

On Generalizing Static Node Embedding to Dynamic Settings: A Practitioner’s Guide

Di Jin
University of Michigan, Ann Arbor
dijin@umich.edu

Ryan A. Rossi
Adobe Research
rossi@adobe.com

Sungchul Kim
Adobe Research
sukim@adobe.com

Danai Koutra
University of Michigan, Ann Arbor
dkoutra@umich.edu

ABSTRACT

Temporal graph embedding has been widely studied due to its superiority in tasks such as prediction and recommendation. Despite the advancement in algorithms and novel frameworks such as deep learning, there has been relatively little work on systematically studying the properties of temporal network models and their cornerstones, the graph time-series representations that are used in these approaches. This paper aims to fill this gap by introducing a general framework that extends an arbitrary existing static embedding approach to handle dynamic tasks, and conducting a systematic study of *seven* base static embedding methods and *six* temporal network models. Our framework generalizes static node embeddings derived from the time-series representation of stream data to the dynamic setting by modeling the temporal dependencies with classic models such as the reachability graph. While previous works on dynamic modeling and embedding have focused on representing a stream of timestamped edges using a time-series of graphs based on a specific time-scale (e.g., 1 month), we introduce the notion of an ϵ -graph time-series that uses a fixed number of edges for each graph, and show its superiority in practical settings over the standard solution. From the 42 methods that our framework subsumes, we find that leveraging the new ϵ -graph time-series representation and capturing temporal dependencies with the proposed reachability or summary graph tend to perform well. Furthermore, the new dynamic embedding methods based on our framework perform comparably and on average better than the state-of-the-art embedding methods designed specifically for temporal graphs in link prediction tasks. We expect our results and findings to be useful for practitioners in the design of new embedding methods and predictive models for temporal networks.

1 INTRODUCTION

Real-world networks that record the interaction between entities have grown rapidly, for example, the Internet [5], various online

social networks (e.g., Facebook, Snapchat), citation networks in academia [14]. Specifically, when nodes and edges continuously change over time with addition, deletion (e.g., a phone call, an email, or physical proximity between two entities), we have a particular type of evolving network structure. Learning an appropriate network representation (embedding) that accurately captures the temporal dynamics and temporal structural properties of these entities is important for many downstream time-series forecasting/prediction tasks such as recommendation and entity resolution. Most recent research efforts devoted in the field follow the common pipeline: given a time-series of graphs, $\mathcal{G} = \{G_1, \dots, G_k, \dots, G_T\}$, modeling the individual graph structures (within-snapshot property) along with the temporal dependency (across-snapshot relation), and deriving node embeddings that incorporate both perspectives. While these works show advantage from various perspectives, the promising performance comes at the cost of time and model complexity, such as introducing extra transition variables to reflect the temporal dependency between snapshots [8], or latent weights on edges between snapshots [18, 24].

In this work, we propose a general framework that simplifies the above process and can generalize *any* static embedding method to a more powerful and predictive dynamic embedding method without introducing transitional variables. The framework consists of three components: (C1) a graph time-series representation, (C2) a temporal network model that appropriately models and weights the temporal dependencies in the graph time-series, and (C3) a base embedding method to learn a time-series of embeddings along with a fusion mechanism to derive the final temporal node embeddings. The framework is highly expressive as any unique combination of C1-C3 gives rise to a new dynamic embedding method.

While previous works on dynamic modeling and embedding have focused on representing the stream of timestamped edges [17] using a time-series of graphs based on a specific time-scale τ (e.g., $\tau = 1$ hour, or 1 month) [7, 8, 14, 24, 25, 29], we instead propose the notion of an ϵ -graph time-series that uses a fixed number of edges for each graph in the time-series. Theoretically, by fixing the number of edges to be ϵ in each graph, we ensure that every graph in the sequence has an equal probability of giving rise to the same exact distribution of higher-order graphlets and other structural patterns¹, and therefore, the new ϵ -graph time-series forces the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
Conference'17, July 2017, Washington, DC, USA

© 2021 Association for Computing Machinery.
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

¹This is in contrast to graphs with different amounts of edges. E.g., given two arbitrary graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ where $|E_1| \ll |E_2|$, then the counts of all $k \in \{3, 4, \dots\}$ -node network motifs (graphlets) in G_2 are almost surely larger than G_1 .

models to avoid capturing simple trivial differences due to edge counts, and instead, allow the models to capture actual *structural changes* to the graphs over time.

We also introduce a number of important temporal models that can be leveraged over any graph time-series representation of the edge stream. The first temporal model is based on the notion of a temporal reachability graph (TRG). TRG is derived by transforming a dynamic graph into a static graph where an edge from u to w indicates a temporal walk. The second temporal model is called a weighted temporal summary graph (TSG). Notably, a weighted temporal summary graph captures the temporal recurrence and recency of links by appropriately weighting links with respect to a function f that assigns larger weights to links that are more recent and recurrent whereas links that occur in the more distant past are assigned lower weights. All temporal models can leverage either the new ϵ -graph or τ -graph time-series representation.

This paper aims to provide a systematic exploration of the most useful graph time-series representations and temporal network models (used to incorporate the temporal dependencies into base embedding methods) in downstream temporal prediction tasks. To the best of our knowledge, this is the first work of this kind. Our primary findings are: (1) node embeddings derived from the ϵ -graphs outperforms the τ -graph time-series in the predictive task with higher stableness, and (2) by composing the static node embedding approaches with classic temporal models such as TRG or TSG, our proposed framework performs comparably or even better than recent dynamic embedding approaches with less complexity. Based on these findings, we hope that this work will benefit practitioners who leverage temporal graph embedding in various applications due to its simplicity in deployment and effectiveness in performance. We also hope that this work could benefit future research on developing and evaluating better dynamic node embedding methods. Our main contributions are as follows:

- **General Framework.** We describe a general framework for leveraging graph stream data and classic temporal network models for prediction-based applications that can generalize any static graph embedding method.
- **Powerful Graph Time-series Representation.** We introduce the notion of a ϵ -graph time-series and show its superiority over the conventional way of discretizing the edge stream based on the application time-scale (e.g., hour, day).
- **Systematic Study.** Our framework allows us to systematically study 42 dynamic node embeddings by combining time-series representations, temporal network models, and static methods. Strikingly, our empirical analysis on 8 real-world networks shows that our framework achieves comparable or better predictive performance than existing state-of-the-art, but more complex, *dynamic* node embedding methods.

2 RELATED WORK

Snapshot-based approaches. Most temporal embedding approaches break down the graph into graph time series based on the application time-scale (1 month, etc.) up to a certain point k , and then derive features from them to make inference on graphs at $k + 1$. One direction is to look into the most recent snapshot, for instance,

Table 1: Qualitative comparison of existing embedding methods on temporal graphs. The graph time-series representation used by the method (application time-scale, or fixed number of edges), the type of temporal model used, and type of embedding fusion used (if any).

