aligned time steps. However, a temporal 77 graph requires all events to be described as a relation between two entities/users. Personal events 78 that involve only a single entity cannot be easily represented by a temporal graph. A common 79 workaround is to introduce event nodes (e.g., representing "login" or "payment declined" as nodes) and connect them to the corresponding users, but this can create high-degree hubs (e.g., a single 81 82 "login" node connected to many users) and is often unnatural in a graph formulation. As argued by Bechler-Speicher et al. [1], forcing an unsuitable data modality into a graph form can fail to encode 83 important information and may not be meaningful. 84

# 3 Problem Setup

85

101

102

103

104

105

117

We consider graph-structured data in which each node is associated with a sequence; equivalently, this can be viewed as a collection of sequences that are not independent but connected through a graph. Formally, let  $G=(\mathcal{V},\mathcal{E})$  denote the graph with nodes  $\mathcal{V}$  and edges  $\mathcal{E}$ .  $\mathcal{V}=\{v_1,\ldots,v_N\}$  denote the set of N nodes. Each node  $v_i$  is associated with a sequence of elements,  $S_i=[e_1,\ldots,e_{M_i}]$ , of length  $M_i$ . Each element  $e_t$  may carry a timestamp, but since our method does not rely on absolute time, we omit timestamps for notational simplicity.

This dual structure presents a modeling challenge: the sequences are inherently temporal and suited for sequential models, while the graph edges capture dependencies among nodes. Neither sequential modeling nor graph-based modeling alone is sufficient, motivating the need for a joint approach that integrates both. In the remainder of this paper, we adopt terminology from user—event data (nodes representing users, sequence elements representing events) to illustrate the setting more concretely, although the formulation itself is general.

# 98 4 Methodology

To capture both sequential and relational structure in the general setting described above, our approach consists of two main components:

- 1. **Model level (Bridge):** an end-to-end architecture that integrates a sequential model with a graph neural network (GNN).
- 2. **Layer level (TOKENXATTN):** a token-wise cross-attention message-passing layer that enables tokens in a sequence to attend to tokens in neighboring sequences.

# 4.1 BRIDGE: End-to-End Model Architecture

Our architecture integrates sequential and graph components in an end-to-end framework. The sequential module first encodes each user's event sequence  $S_i$  into an embedding  $X_i \in \mathbb{R}^{M_i \times d}$ , which is then treated as the node feature input to the graph model. In standard message passing, the node feature is a d-dimensional vector; in such a case we compress  $X_i$  along the first dimension via e.g., mean or max pooling to obtain a node feature  $Z_i \in \mathbb{R}$ . The graph model performs message passing over the relational user graph, and its output is added back to the original sequence embedding via a residual connection, yielding updated node representations for downstream tasks.

Variants. Algorithm 1 specifies a family of sequence–graph models that we denote BRIDGE. The graph step can be instantiated with standard GNN layers (GCN, GAT, TransformerConv) or with the token-wise cross-attention layer introduced next. We refer to these variants as BRIDGE-GCN, BRIDGE-GAT, BRIDGE-TransformerConv, and BRIDGE-TOKENXATTN.

## 4.2 TOKENXATTN: Token-Level Cross-Attention

In the previous subsection, we mentioned that conventional GNN layers [15, 24, 31] assume each node is represented by a single feature vector. However, in our setting, each user corresponds to a sequence of events with inherent temporal structure. Compressing an entire sequence into one vector  $Z_i := \operatorname{Compress}(X_i) \in \mathbb{R}^d$ , whether by pooling or taking the final event's embedding, inevitably discards the temporal information. We propose keeping the full sequence embedding for message passing, i.e.,  $Z_i = X_i \in \mathbb{R}^{M_i \times d}$ , as different events in a neighbor's history have different effects on the target user (e.g., the most recent action is often most influential). To preserve this granularity,

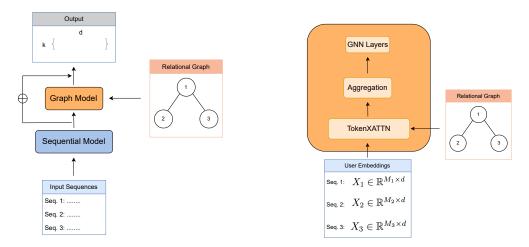


Figure 2: Illustration of our proposed model. (*Left*) BRIDGE, the end-to-end sequence–graph architecture, where the sequential module encodes node sequences and the resulting embeddings serve as node features for the graph module. (*Right*) TOKENXATTN, a zoomed-in view of the tokenwise cross-attention layer that enables event-level message passing between neighboring sequences.

# Algorithm 1 BRIDGE

```
Require: Graph G=(\mathcal{V},\mathcal{E}); sequences \{S_i\}_{i=1}^N; Number of message passing layers L Ensure: Final node representations \{h_i\}_{i=1}^N
 1: for v_i \in \mathcal{V} do
 2:
           X_i \leftarrow \text{SeqEncoder}(S_i)
 3: end for
 4: for \ell = 1, ..., L do
           for v_i \in \mathcal{V} do
 5:
                H_i^{(\ell)} \leftarrow \text{MessagePassing}(i, X_i, \{X_j\}_{j \in \mathcal{N}(i)})

⊳ GCN, GAT, TOKENXATTN, etc.

 6:
                X_i \leftarrow X_i + H_i^{(\ell)}
 7:
                                                                                                                                   end for
 8:
 9: end for
10: for v_i \in \mathcal{V} do
           h_i \leftarrow \text{vec}(\text{Compress}(X_i))
12: end for
```

we retain event-level embeddings: each user i is represented by a matrix  $X_i \in \mathbb{R}^{M_i \times d}$ , where  $M_i$  is the sequence length. Building on this representation, we introduce a token-wise cross-attention mechanism (Algorithm 2) that enables token-level message passing.

General neighbor aggregation. After computing the per-neighbor messages  $H_{i \leftarrow j}$ , the final step is to aggregate them into an updated representation for user i. We denote this step by a flexible aggregation function  $f_{\rm agg}$ , which combines neighbor messages using edge weights  $w_{ij}$  and an operator AGG  $\in$  {sum, mean}:

$$H_i = f_{\text{agg}}(\{H_{i \leftarrow j}\}_{j \in \mathcal{N}(i)}) = \text{AGG}_{j \in \mathcal{N}(i)}(w_{ij} H_{i \leftarrow j}), \qquad H_i \in \mathbb{R}^{M_i \times d_h}.$$

Different GNN variants define  $w_{ij}$  differently. For example, mean aggregation sets  $w_{ij} = \frac{1}{|\mathcal{N}(i)|}$ , GCN uses  $w_{ij} = \frac{1}{\sqrt{\deg(i)\deg(j)}}$ , and GAT learns  $w_{ij}$  through edge attention mechanisms.

Finally, to capture multi-hop dependencies, we stack standard GNN layers on top of TOKENXATTN, enabling information to propagate beyond immediate neighbors and reach higher-order neighbors.

# 5 Experiments

125

126

127

137

We validate our methods on two tasks, friendship prediction and fraud detection. For both experiments,we report the average over three runs.

## Algorithm 2 Token-Level Cross-Attention (TOKENXATTN)

```
Require: Node i; token matrix X_i \in \mathbb{R}^{M_i \times d}; neighbor tokens \{X_j\}_{j \in \mathcal{N}(i)}; weights
        W_O, W_K, W_V \in \mathbb{R}^{d \times d_h}; edge weights \{w_{ij}\}_{j \in \mathcal{N}(i)}; f_{\text{agg}} \in \{\text{sum}, \text{mean}\}
Ensure: H_i \in \mathbb{R}^{M_i \times d_h}
  1: Q_i \leftarrow X_i W_Q
  2: for all j \in \mathcal{N}(i) do
               K_j \leftarrow X_j W_K; \quad V_j \leftarrow X_j W_V

A_{i \leftarrow j} \leftarrow \operatorname{softmax}(Q_i K_j^{\top} / \sqrt{d_h})
                                                                                                                                                                              \triangleright \in \mathbb{R}^{M_i \times M_j}
                                                                                                                                                                               \triangleright \in \mathbb{R}^{M_i \times d_h}
  5:
                H_{i \leftarrow j} \leftarrow A_{i \leftarrow j} V_j
  6: end for
  7: if f_{\text{agg}} = \text{sum then}
               H_i \leftarrow \sum_{j \in \mathcal{N}(i)} w_{ij} H_{i \leftarrow j}
  9: else
                                                                                                                                                                           \triangleright f_{\text{agg}} = \text{mean}
               H_i \leftarrow \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} w_{ij} H_{i \leftarrow j}
11: end if
12: return H_i
```

## 5.1 Friendship Prediction

140

141

142

143

144

145

146

147

148

149

150

151

152

153

154

155

156 157

158

159

160 161

162 163

164

165

166

167

168

169

170

171

172

173

**Experimental Setup.** We perform friendship prediction on the Brightkite dataset [4, 16] and Amazon reviews [18]. Brightkite is a location-based social networking platform, and the dataset contains location check-ins and friendship relationships among Brightkite users. Each check-in is originally recorded in latitude and longitude, which we converted into Geohash-8 representation [20]. Example geohashes include 9v6kpmr1, gcpwkeq6, and u0yhxgm1, where nearby locations share similar prefixes. Each user has a sequence of geohash checkins, where we further split each geohash8 into 4 tokens, where the tokens represent increasingly higher resolution (smaller areas) of the location. For Amazon, we use the *Musical*, *Clothing*, and *Electronics* categories. Each user's sequence consists of product IDs and ratings; the "friendship" graph is constructed via co-review ties—two users are connected if they have co-reviewed at least three products.