	REPRESENTATION		TEMPORAL MODEL		
	Time (τ)	Edge (ϵ)	Snapshot	Weighting	Emb. Fusion
DANE [16]	✓	✗	✓	✗	✗
DynGem [8]	✓	✗	✓	✗	✓
TIMERS [28]	✓	✗	✓	✗	✗
Dynagraph2vec [7]	✓	✗	✓	✗	✓
tNodeEmbed [25]	✓	✗	✓	✓	✓
EvolveGCN [18]	✓	✗	✓	✓	✓
DySAT [24], DyHATR [27]	✓	✗	✓	✗	✓
our framework [24]	✓	✓	✓	✓	✓

DANE [16] proposes to embed both nodes and the associated attributes in the graph by minimizing the loss of reconstruction of the snapshot at a given times point k : $\frac{1}{2} \sum_{i,j} A_{ij}^{(k)} \|y_i - y_j\|^2$, and update the embeddings for snapshot at $k + 1$ based on the change of graph structure and node attributes. DynGEM [8] adopts the deep auto-encoder to generate the nonlinear embeddings from the snapshot at k while addressing stability. TIMERS [28] models the relative changes in adjacency matrices between snapshots and leverages incremental SVD to derive embeddings. A more popular direction is to track back a certain number of snapshots from the time point k by deriving node embeddings from each individual tracked snapshot and then merging them through specific operation. Dynagraph2vec [7] leverage totally l snapshots to predict the snapshot at $k + 1$. It leverages various deep architectures (*i.e.*, auto-encoder, RNN) to derive latent features by minimizing loss of reconstruction error: $\|f(A_{k-l+1}, \dots, A_k) - A_{k+1}\|_F^2$. tNodeEmbed [25] is an end-to-end framework based on node embeddings derived from individual snapshots using static methods. The embeddings are merged by minimizing the loss of specific tasks (*i.e.*, link prediction and node classification) through LSTM. DySAT [24] leverages the notion of self-attention to compute node representations by jointly employing graph structural property and temporal dynamics. Similarly, DyHATR [27] proposes a hierarchical attention model to capture both the heterogeneity and temporal attention using GRU/LSTM to model the temporal evolution. EvolveGCN [18] uses GCN to generate node embeddings for the past snapshots, and learns the hidden parameters for the next using GRU/LSTM. Unlike the above methods that jointly explore the graph structural changes with the evolution of the #edges, our proposed ϵ -graph time series does not require the specification of time-scales.

Sequential-interaction-based approaches. There is another line of works that studies the sequential interaction between nodes in the graph. CTDNE [17] is the first approach to learn embeddings directly from the stream of timestamped edges at the finest temporal granularity. In that work, they proposed the notion of temporal walks and used it for embeddings [17]. More recently, node2bits [11] expanded on this idea by incorporating features in the temporal walks and hashing them. Alternatively, some other work has modeled the node-specific temporal dynamics as the point process where the probability of interaction is represented through different intensity functions. HTNE [30] proposes to model the node evolution through the Hawkes process. JODIE[13] models the sequential interaction in bipartite graphs to predict the change of

Table 2: Network statistics and properties

Data	$ V $	$ E $	Type	Timespan
enron	151	50,572	Unipartite	38 months
bitcoin	3,783	24,186	Unipartite	63 months
wiki-elec	7118	107,071	Unipartite	47 months
stackoverflow	24,818	506,550	Tripartite	79 months
fb-forum	899	33,720	Unipartite	24 weeks
reallity-call	6,809	52,050	Unipartite	16 weeks
wiki-edit	8,227	157,474	Bipartite	32 days
contacts-dublin	10,972	415,912	Unipartite	69 days

embedding trajectory over time instead of interaction probability. CTDNE, HTNE and JODIE are designed to handle continuously sequential data, which is not the scope of this paper.

3 DATA

In this study we adopt a variety of real-world temporal networks from SNAP [15] and NR [20], the brief description is as follows.

- enron² records email exchanging between employees of Enron from May, 1999 to June, 2002.
- bitcoin³ is a who-trusts-whom network of people who trade using bitcoins from Nov, 2010 to Feb., 2017. We study the user connectivity by dropping the edge signs.
- wiki-elec² contains the voting history based on the Wikipedia page edit history from Mar., 2004 to Jan., 2008.
- stackoverflow³ is a temporal network consisting of three types of interactions on the stack exchange web site Math Overflow: a user answers questions, a user comments on questions, and a user comments on answers.
- wiki-edit⁴ is a public bipartite dataset containing one month of edits made by users in the Wikipedia page.
- fb-forum² is the Facebook-like Forum network that records users’ activity in the forum.
- contacts-dublin² is a human contact network where nodes represent humans and edges between them represent proximity (i.e., contacts in the physical world).
- reality-call² is a subgraph of the reality mining study. Nodes are participants and edges are phone calls.

We summarize the graph statistics and temporal timespans in Table 2, and analyze the sequential graph statistics of three graphs over time. As the timespans vary from 32 days to 79 months, we adopt the time-scale following Table 2 to get the sequential graph time-series. We visualize 2 graph statistics, the number of edges $|E|$ and the average degree on 3 datasets with different time-scales in Figure 1, which are contacts-dublin (day), wiki-elec (month), and fb-forum (week). In the figure, we also visualize the same graph statistics using a different time-series representation by fixing the number of edges in each snapshot to $\frac{|E|}{T}$, where T denotes the timespan following the corresponding time-scale. For example, for wiki-elec, this number is $\frac{107,701}{47}$ in each snapshot. From Figure 1, we compare the temporal patterns of the two time-series and it can be seen following the fixed edge count in each snapshot gives more stable temporal pattern using both graph statistics. We discuss this new graph time-series in detail in Section 5.1. Besides, in this work,

²<http://networkrepository.com>

³<https://snap.stanford.edu/data/>

⁴<https://github.com/srijankr/jodie>

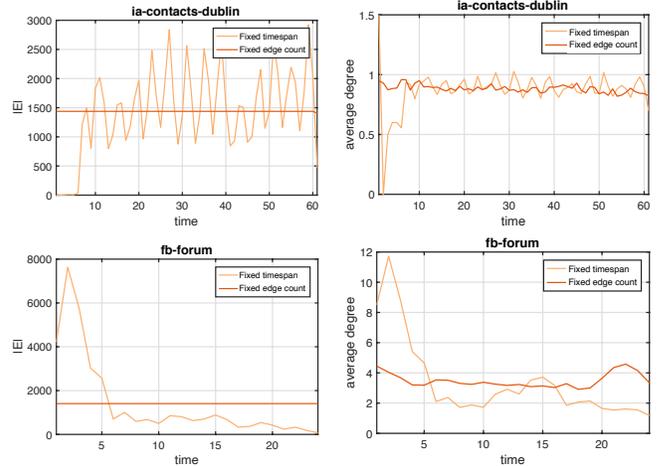


Figure 1: Graph properties (#edge and average degree) over two time-series representation (fixed timespans vs. fixed edge count). Fixing the edge number gives more stable temporal patterns while fixing the timespans shows higher fluctuation.

we focus on exploring the impacts of graph structures and temporal dependency between snapshots to the predictive tasks, thus we do not leverage node features such as geographic location or content.