#### **Compared Methods.** We compare three categories of baselines with BRIDGE:

- 1. *Graph Only models.* GCN [15], GAT [31], and TransformerConv [24] operate on a static user friendship graph using single-vector node representations, ignoring the user sequences.
- 2. *Temporal Graph models*. TGN [23], DyRep [28], and TNCN [44] are dynamic graph learning methods that model event sequences as timestamped edges. Here, we convert each Geohash8 into a node (we do not perform a tokenized split of the geohash) and connect a user to the geohash node every time they check in. The friendship info does not have a timestamp, so we randomly add a time to it within each user's check-in timeline.
- 3. *Two-Stage Training*. GCN + S, GAT + S, and TransformerConv + S extend the static graphs by incorporating sequence embeddings as node features. These embeddings are obtained from a BERT model pretrained with masked language modeling on the user sequences.
- 4. BRIDGE. We stack a BERT sequence encoder with different GNN backbones, trained jointly in an end-to-end manner. In particular, the proposed TOKENXATTN layer achieves the best performance across all metrics. Unlike the Two-Stage Training, BRIDGE does not use frozen embeddings but learns the BERT and GNN components jointly on the target task.

**Evaluation Metrics.** For the friendship prediction task, we follow prior works [9, 13, 39] and use ranking-based metrics: mean reciprocal rank (MRR) and Hits-at-k, for k in {1,3,5,10}. At evaluation time, we rank the ground-truth friendship of each user against 100 negative samples, generated by uniformly sampling users not already connected to the target user in the training set.

**Results.** From Table 1, we observe a clear hierarchy of performance. Augmenting static graph methods with sequence embeddings as node features (*Two-Stage*) consistently improves their results. Our proposed variants of BRIDGE achieve the best scores across all metrics, substantially outperforming methods in the other three categories. By contrast, temporal graph models do not perform well, likely because encoding geohash as a node and user checkin sequences as time-stamped edges fails to capture fine-grained sequential patterns and can introduce spurious neighbor relations (e.g.,

treating any two users who check in at the same location as second-degree neighbors). In addition, by converting a Geohash into a node, we lose all the hierarchical information in the Geohash (i.e. first digit of a Geohash conveys larger area information than the second, third, etc.), whereas in our sequence model, we still retain this information by tokenizing a Geohash8 into 4 tokens. Finally, as the friendship does not have time information, we need to incorporate randomized timestamps into the temporal graph model, which may also cause performance degradation.

#### 5.2 Fraud Detection

Experimental Setup. We evaluate fraud detection on Amazon review data [18]. Each user's sequence comprises product IDs and ratings; for this task, we also encode the associated review text. Each review may receive "helpful" or "unhelpful" votes from other users. Following Dou et al. [8], we use this helpfulness information as a proxy for fraud review incidents. Specifically, we label a user with at least 5 reviews as fraudulent if more than 70% of their reviews are marked unhelpful. Relational edges are defined by co-review patterns: two users are connected if they have co-reviewed at least three products.

**Compared Methods.** Similar to the friendship prediction task, we evaluate static graph and temporal graph methods on the relational graph, using the resulting node embeddings for classification. We further compare against a text-based baseline that uses pretrained review embeddings.

- 1. *Graph Only models.* GCN [15], GAT [31], and TransformerConv [24] operate on a static user co-review graph with single-vector node representations, ignoring user sequences.
- 2. *Temporal Graph models*. We run TGN [23] and DyRep [28]. TNCN [44] is not directly applicable here, as it is designed for link prediction tasks. In this setup, the combination of product and rating is converted to a node. If a user rates the product with a particular rating, we connect the user node to the product rating node with a timestamped edge.
- 3. **SBert Seq.** We encode each review with Sentence-BERT<sup>1</sup> to obtain a sequence of embeddings per user. For classification, we average the embeddings to form a single user representation, which is passed to a classifier.
- 4. **BRIDGE.** Our approach combines Sentence-BERT embedding of review texts with a learnable rating embedding. The rating embedding and GNN backbone are trained jointly for fraud detection, allowing the model to integrate both textual and relational signals.

**Evaluation Metrics.** For fraud detection, we report two classification metrics: **Max F1** (best F1 across thresholds) and **PR-AUC** (area under the precision–recall curve).

**Results.** From Table 2, a consistent ordering emerges: graph-only < temporal graph < sequence-only (SBERT) < BRIDGE. Temporal structure helps over static graphs but still trails a simple SBERT average, highlighting the importance of textual information for this task. All BRIDGE variants outperform the baselines, showing clear gains from coupling event sequences with graph context.

#### 6 Conclusion

We presented BRIDGE, a unified end-to-end architecture that learns from sequences and graphs jointly, eliminating the common handoff where sequence encoders are frozen and passed to a separate GNN. In BRIDGE, the sequential and relational components are trained together under a single objective, allowing gradients to flow across both parts of the model and yielding task-aligned representations. This simple design choice, treating sequence and graph modeling as one training problem, consistently outperforms static GNNs, temporal graph models, and sequence-only baselines across friendship prediction and fraud detection, demonstrating its effectiveness. We further proposed Tokenxattn, a token-level cross-attention layer that enables interaction between events across neighboring sequences. Rather than viewing sequences and graphs as separate problems, we argue they are complementary components of the same dataset—Bridge unifies them under a single architecture, with Tokenxattn enabling event-level interactions across them. Looking forward, we believe this perspective opens the door to new models that can seamlessly integrate multimodal sequential and relational signals across a wide range of domains, from social networks to recommendation systems and beyond.

https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2

#### References

- [1] Maya Bechler-Speicher, Ben Finkelshtein, Fabrizio Frasca, Luis Müller, Jan Tönshoff, Antoine Siraudin, Viktor Zaverkin, Michael M. Bronstein, Mathias Niepert, Bryan Perozzi, Mikhail Galkin, and Christopher Morris. Position: Graph learning will lose relevance due to poor benchmarks, 2025. URL https://arxiv.org/abs/2502.14546.
- [2] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece
   Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi,
   Marco Tulio Ribeiro, and Yi Zhang. Sparks of artificial general intelligence: Early experiments
   with gpt-4, 2023. URL https://arxiv.org/abs/2303.12712.
- Zhikai Chen, Haitao Mao, Hang Li, Wei Jin, Haifang Wen, Xiaochi Wei, Shuaiqiang Wang,
   Dawei Yin, Wenqi Fan, Hui Liu, and Jiliang Tang. Exploring the potential of large language
   models (Ilms)in learning on graphs. ACM SIGKDD Explorations Newsletter, 25:42 61, 2023.
   URL https://api.semanticscholar.org/CorpusId:259375824.
- [4] Eunjoon Cho, Seth A. Myers, and Jure Leskovec. Friendship and mobility: user movement in location-based social networks. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '11, page 1082–1090, New York, NY, USA, 2011. Association for Computing Machinery. ISBN 9781450308137. doi: 10.1145/2020408.2020579. URL https://doi.org/10.1145/2020408.2020579.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V. Le, and Ruslan Salakhutdinov.
   Transformer-xl: Attentive language models beyond a fixed-length context, 2019. URL https://arxiv.org/abs/1901.02860.
- [6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, May 2019. URL http://arxiv.org/abs/1810.04805. arXiv:1810.04805 [cs].
- [7] Chaoyue Ding, Shiliang Sun, and Jing Zhao. MST-GAT: A Multimodal Spatial-Temporal
   Graph Attention Network for Time Series Anomaly Detection. *Information Fusion*, 89:527–536,
   January 2023. ISSN 15662535. doi: 10.1016/j.inffus.2022.08.011. URL http://arxiv.org/abs/2310.11169. arXiv:2310.11169 [cs].
- Yingtong Dou, Zhiwei Liu, Li Sun, Yutong Deng, Hao Peng, and Philip S Yu. Enhancing graph neural network-based fraud detectors against camouflaged fraudsters. In *Proceedings of the* 29th ACM international conference on information & knowledge management, pages 315–324,
   2020.
- Julia Gastinger, Shenyang Huang, Mikhail Galkin, Erfan Loghmani, Ali Parviz, Farimah Poursafaei, Jacob Danovitch, Emanuele Rossi, Ioannis Koutis, Heiner Stuckenschmidt, Reihaneh
   Rabbany, and Guillaume Rabusseau. Tgb 2.0: A benchmark for learning on temporal knowledge graphs and heterogeneous graphs, 2024. URL https://arxiv.org/abs/2406.09639.
- [10] Zirui Guo, Lianghao Xia, Yanhua Yu, Yuling Wang, Zixuan Yang, Wei Wei, Liang Pang, Tat-Seng Chua, and Chao Huang. Graphedit: Large language models for graph structure learning. ArXiv, abs/2402.15183, 2024. URL https://api.semanticscholar.org/CorpusId: 267897533.
- 266 [11] William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs, 2018. URL https://arxiv.org/abs/1706.02216.
- <sup>268</sup> [12] Chao Huang, Xubin Ren, Jiabin Tang, Dawei Yin, and Nitesh V. Chawla. Large language models for graphs: Progresses and directions. *Companion Proceedings of the ACM Web Conference* 2024, 2024. URL http://dl.acm.org/citation.cfm?id=3641251.
- 271 [13] Shenyang Huang, Farimah Poursafaei, Jacob Danovitch, Matthias Fey, Weihua Hu, Emanuele 272 Rossi, Jure Leskovec, Michael Bronstein, Guillaume Rabusseau, and Reihaneh Rabbany. 273 Temporal graph benchmark for machine learning on temporal graphs, 2023. URL https: 274 //arxiv.org/abs/2307.01026.