4 PRELIMINARIES

We summarize symbols and notations used in this work in Table 3. Some important notions are given as follows.

DEFINITION 1 (TEMPORAL GRAPH). Let V be a set of vertices, and $E \subseteq V \times V \times \mathbb{R}^+$ be the set of temporal edges between vertices in V . Each edge (u, v, t) has a unique time $t \in \mathbb{R}^+$.

When edges represent contacts—a phone call, an email, or physical proximity—between two entities at a specific point in time, we have a particular type of evolving network structure [2]. A temporal walk in such a network represents a sequence of contacts that obeys time. That is, if each edge represents a contact between two entities, then a path represents a feasible route for a piece of information.

DEFINITION 2 (TEMPORAL WALKS). A temporal walk from u to w in $G = (V, E)$ is a sequence of edges e_1, \dots, e_k such that $e_1 = (u_1, u_2, t_1), \dots, e_k = (u_k, u_{k+1}, t_k)$ where $t_j < t_{j+1}$ for all $j = 1$ to k . We say that u is temporally connected to w if there exists such a temporal walk.

This definition echoes the standard definition of a path, but adds the additional constraint that paths must respect time, i.e., follow the directionality of time. Temporal walks are inherently asymmetric because of the directionality of time. The notion of temporal walks has been recently used in embedding methods [17].

5 FRAMEWORK

The framework in this paper provides a fundamental basis for studying different temporal network representations and the utility of these for generalizing existing static embedding methods to temporal network data. As shown in Figure 2, our proposed framework has the following main components. First, given the continuous

Table 3: Summary of notation

Symbol	Definition
$\mathcal{G} = \{G_k\}$	a graph time-series with snapshots indexed by k .
$G_k = (V_k, E_k)$	a directed and weighted temporal network from \mathcal{G} with $ V_k $ nodes and $ E_k $ temporal edges
A_k	adjacency matrix for graph G_k in \mathcal{G} .
$G_R = (V, E_R)$	the weighted temporal reachability graph
N_i^R	the set of nodes that are temporally reachable from node i
τ	window size representing the timespan of edges
ϵ	window size representing the number of edges
α	the decay factor in the temporal summary graph model
θ	the decay factor in the temporal embedding smoothing
f	arbitrary base embedding method
d	dimensionality of the embedding
Z	$ V \times d$ embedding matrix

stream of timestamped edges, we derive the time-series of graphs (Section 5.1). Then, we use one of the temporal network models to incorporate the temporal dependencies of the graph-based time-series (Section 5.2). Third, our framework generalizes existing embedding methods and effectively enables the new dynamic variants of these methods to learn more accurate and appropriate time-dependent embeddings. (Section 5.3). We also briefly describe a few general methods that can be used over the resulting time-dependent embeddings for improving temporal prediction tasks (Section 5.3.2).

5.1 Graph Time-Series Representations

We formally introduce two approaches for deriving a time-series of graphs from the stream of timestamped edges. For clarity, we use k to index the snapshots in the time-series in this section to avoid mixing with the timestamp t associated with an edge e .

5.1.1 τ -graph time-series. The τ -graph time-series representation is used by the vast majority of previous work [7, 10].

DEFINITION 3 (τ -GRAPH TIME-SERIES). Given a temporal network $G=(V, E)$ representing a continuous edge stream with time-stamped edges E , we define a graph time-series $\mathcal{G}^\tau = \{G_1, \dots, G_k, \dots\}$ such that G_1 consists of all edges within the first time scale (period) s , G_2 consists of all edges within the next time period s , and so on. Thus, each graph contains edges within a specific period of time. More formally, let t_0 denote the timestamp of the first edge in the temporal network (stream of timestamped edges) and τ is the application time-scale (e.g., 1 month), then

$$E_k = \{(i, j, t) \in E \mid t_0 + k\tau > t \geq t_0 + (k-1)\tau\} \quad (1)$$

Temporal models that use a time-based graph time-series are denoted with the suffix “ $-\tau$ ”.

5.1.2 ϵ -graph time-series. While most work uses the previous approach for deriving the graph time-series, we introduce a new alternative based on the idea of using a fixed number of edges. In particular, we propose a new approach that derives a time-series of graphs $\mathcal{G}^\epsilon = \{G_1, \dots, G_k, \dots\}$ such that each G_k consists of ϵ edges (Definition 4) and therefore $|E_k| = \epsilon, \forall k$. More formally,

DEFINITION 4 (ϵ -GRAPH TIME-SERIES). Given a temporal network $G = (V, E)$ representing a continuous edge stream E with timestamped edges and let ϵ denote a fixed number of temporal edges in the stream (ordered by time), we define a graph time-series $\mathcal{G}^\epsilon = \{G_1, \dots, G_k, \dots\}$ such that $|E_k| = \epsilon$, for all $k = 1, 2, \dots$. Hence,

$G_1 = (E_1, V)$ consists of the first ϵ edges $E_1 = \{e_1, e_2, \dots, e_\epsilon\}$ whereas G_2 consists of the next ϵ edges $E_2 = \{e_{\epsilon+1}, \dots, e_{2\epsilon}\}$, and so on. More formally, E_k is defined as follows:

$$E_k = \bigcup_{i=(k-1)\epsilon+1}^{k\epsilon} e_i = \{e_{(k-1)\epsilon+1}, \dots, e_{k\epsilon}\} \quad (2)$$

Temporal models that use the notion of an ϵ -graph time-series are denoted with the suffix “ $-\epsilon$ ”.

Note in both cases $E_1 \cup \dots \cup E_k \cup \dots = E$. Since the proposed ϵ -graph time-series controls for the number of edges over time, embedding methods can more appropriately model and capture the actual change in the structural properties and subgraph patterns over time, as opposed to just the frequency of edges that is captured by the τ -graph time-series representation used in previous work. Another advantage of the ϵ -graph time-series representation is that it preserves the sequential order of timestamped edges *without* suffering from the structural instability of the graph due to the sometimes drastic difference in edge counts from one time to the next. As observed in Fig. 1, while the ϵ -graph time-series representation has a fixed number of edges over time, conventionally-used τ representation can significantly deviate with large spikes even between consecutive graphs in the series. Finally, it is obvious that if a graph time-series representation is unable to capture the simplest 1st-order subgraph structures (edges), then by definition it cannot capture higher-order subgraph structures that are built on such lower-order ones. Hence, the proposed ϵ -graph time-series representation effectively models the *structural changes* between graphs whereas the τ -graph time-series captures changes in *edge frequencies* for a fixed application-specific time-scale such as 1 day or 1 hour.