- 275 [14] Thomas N. Kipf and Max Welling. Variational graph auto-encoders, 2016. URL https: //arxiv.org/abs/1611.07308.
- Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks, 2017. URL https://arxiv.org/abs/1609.02907.
- 279 [16] Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection. http://snap.stanford.edu/data, June 2014.
- [17] Jiajin Liu, Dongzhe Fan, Jiacheng Shen, Chuanhao Ji, D. Zha, and Qiaoyu Tan. Graph mllm: Harnessing multimodal large language models for multimodal graph learning.
   ArXiv, abs/2506.10282, 2025. URL https://api.semanticscholar.org/CorpusId:
   279318592.
- [18] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton van den Hengel. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '15, page 43–52, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450336215. doi: 10.1145/2766462.2767755. URL https://doi.org/10.1145/2766462.2767755.
- [19] Giang Hoang Nguyen, John Boaz Lee, Ryan A Rossi, Nesreen K Ahmed, Eunyee Koh, and
   Sungchul Kim. Continuous-time dynamic network embeddings. In *Companion Proceedings of the The Web Conference 2018*, pages 969–976, 2018.
- 293 [20] Gustavo Niemeyer. Geohash, 2008.
- 294 [21] Bryan Perozzi, Bahare Fatemi, Dustin Zelle, Anton Tsitsulin, Mehran Kazemi, Rami Al-295 Rfou, and Jonathan J. Halcrow. Let your graph do the talking: Encoding structured data for 296 llms. ArXiv, abs/2402.05862, 2024. URL https://api.semanticscholar.org/CorpusId: 297 267547812.
- Zubin Ren, Jiabin Tang, Dawei Yin, Nitesh V. Chawla, and Chao Huang. A survey of large language models for graphs. Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2024. URL https://api.semanticscholar.org/CorpusId: 269761536.
- Emanuele Rossi, Ben Chamberlain, Fabrizio Frasca, Davide Eynard, Federico Monti, and Michael Bronstein. Temporal Graph Networks for Deep Learning on Dynamic Graphs, October 2020. URL http://arxiv.org/abs/2006.10637. arXiv:2006.10637 [cs].
- Yunsheng Shi, Zhengjie Huang, Shikun Feng, Hui Zhong, Wenjin Wang, and Yu Sun. Masked label prediction: Unified message passing model for semi-supervised classification, 2021. URL https://arxiv.org/abs/2009.03509.
- 308 [25] Zeren Shui, Ge Liu, Anoop Deoras, and George Karypis. Sequence-graph duality: user modeling with self-attention for sequential recom-Unifying 309 https://www.amazon.science/publications/ mendation, URL 310 sequence-graph-duality-unifying-user-modeling-with-self-attention-for-sequential-recommenda 311
- Yanchao Tan, Hang Lv, Pengxiang Zhan, Shiping Wang, and Carl Yang. Graph-oriented instruction tuning of large language models for generic graph mining. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. URL https://api.semanticscholar.org/CorpusId:268296710.
- Yijun Tian, Huan Song, Zichen Wang, Haozhu Wang, Ziqing Hu, Fang Wang, N. Chawla, and Panpan Xu. Graph neural prompting with large language models. In *AAAI Conference on Artificial Intelligence*, 2023. URL https://arxiv.org/pdf/2309.15427.pdf.
- Rakshit Trivedi, Mehrdad Farajtabar, Prasenjeet Biswal, and Hongyuan Zha. DyRep: Learning Representations over Dynamic Graphs. September 2018. URL https://openreview.net/forum?id=HyePrhR5KX.

- [29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper\_files/paper/2017/hash/ 3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.
- [30] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua
   Bengio. Graph Attention Networks, February 2018. URL http://arxiv.org/abs/1710.
   10903. arXiv:1710.10903 [stat].
- 330 [31] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua 331 Bengio. Graph attention networks, 2018. URL https://arxiv.org/abs/1710.10903.
- Hongjun Wang, Jiyuan Chen, Tong Pan, Zheng Dong, Lingyu Zhang, Renhe Jiang, and Xuan Song. STGformer: Efficient Spatiotemporal Graph Transformer for Traffic Forecasting, October 2024. URL http://arxiv.org/abs/2410.00385. arXiv:2410.00385 [cs] version: 1.
- Xiang Wang, Xiangnan He, Yixin Cao, Meng Liu, and Tat-Seng Chua. Kgat: Knowledge graph attention network for recommendation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '19, page 950–958. ACM, July 2019. doi: 10.1145/3292500.3330989. URL http://dx.doi.org/10.1145/3292500.3330989.
- [34] Yuxiang Wang, Xinnan Dai, Wenqi Fan, and Yao Ma. Exploring graph tasks with pure
   llms: A comprehensive benchmark and investigation. ArXiv, abs/2502.18771, 2025. URL
   https://api.semanticscholar.org/CorpusId:276617594.
- Shiwen Wu, Fei Sun, Wentao Zhang, Xu Xie, and Bin Cui. Graph neural networks in recommender systems: A survey, 2022. URL https://arxiv.org/abs/2011.02260.
- [36] Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, Xiaojun Chang, and Chengqi Zhang.
   Connecting the Dots: Multivariate Time Series Forecasting with Graph Neural Networks. In
   Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery
   & Data Mining, KDD '20, pages 753–763, New York, NY, USA, August 2020. Association
   for Computing Machinery. ISBN 978-1-4503-7998-4. doi: 10.1145/3394486.3403118. URL
   https://dl.acm.org/doi/10.1145/3394486.3403118.
- [37] Da Xu, Chuanwei Ruan, Evren Korpeoglu, Sushant Kumar, and Kannan Achan. Inductive
   Representation Learning on Temporal Graphs, February 2020. URL http://arxiv.org/
   abs/2002.07962. arXiv:2002.07962 [cs].
- 353 [38] Ruosong Ye, Caiqi Zhang, Runhui Wang, Shuyuan Xu, and Yongfeng Zhang. Language is all a graph needs. In *Findings*, 2023. URL https://api.semanticscholar.org/CorpusId: 260887732.
- [39] Lu Yi, Jie Peng, Yanping Zheng, Fengran Mo, Zhewei Wei, Yuhang Ye, Yue Zixuan, and
   Zengfeng Huang. Tgb-seq benchmark: Challenging temporal gnns with complex sequential
   dynamics, 2025. URL https://arxiv.org/abs/2502.02975.
- [40] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L. Hamilton, and Jure Leskovec. Graph convolutional neural networks for web-scale recommender systems. In
   Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '18, page 974–983. ACM, July 2018. doi: 10.1145/3219819.3219890. URL http://dx.doi.org/10.1145/3219819.3219890.
- [41] Bing Yu, Haoteng Yin, and Zhanxing Zhu. Spatio-Temporal Graph Convolutional Networks:
   A Deep Learning Framework for Traffic Forecasting. In Proceedings of the Twenty-Seventh
   International Joint Conference on Artificial Intelligence, pages 3634–3640, July 2018. doi:
   10.24963/ijcai.2018/505. URL http://arxiv.org/abs/1709.04875. arXiv:1709.04875
   [cs].
- Mengqi Zhang, Shu Wu, Xueli Yu, Qiang Liu, and Liang Wang. Dynamic Graph Neural Networks for Sequential Recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 35(05):4741–4753, May 2023. ISSN 1041-4347. doi: 10.1109/TKDE.2022.

- 3151618. URL https://www.computer.org/csdl/journal/tk/2023/05/09714053/ 1BOXPB9Fgk0. Publisher: IEEE Computer Society.
- Muhan Zhang and Yixin Chen. Link prediction based on graph neural networks, 2018. URL https://arxiv.org/abs/1802.09691.
- 376 [44] Xiaohui Zhang, Yanbo Wang, Xiyuan Wang, and Muhan Zhang. Efficient neural common
   377 neighbor for temporal graph link prediction, 2024. URL https://arxiv.org/abs/2406.
   378 07926.
  - [45] Huachi Zhou, Jiahe Du, Chuang Zhou, Chang Yang, Yilin Xiao, Yuxuan Xie, and Xiao Huang. Each graph is a new language: Graph learning with llms. In *Annual Meeting of the Association for Computational Linguistics*, 2025. URL https://api.semanticscholar.org/CorpusId:275757427.