5.2 Temporal Network Models

Now we introduce temporal network models that incorporate the temporal dependencies into the graph time-series representations to learn more effective time-dependent embeddings.

5.2.1 Snapshot Graph (SG) Model. This model simply leverages the $\{\tau, \epsilon\}$ -graph time-series representation directly without encoding any additional temporal information into the representation. Hence, the temporal information (edge timestamps) associated with the edges in any graph $G_k \in \mathcal{G}$ is effectively ignored/discarded. For example, e_1 and e_2 are considered to occur simultaneously if they fall into the same snapshot, even though e_2 comes later than e_1 in the actual time-series. Therefore, this model incorporates the temporal dependencies at the level of the graph, *i.e.*, we only know that edges in G_{k-1} occurred before G_k .

5.2.2 Temporal Summary Graph (TSG) Model. The temporal summary graph model incorporates the temporal dependencies by deriving a weighted summary graph from the graph-based time series \mathcal{G} [23] where the more recent edges are assigned larger weights than those in the distant past. More formally, let $A_1, A_2, \dots, A_k, \dots, A_T$ be a time-series of adjacency matrices of the graph time-series constructed using either Definition 3 or Definition 4. Furthermore, let $A_k(i, j)$ denote the (i, j) entry of A_k . We define the general *weighted*

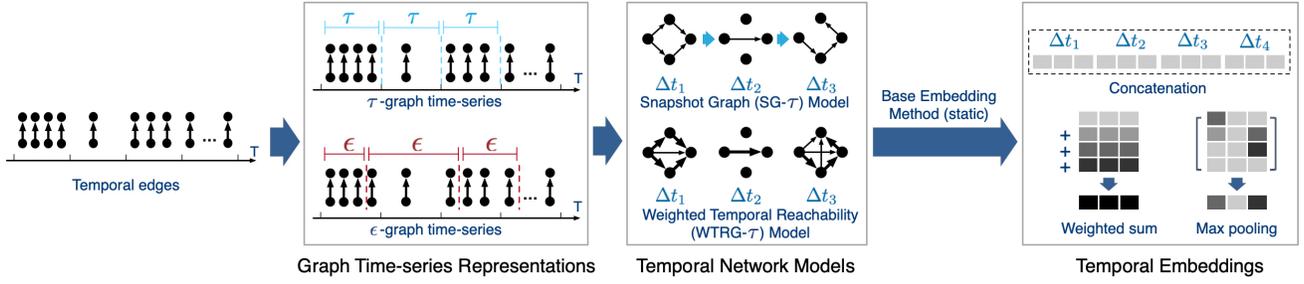


Figure 2: Framework Overview. In the first component of the framework (Sec. 5.1), we derive a time-series of graphs from the stream of timestamped edges using either an application-specific time-scale τ (e.g., 1 day) or a fixed number of edges ϵ for each graph in the time-series. Next, given the $\{\tau, \epsilon\}$ -graph time-series representation, we incorporate the temporal dependencies and weights with a temporal network model from Sec. 5.2. Finally, we use an arbitrary base embedding method to learn a time-series of embeddings and then leverage a temporal fusion mechanism to obtain the final temporal embeddings (Sec. 5.3).

temporal summary graph (TSG) model as

$$S = \sum_{k=1}^T f(A_k, \alpha) \quad (3)$$

where f is a decay function for temporally weighting the edges (nonzeros), α is the decay factor ranging in $(0, 1)$, T is the total number of graphs in the time-series, and S is the weighted temporal summary graph. In this work, we define f as an exponential decay function [23], then we obtain

$$S = \sum_{k=1}^T (1 - \alpha)^{T-k} A_k \quad (4)$$

and the weight for an edge (i, j) is simply $S(i, j) = \sum_{k=1}^T (1 - \alpha)^{T-k} A_k(i, j)$. Alternatively, instead of using all available graphs in the initial time-series, we can use only the L most recent graphs. For example, suppose $\mathcal{G}^\epsilon = \{G_k\}_{k=1}^T = \{G_1, \dots, G_T\}$ is an ϵ -graph time-series with T graphs. Instead of using all T graphs, we can leverage only the most recent L graphs, hence,

$$\mathcal{G}^\epsilon = \{G_k\}_{k=T-L+1}^T = \{G_{T-L+1}, \dots, G_T\} \quad (5)$$

The idea of leveraging only the most recent graphs in the time-series was first explored in [23] and can be applied to any of the proposed temporal models in this section.

5.2.3 Temporal Reachability Graph (TRG) Model. The temporal reachability graph (TRG) is a graph derived from the timestamped edge stream where a link is added between two nodes if they are temporally connected. More formally, an edge (u, v) in the TRG model indicates the existence of a temporal walk from u to v in the original graph. The formal definition is given as follows.

DEFINITION 5 (TEMPORAL REACHABILITY GRAPH). Given an interval $\mathcal{I} \in \mathbb{R}^+$, the temporal reachability graph $G_R = (V, E_R)$ is defined as a directed graph where the edge $(u, v) \in E_R$ denotes the existence of a temporal walk leaving u and arriving v within that interval. We denote the number of edges in \mathcal{I} as ω (which could be defined based on $\{\tau, \epsilon\}$ -graph time-series).

A TRG is a static unweighted graph where each edge indicates a temporally-valid walk reaching from the source to the destination. However, it does not capture the strength of reachability. For example in Fig. 3a, the walk $\{A, B, C\}$ takes two timestamps while $\{A, B, D\}$ takes four. Intuitively D is harder to reach than C from node A due to less temporal continuity. Vanilla TRG fails to capture such property since all the edges are equally important (shown

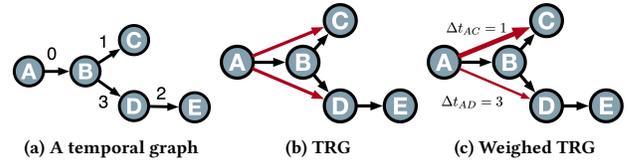


Figure 3: A toy temporal graph (a) and its temporal reachability modeling TRG (b) and WTRG (c). (b) An edge in the vanilla TRG represents a temporally-valid walk. The red edges represents the length-2 walks $\{A, B, C\}$ and $\{A, B, D\}$ in the original graph (c) WTRG extends TRG by assigning weights to indicate the temporal closeness e.g., $\{A, B, C\}$ has higher weights than $\{A, B, D\}$ as C is temporally closer to A than D ($\Delta t_{AC} < \Delta t_{AD}$), which reflects stronger temporal continuity.

in 3b). This would potentially affect the proximity-based embedding methods as they are based on the closeness of nodes in the graph. To overcome this drawback, we propose an extension of TRG called Weighted TRG (WTRG) that encapsulates the strength of reachability in the graph weights. We define the strength of reachability between a pair of nodes (i, j) as a function of both the number of temporally-valid paths and the timestamp difference. The weighting function is given as follows.

$$g_{i,j} = \sum_{w \in \mathcal{W}} e^{-(\Delta t_{i,j}|w)} \quad (6)$$

where w is a specific temporally-valid walk from i to j , and $\Delta t_{i,j}$ denotes the temporal delay reaching from i to j along that walk. We depict the process of deriving WTRG in Algorithm 1. The cornerstone of the algorithm is the temporally-reachable neighborhood N_i^R that records nodes that can be reached by i as well as the latest timestamps associated with the temporal paths. Formally, we define N_i^R as follows.

DEFINITION 6 (TEMPORALLY REACHABLE NEIGHBORHOOD). Given a node i , its temporally reachable neighborhood N_i^R is defined as the set of tuples $\{(j, t_j)\}$ where j is the node reachable from i following a temporally-valid walk and t_j is the timestamp of the edge reaching j in that walk.

Given an input temporal edge (i, j, t) , Algorithm 1 loops through reachable neighbors in N_i^R to add edges in E_R and updates the weights based on Eq. (6) (line 5-8). It also adds (i, j) to the WTRG

Algorithm 1 Weighted Temporal Reachability Graph

```

1: procedure TEMPORALREACH( $G = (V, E)$ )
2:   Set  $E_R = \emptyset$ , sort  $E_T$  in reverse time order
3:   while next edge  $(i, j, t) \in E$  do
4:     for  $(k, t_k) \in N_j^R$  do
5:        $E_R \leftarrow E_R \cup \{(i, k)\}$ 
6:        $g_{i,k} = g_{i,k} + e^{-(t_k-t)}$ 
7:        $N_i^R \leftarrow N_i^R \cup \{(k, t_k)\}$ 
8:      $E_R \leftarrow E_R \cup \{(i, j)\}$ 
9:      $g_{i,j} = g_{i,j} + 1$   $\triangleright \Delta t_{i,j} = 0$  as  $i, j$  are adjacent
10:     $N_i^R \leftarrow N_i^R \cup \{(j, t)\}$ 
11:   end while
12:   return  $G_R = (V, E_R, g)$ 

```

as well as the immediate weight (line 9-11). Overall, the computational complexity of the algorithm is $O(|E| \max d(N^R))$, where $\max d(N^R)$ is the maximum degree of a node in WTRG. While the derived WTRG can be dense with huge amounts of reachable neighbors, we show that this number is bounded by ω , which is the size of the interval associated with the WTRG. Accordingly, the computational complexity of the algorithm is denoted as $O(|E|\omega)$.

PROPERTY 1. *The number of edges in G_R is bounded by the number of temporally-valid walks in G .*

Based on Def. 5, an edge $(u, v) \in E_R$ indicates a temporally-valid walk reaching from u to v in G . However, this edge corresponds to multiple unique temporal walks with different intermediate nodes and associated timestamps, therefore, $|E_R|$ is no more than the number of temporally-valid walks in G . Let N_i^R denote the temporally reachable nodes of i , $\Delta(G_R) = \max\{d(N_1^R), \dots, d(N_n^R)\}$ is the maximum degree of a node in G_R , and ω is the window size. Then

$$|N_i^R| \leq \Delta(G_R) \leq \omega \quad (7)$$

According to Def. 5, a TRG is comprised by edges within the interval with size ω . These edges comprise upto ω different temporal walks originating from a specific node i . Therefore, based on Property 1, the number of edges originating from node i is bounded by the number of temporally-valid walks, which is ω .

To combine the embeddings over the graph time-series, we follow Algorithm 2.

5.3 Temporal Embeddings

5.3.1 Base embedding methods. Given the graph time-series representation and temporal model (Section 5.1-5.2), the proposed approach can leverage any existing static embedding method to derive time-dependent node embeddings that capture the important temporal dependencies between the nodes as well as the temporal structural (role-based) and proximity-based properties [22]. We use the proposed framework to generalize a wide variety of static base embedding methods including both community-based and role-based structural node embedding methods [22]. Namely, they are: (1) LINE [26], (2) Node2vec [9], (3) Graph2Gaussian [3], (4) struc2vec [19], (5) Role2vec [1], (6) Graphwave [6], and (7) multilens [12]. We provide the detailed configuration of each individual method in Section 6.1.3 for the experiments. Among these static

Algorithm 2 General Framework for Temporal Embeddings

Input: ϵ or τ for deriving the graph time-series representation, base embedding method f (e.g., GraphWave, role2vec)

```

1: Construct a graph time-series  $\mathcal{G} = \{G_1, G_2, \dots, G_T\}$  using a graph time-series representation  $\{\tau, \epsilon\}$  from Section 5.1.
2: Initialize  $Z_0$  to all zeros
3: for each  $G_k \in \mathcal{G}$  do  $\triangleright$  for  $k = 1, 2, \dots$ 
4:   Use Alg. 1 to derive the temporal reachability graph for  $G_k$ 
5:   Compute node embedding matrix  $Z_k$  using the base embedding method  $f$  with the temporal reachability graph from Alg. 1
6:   Concatenate or aggregate (using sum, mean, etc.) the embedding matrix, e.g.,  $\bar{Z}_k = (1 - \theta)\bar{Z}_{k-1} + \theta Z_k$  where  $\bar{Z}_k$  is the temporally weighted embedding using the above exponential weighting kernel  $\mathbb{K}(\cdot)$  and  $0 \leq \theta \leq 1$  is a hyperparameter controlling the importance of past information relative to more recent (Section 5.3.2).
7: return  $\bar{Z}_k$  (temporally weighted embeddings using  $\mathbb{K}$  and  $\theta$ ) or  $Z = [Z_1 Z_2 \dots Z_T]$  (concatenated embeddings)

```

methods, approaches (1-3) are community/proximity-based and (4-6) are role-based. Method (7) is a hybrid that is based on structural similarity of node-central subgraphs.

5.3.2 Temporal fusion. Given the time-series of node embeddings $\{Z_k\}_{k=1}^T$, we explore two temporal fusion techniques.

Concatenation: Given a time-series of embeddings, one simple approach to obtain a final embedding is to concatenate the embeddings as follows: $Z = [Z_1 \dots Z_T]$. We could further weight the embeddings based on temporal recency, i.e., under-weighting node embeddings that occur in the distant past since they are not as important as the more recent ones for prediction.

Temporally weighting: This technique aggregates (e.g., sum, mean) the embedding matrix, e.g., $\bar{Z}_k = (1 - \theta)\bar{Z}_{k-1} + \theta Z_k$ where \bar{Z}_k is the temporally weighted embedding using the above exponential weighting kernel $\mathbb{K}(\cdot)$. $0 \leq \theta \leq 1$ is a hyperparameter controlling the importance of past information relative to more recent.