Table 1: Link ranking performance across four datasets (higher is better). **H@k** denotes the fraction of test queries for which at least one ground-truth link appears among the top-k predictions.

	Brightkite					Amazon-Clothing						
Method	MRR	H@1	H@3	H@5	H@10	H@50	MRR	H@1	H@3	H@5	H@10	H@50
Graph Only (Static)												
GCN	0.6962	0.5832	0.7759	0.8388	0.9006	0.9672	0.1909	0.1103	0.1953	0.2535	0.3496	0.6423
GAT	0.6683	0.5538	0.7427	0.8098	0.8833	0.9687	0.1451	0.0884	0.1180	0.1577	0.2574	0.6444
TransformerConv	0.6876	0.5706	0.7702	0.8364	0.9009	0.9699	0.1951	0.1159	0.1943	0.2503	0.3512	0.6456
Two-Stage (Static + Sequence)												
GCN + S	0.6917	0.5754	0.7739	0.8400	0.9027	0.9697	0.1632	0.0827	0.1571	0.2094	0.3233	0.6452
GAT + S	0.6850	0.5700	0.7643	0.8292	0.8963	0.9724	0.1829	0.1193	0.1757	0.2189	0.2970	0.6407
TransformerConv + S	0.7298	0.6239	0.8073	0.8640	0.9170	0.9719	0.1896	0.1040	0.1890	0.2518	0.3695	0.7231
Temporal Graphs												
TGN	0.4056	0.2684	0.4579	0.5548	0.6929	0.9721	0.1616	0.0576	0.1536	0.2330	0.3916	0.8200
DyRep	0.2994	0.1720	0.3267	0.4216	0.5717	0.9519	0.1309	0.0345	0.1155	0.1823	0.3251	0.9171
TNCN	0.5347	0.4251	0.5799	0.6559	0.7564	0.9655	0.3108	0.2230	0.3267	0.3846	0.4683	0.8407
BRIDGE												
Bridge-GCN	0.9353	0.9163	0.9470	0.9567	0.9689	0.9920	0.4554	0.2983	0.5175	0.6993	0.8357	0.8958
Bridge-GAT	0.7935	0.7107	0.8568	0.8970	0.9336	0.9799	0.6992	0.6351	0.7441	0.7778	0.8051	0.8684
Bridge-TransformerConv	0.7902	0.7043	0.8539	0.8964	0.9365	0.9912	0.4370	0.3133	0.4419	0.5537	0.8207	0.9020
BRIDGE-TokenXAttn	0.9458	0.9304	0.9545	0.9633	0.9735	0.9947	0.4037	0.3466	0.4301	0.4569	0.4807	0.5627
		Amazon-Electronics					Amazon-Musical					
Method	MRR	H@1	H@3	H@5	H@10	H@50	MRR	H@1	H@3	H@5	H@10	H@50
Graph Only (Static)												
GCN	0.5312	0.3968	0.6135	0.7007	0.7876	0.9033	0.1463	0.0727	0.1399	0.1988	0.2965	0.6657
GAT	0.4411	0.2961	0.5120	0.6223	0.7453	0.9009	0.1256	0.0622	0.1005	0.1438	0.2410	0.7018
TransformerConv	0.5665	0.4419	0.6446	0.7194	0.7967	0.9173	0.2489	0.1516	0.2571	0.3320	0.4625	0.7463
Two-Stage (Static + Sequen	ıce)											
GCN + S	0.5156	0.3778	0.5997	0.6907	0.7782	0.8996	0.1861	0.1133	0.1738	0.2177	0.3237	0.7429
GAT + S	0.5061	0.3692	0.5878	0.6738	0.7636	0.9101	0.1294	0.0689	0.1088	0.1516	0.2382	0.6680
TransformerConv + S	0.5865	0.4652	0.6630	0.7358	0.8084	0.9276	0.2550	0.1594	0.2604	0.3159	0.4742	0.8279
Temporal Graphs												
TGN	0.3549	0.2259	0.3971	0.4889	0.6196	0.9264	0.1358	0.0562	0.1314	0.1935	0.2712	0.6989
DyRep	0.1731	0.0687	0.1733	0.2463	0.3878	0.8990	0.1237	0.0585	0.0980	0.1446	0.2521	0.6726
TNCN	0.4424	0.3446	0.4863	0.5475	0.6185	0.8547	0.1691	0.0848	0.1720	0.2234	0.3262	0.7085
Bridge												
Bridge-GCN	0.7466	0.6864	0.7851	0.8174	0.8517	0.9112	0.4453	0.2959	0.5514	0.6419	0.7085	0.8479
Bridge-GAT	0.6910	0.6062	0.7520	0.7926	0.8260	0.9273	0.5045	0.3853	0.5775	0.6757	0.7229	0.8306
BRIDGE-TransformerConv	0.6984	0.6081	0.7608	0.8074	0.8492	0.9619	0.4819	0.3787	0.5180	0.5786	0.7168	0.8823
BRIDGE-TokenXAttn	0.7318	0.6724	0.7685	0.8069	0.8353	0.8996	0.5870	0.5075	0.6341	0.6702	0.7157	0.8368