6 EXPERIMENTS

In this section, we systematically investigate the effectiveness of each component in the framework, i.e., the different graph time-series representations (Section 6.3, temporal network models (Section 6.4), and the new dynamic node embedding methods generalized using the proposed framework (Section 6.5). More specifically, we aim to explore the following research questions:

- **Q1** Does the τ -graph time-series representation that previous methods use perform better than the proposed ϵ -graph time-series representation?
- **Q2** How effective is the proposed WTRG model comparing with the vanilla TRG model? What temporal models are most useful for incorporating temporal dependencies into static embedding methods?
- **Q3** Are the dynamic embedding methods generalized via the framework useful for temporal prediction? How do they compare to the state-of-the-art dynamic methods?

6.1 Experimental Setup

6.1.1 Data. We learn node embeddings from the graph time-series starting from roughly $\frac{1}{3}$ of the timespans. For example, for the

bitcoin dataset, we train the classifier based on node embeddings derived from month 20 to month 25 out of 63 months, inclusive. This ensures that there are sufficient edges for training. For all datasets, we perform training on the first 6 graphs and predict links on the 7th graph. Depending on the time-scale shown in Table 2, they represent 6 months (enron, bitcoin, wiki-elec and overflow), weeks (fb-forum and reality-call), or days (wiki-edit and contact-dublin). We create evaluation examples from the links in the 7th graph and an equal number of randomly sampled pairs of unconnected nodes as negative samples [24].

6.1.2 Model configuration and variants. We consider the task of link prediction over time and systematically compare the performance of different temporal network models and representations. Given a set of timestamped edges up to timestamp T , i.e., $\mathcal{G} = \{G_1, \dots, G_T\}$, the temporal link prediction task aims to predict the future links that will form in G_{T+1} . We first follow the conventional setup to construct the τ -graph time-series $\mathcal{G}^\tau = \{G_1, \dots, G_T\}$ for model training and G_{T+1} for testing, where each snapshot $G_k (k \in \{1, 2, \dots, T\})$ represents edges that occur within a consistent time scale shown in Table 2. Then we construct the ϵ -graph time-series representation \mathcal{G}^ϵ . For fair comparison, we set $\epsilon = |E_{T+1}|$ to ensure the trained models based on both ϵ - and τ -based temporal networks are used to predict links in the same hold-out test set G_{T+1} . Thus, graphs in the ϵ -graph time-series $\mathcal{G}^\epsilon = \{G_1, \dots, G_T\}$ and G_{T+1} are also consistent with respect to the ϵ representation, where $|E_1| = |E_2| = \dots = |E_{T+1}|$.

For each $\{\epsilon, \tau\}$ -graph time-series representation, we select a temporal network model from $\{SG, TSG, WTRG\}$ and a base embedding method using the framework. Therefore, we have totally 6 dynamic variants: $\{SG-\epsilon, TSG-\epsilon, WTRG-\epsilon, SG-\tau, TSG-\tau, WTRG-\tau\}$. To train the classifier, we applying these dynamic variants to derive node embeddings and feed them to the logistic regression model for prediction with regularization strength 1.0 and stopping criteria 10^{-4} . Following [4], we concatenate the node embeddings z_i and z_j to obtain an edge embedding $z_{ij} = [z_i z_j]$. For temporal fusion, we use the temporally weighting technique from Section 5.3.2 with $\theta = 0.8$ for dimensional consistency. The TSG decay parameter α is set to 0.8 for computational fairness. For all experiments, we perform 3 runs and report the average.

6.1.3 Base and Dynamic method configuration. We configured all the base methods to achieve the best performance according to the respective papers. For all the static methods based on random walks (i.e., node2vec, struc2vec), we perform 20 walks with the maximum walk length $L = 20$. For node2vec, we perform grid search over $p, q \in \{0.25, 0.50, 1, 2, 4\}$ as mentioned in [9] and report the best performance. For LINE and Multi-Lens, we incorporate 2nd-order proximity in the graph. For role2vec, we leverage the node degree as the feature for roles. For Graphwave, we perform the method to automatically select the scaling parameter with exact heat kernel matrix calculation. We set the final embedding with dimension $K = 128$ for evaluation, and leverages the weighted summation fusion approach so that the embedding dimensions of individual graphs are fixed to be the same.

For the state-of-the-art dynamic embedding methods, we follow the configuration given by the paper/code repository. Specifically, for CTDNE, we set #walks= 10, the walking length $L = 20$. For

node2bits, we perform short-term temporal random-walk with scope to be 3. The the #walks and the walking length are set to be the same as CTDNE. For DANE, we leverage both the offline computation model to derive node embeddings based on the first 6 graphs, and the online model to derive node embeddings for the 6th graph based on the first 5. We set the intermediate embedding dimensions to be 100 for both models and report the best performance. For TIMERS, we set the tolerance threshold value that is used to restart the optimal SVD calculation to be 0.17 as provided in the code repository. For DyAE/DyAERNN, we leverage the 2-layer auto-encoder/decoder with 400 and 200 units, respectively. We set the regularization hyperparameter to be 10^{-6} , bounding ratio for number of units in consecutive layers to be 0.3 as suggested in the paper, and perform grid search in the range of $\pm 10\%$ of the default value. In the learning stage, the sgd learning rate is set to be 10^{-6} with minibatch size to be 100. Lastly, for DySAT, we leverage the base model and perform grid search on the default hyperparameters in the range of $\pm 10\%$ of the default values.

For reproducibility, we will make the source code and the complete experimental results public upon paper acceptance.

6.2 WTRG vs. TRG

We first study the effectiveness of WTRG model over the vanilla TRG model. As WTRG incorporates the strength of reachability in edge weights, we consider embedding methods that handles weighted graphs, namely, they are node2vec, struc2vec and multilens. We run both methods on two datasets using both TRG and WTRG with τ -graph time series as shown in Table 4.

Table 4: Performance of WTRG over TRG on τ -graph time series

		bitcoin		wiki-elec	
Method	Metric	TRG	WTRG	TRG	WTRG
node2vec	AUC	0.9214	0.9239	0.7348	0.7344
	ACC	0.8294	0.8412	0.6171	0.6144
	F1	0.8285	0.8408	0.5909	0.5889
struc2vec	AUC	0.9274	0.9301	0.7840	0.7933
	ACC	0.7959	0.8109	0.6583	0.6703
	F1	0.7925	0.8081	0.6388	0.6534
multilens	AUC	0.9226	0.9389	0.8106	0.8143
	ACC	0.8656	0.8793	0.7438	0.7539
	F1	0.8655	0.8792	0.7385	0.7493

The first observation from Table 4 is that structure-based embedding methods tend to outperform node2vec, the proximity-based method. In addition, we observe that WTRG improves most embedding methods in link prediction, except for node2vec on wiki-elec dataset. One possible reason is that the random walker in WTRG are more likely to visit nodes that are close in time, and thus limiting the derived embeddings to incorporate distant neighborhood information. We put this deep study of WTRG in the future work. Nevertheless, for embedding methods that are based on structural information, WTRG outperforms TRG by 0.8% in AUC, 1.3% in ACC, and 1.4% in F1 score on average.