# A Additional Related Works

# A.1 Sequence Modeling

383

384

385

386

387

388

389

390

379

380

381 382

Transformer-based sequence models have delivered state-of-the-art performance across a wide range of language and vision tasks by learning contextual representations over token sequences [2, 5, 6, 29]. Given an input sequence, these models map each token to a continuous embedding and use self-attention to capture local and long-range dependencies within the sequence. While highly effective for modeling the internal structure of a single sequence, this paradigm typically treats each sequence in isolation and does not leverage information from *related* sequences.

Table 2: Fraud detection performance across three Amazon categories.

	Amazo	n-Movies	Amazon-	Electronics	Amazon-Clothing			
Method	Max F1	PR AUC	Max F1	PR AUC	Max F1	PR AUC		
Graph Only (Static)								
GCN	0.3374	0.2037	0.1197	0.0640	0.0485	0.0236		
GAT	0.3390	0.2146	0.1192	0.0638	0.0497	0.0240		
TransformerConv	0.3356	0.2129	0.1190	0.0639	0.0462	0.0216		
Temporal Graphs								
TGN	0.4438	0.3548	0.1449	0.0897	0.0799	0.0369		
DyRep	0.4412	0.3117	0.1247	0.0676	0.0545	0.0270		
Sequence Only (Review Text Embedding)								
SBert Seq	0.7809	0.8559	0.4218	0.3867	0.3423	0.3195		
BRIDGE								
Bridge-GCN	0.7956	0.8565	0.4524	0.4343	0.3598	0.3481		
Bridge-GAT	0.8039	0.8704	0.4617	0.4518	0.3840	0.3629		
BRIDGE-TransformerConv	0.7984	0.8694	0.4789	0.4656	0.3702	0.3539		
BRIDGE-TokenXAT	0.7969	0.8771	0.5016	0.5228	0.3606	0.3442		

#### A.2 Graph Modeling

Graph models capture relational structures among entities (nodes). In particular, graph neural networks (GNNs) propagate and aggregate information over local neighborhoods to learn node, edge, or graph level embeddings [15, 24, 30]. These methods have demonstrated strong performance across tasks such as node classification[11, 15], link prediction [14, 43], and recommendation [33, 35, 40]. Standard GNNs model each node in the graph as a single feature vector and do not account for sequential information within nodes. Bridging graph-based relational modeling with sequence-aware representations is therefore an important step toward richer models that capture both intra-sequence patterns and inter-sequence relations.

#### A.3 Graph Large Language Models

An emerging line of research seeks to leverage the generalization capabilities of foundation models for graph learning. *Graph large language models (Graph LLMs)* achieve this by integrating pretrained LLMs with graph-structured data through prompting [3, 21, 27], instruction tuning [10, 26, 34, 38], or hybrid architectures [12, 17, 22]. The promise of Graph LLMs lies in bridging reasoning over relational data with the flexibility of natural language interfaces. In line with this view, Zhou et al. [45] even conceptualizes that each graph can be seen as a new language for the LLMs to learn. Our approach, however, differs: rather than adapting pretrained LLMs to graphs, we introduce a unified architecture that integrates sequence modeling with GNNs to jointly capture temporal dynamics and relational structure.

#### 410 B Limitations and Discussion

- While BRIDGE provides a unified way to learn from sequences and graphs, several limitations remain.
- Use of timestamps. BRIDGE preserves event order but does not inject absolute timestamps or inter-event gaps. Likewise, TOKENXATTN attends over tokens without explicit time encoding. Tasks where absolute time or time interval between events are informative may benefit from adding relative or absolute time features, which we leave to future work.
- Complexity of TOKENXATTN. For preserving the granularity of event-level information, the complexity of TOKENXATTN is quadratic with respect to the sequence length. Some potential efficiency improvements can be applying a more efficient attention mechanism, such as linear attention, or reducing the sequence length through patching.