As we observe that the WTRG model tends to outperform the vanilla TRG model, we use the WTRG model across the experiments from Section 6.3 to Section 6.5.

Table 5: Mean rank (and std.) of the temporal network models across all base embedding methods and graphs based on AUC, ACC and F1, lower is better. The top-3 temporal network models are based on the new ϵ -graph time-series representation (fixed #edges)

TEMPORAL MODEL	AUC	MEAN RANK (MR) ACC	F1
WTRG- ϵ	2.30 \pm 2.16	2.73 \pm 1.95	2.66 \pm 1.90
TSG- ϵ	2.43 \pm 1.84	2.61 \pm 2.19	2.70 \pm 2.24
SG- ϵ	2.57 \pm 1.64	2.66 \pm 1.86	2.59 \pm 1.88
WTRG- τ	2.80 \pm 1.96	2.80 \pm 1.98	2.86 \pm 1.99
SG- τ	2.95 \pm 1.91	2.84 \pm 1.87	2.82 \pm 1.83
TSG- τ	3.63 \pm 1.75	3.46 \pm 1.87	3.45 \pm 1.80
SG	4.32 \pm 1.88	3.89 \pm 1.91	3.93 \pm 1.94

6.3 Fixed #edges (ϵ) vs. time-scale (τ)

In this section, we investigate the effectiveness of different graph time-series representations (Q1). Due to the massive amount of experimental results, we first define 2 evaluation metrics for this experiment. These newly proposed measurement are for readers to have a clear overview of the comparison results across all components in the proposed framework across all the datasets.

We first evaluate the general performance of each temporal model through the mean ranking (and std) across all datasets and embedding methods in terms of the AUC, ACC and F1 score. We leverage the following metrics to better interpret the results. Let $y_{jk} \in \mathbb{R}^{|\mathcal{M}|}$ denote the vector of AUC (or ACC, F1) scores of the temporal models \mathcal{M} for an embedding method $f_j \in \mathcal{F}$ and graph dataset k . Further, let $\pi(y_{jk}, M_i)$ denote the rank of the temporal model $M_i \in \mathcal{M}$ for a given embedding method f_j and graph dataset $d_k \in \mathcal{D}$. The mean rank is computed as

$$MR_i = \frac{1}{|\mathcal{D}||\mathcal{F}|} \sum_{d_k \in \mathcal{D}} \sum_{f_j \in \mathcal{F}} \pi(y_{jk}, M_i) \quad (8)$$

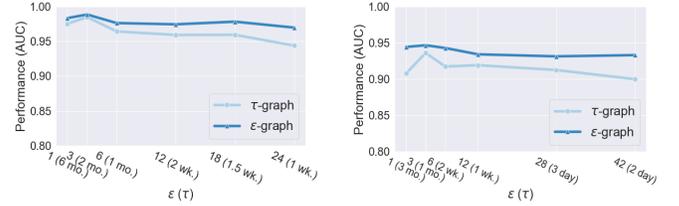
Therefore, smaller values of MR indicates better performance of that model. We report the results in Table 5.

In addition to the general performance, we also provide an intuitive ranking of the temporal network models based on the number of times each model performs best following [21]. This metric reflects the occurrence of a model to be optimal. We define this score s_i of temporal model M_i as

$$s_i = \sum_{d_k \in \mathcal{D}} \sum_{f_j \in \mathcal{F}} \mathbb{I}\{\pi(y_{jk}, M_i) = 1\} \quad (9)$$

where $\mathbb{I}\{\pi(y_{jk}, M_i) = 1\}$ returns 1 if $\pi(y_{jk}, M_i) = 1$ and 0 otherwise. $\mathbb{I}\{\pi(y_{jk}, M_i) = 1\}$ indicates that the temporal model M_i performs best for the given graph dataset d_k and base embedding method f_j . Thus, s_i denotes the total score of model M_i based on the number of times temporal model M_i appeared first in the ranking across all base embedding methods and graph datasets.

Performance. Based on the results shown in Table 5, our first observation is that the top-3 temporal models are those that use the proposed ϵ -graph time-series representation. These models perform comparably well in terms of AUC, ACC and F1 and are in general better than τ -graph time-series representation used in previous work. This finding indicates the general effectiveness of ϵ -graph time-series in representing the temporal network. We also compute an overall score by summing over each s_i for all evaluation



(a) Prediction performance on bitcoin (b) Prediction performance on fb-forum

Figure 4: Sensitivity analysis. Link prediction performance on τ - and ϵ -graph. The τ -graphs are created based on different timescales, the ϵ -graphs are created via equal division. mo.: month. wk.: week.

criterion (bottom row in Table 6). We observe that the top models are ϵ -based, which demonstrate the effectiveness of ϵ -graph time-series in capturing the graph structural changes over edge frequency changes in prediction tasks.

Sensitivity Analysis. We also conduct a parameter sensitivity analysis on two datasets, bitcoin and fb-forum, to evaluate the impact of different values of τ and ϵ on the overall performance. For bitcoin, we train our model on temporal data spanning 6 months and test on the 7th month. Forfb-forum, we train on 12 weeks and test on the 13rd week. We create the τ -graph time series using different scales (e.g., months, weeks, days), and generate the same number of ϵ -graphs with equal number of edges. Based on the result shown in Figure 4, we observe that our model running on the ϵ -graph time series consistently outperforms the τ -graph time series, while being more robust. This also indicates that in practice, the ϵ -graph time series can be used as an alternative to create temporal graph snapshots for various mining tasks, especially in preliminary graph analysis when the optimal timescale is undetermined.

RESULT 1. Overall, the proposed ϵ -graph time-series representation based on a fixed number of edges outperforms the time-scale τ -graph time-series across different scales, while being more robust.

6.4 Temporal Model Comparison

To answer Q2, we follow the formulation in Section 6.3 to quantitatively evaluate and rank the temporal models according to their effectiveness in prediction. We show the complete performance of temporal network models that perform the best following Equation (9) with respect to individual datasets in Table 6 to supplement the mean ranking in Table 5.

Notably, the TSG- ϵ model has the highest # of first ranks across all datasets, especially on datasets with short timespans (i.e., wiki-edits and contacts). It also has the highest # of first ranking, in terms of ACC and F1, which indicates that this model is generally promising but at the same time less stable than the other ϵ -based models. We also confirm this finding in Table 5 as it shows relatively higher variance of ranking. WTRG- ϵ performs the second best and is a close second to TSG- ϵ on datasets with long timespans. This is potentially due to how they model the temporal recency: TSG models the past information with exponential decay (4), while WTRG models it with the absolute temporal difference (6). Thus, TSG could still capture the temporal evolution within a relatively short period of time. Nevertheless, both TSG and WTRG perform well on all datasets even though spikes and fluctuation are observed such as fb-forum (Figure 1).

Table 6: Temporal model performance across the temporal graphs. Each (i, j) is the # of times temporal model $M_j \in \mathcal{M}$ in graph G_i performed best comparing to the other models across all base embedding methods $f \in \mathcal{F}$ and evaluation criterion. We bold the temporal model that performs best overall for each graph.

	TSG- ϵ	WTRG- ϵ	SG- ϵ	SG- τ	WTRG- τ	SG	TSG- τ
bitcoin	6	6	4	5	0	0	0
stackoverflow	1	4	3	3	9	0	1
enron	4	1	1	3	8	4	0
wiki-elec	2	6	7	6	0	0	0
fb-forum	10	10	0	1	0	0	0
wiki-edit	7	3	2	3	2	2	2
reality-call	1	0	2	4	6	4	4
contacts-dublin	9	2	8	1	1	0	0
overall score	40	32	27	26	26	10	7

It can be seen that there is not a single temporal model that prevails across all datasets. On the other hand, the WTRG model tends to perform well regardless of the timescales in graph representation, while TSG model tends to perform well on graphs with short timespans. Besides, the temporal models that are combined with the proposed ϵ -graph time-series representation tend to outperform their other τ -counterparts, which is consistent with our previous findings from Section 6.3.

RESULT 2. Out of all models, WTRG- ϵ and TSG- ϵ tend to perform the best. Empirically, WTRG- ϵ is more stable overall (Table 5) and TSG- ϵ performs well on datasets with short timespans (Table 6).

6.5 Dynamic Embeddings: Variants vs. State-of-the-art

To answer Q3, we first use the framework to derive new dynamic embedding methods (by selecting the representation, temporal model, base embedding method, etc.), then we compare their performance to the state-of-the-art dynamic embedding methods on all 8 datasets. One would presumably expect that the state-of-the-art methods for dynamic node embeddings will outperform the dynamic embedding methods generalized by our framework. This is because the state-of-the-art methods are typically more complex and have been designed specifically for learning such dynamic node embeddings. We use 9 recent state-of-the-art dynamic methods during 2017 ~ 2020 as baselines, including CTDNE [17], node2bits [11], DANE [16], DynGem [8] TIMERS [28], DynAE/DynAERNN [7], DySAT [24], DyHATR [27], and EvolveGCN[18].

Figure 5 shows the mean AUC for each method where the average is taken over all graphs investigated. As representative dynamic embedding methods from the proposed framework, we use 4 dynamic embedding variants of struc2vec (s2v-TSG- ϵ/τ , s2v-WTRG- ϵ/τ) and 4 variants of MultiLENS (ML-TSG- ϵ/τ , ML-WTRG- ϵ/τ). Strikingly, we observe that the dynamic embedding methods from the framework perform comparably or even better than the state-of-the-art methods that are designed particularly for temporal graphs and time-series prediction. In particular, ML-TSG- ϵ performs best with a mean gain of 12.34% followed by s2v-TSG- ϵ with a gain in AUC of 10.97%. Also, we note that our proposed framework is computationally efficient. Taking multilens as an base embedding method, our proposed framework has the computational complexity $\mathcal{O}(|E||\mathcal{G}|)$ where $|\mathcal{G}|$ is the number of graphs in the time-series,

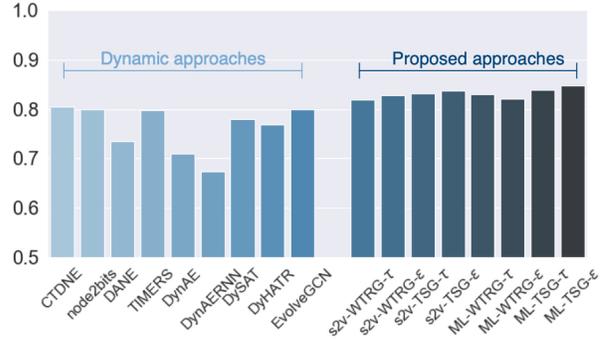


Figure 5: Predictive results comparing the dynamic embedding methods to our framework. The variants of our proposed framework approximate well to approaches specifically designed for temporal graphs with comparable or even better performance (ML = multilens, s2v = struc2vec).

and the number of parameters to learn is $\mathcal{O}(|V||\mathcal{G}|)$. This is also validated empirically throughout our experiment.

RESULT 3. The dynamic embeddings derived from our framework leveraging conventional static embedding methods (Section 5) perform better than state-of-the-art dynamic embedding methods.

Notably, in this experiment we do not aim to show substantial improvement of our framework over all the dynamic approaches, especially those based on deep learning, since node features are not used. Instead, these results show that our proposed framework could capture the graph structures and temporal dependency at least as good as those recent dynamic approaches with less complexity (i.e., no transitional or latent variables). Unlike methods that are based on complex models as “black boxes”, the components of our proposed framework are easy-to-understand, which further motivates its usage for practitioners in predictive applications.

7 CONCLUSION

Despite the recent increasing interest in temporal networks in the field of representation learning, there has been relatively little work that systematically studies the properties of temporal network models and their cornerstones, the graph time-series representations. This works attempts to fill this gap by proposing a general yet powerful framework. Specifically, we introduce the notion of ϵ -graph time-series to address data imbalance that arise with the traditional way of deriving a graph time-series based on a—sometimes arbitrary—time-scale (e.g., 1 day or 1 week). We find that the ϵ -graph time-series is beneficial to most existing embedding methods in temporal link prediction. Furthermore, our proposed framework gives rise to new dynamic embedding methods by combining these graph time-series representations, new temporal models, and base static embedding methods. We find that although there is no single temporal model (or embedding method) that could prevail on any dataset, our proposed WTRG model and TSG model along with the ϵ time-series tend to perform the best across all datasets studied. We further show that these dynamic embedding approaches from our framework outperform recent, powerful dynamic embedding methods. We expect that the findings of this work will be useful for practitioners in various applications or developing novel and better embedding methods for temporal networks.

REFERENCES

- [1] N. K. Ahmed, R. A. Rossi, R. Zhou, J. B. Lee, X. Kong, T. L. Willke, and H. Eldardiry. Learning role-based graph embeddings. In *IJCAI StarAI*, 2018.
- [2] S. Bhadra and A. Ferreira. Complexity of connected components in evolving graphs and the computation of multicast trees in dynamic networks. *ADHOC-NOW*, pages 259–270, 2003.
- [3] A. Bojchevski and S. Günnemann. Deep gaussian embedding of graphs: Unsupervised inductive learning via ranking. In *ICLR*, 2018.
- [4] S. Cao, W. Lu, and Q. Xu. Grarep: Learning graph representations with global structural information. In *CIKM*, pages 891–900. ACM, 2015.
- [5] K. G. Coffman and A. M. Odlyzko. Growth of the internet. In *Optical fiber telecommunications IV-B*, pages 17–56. Elsevier, 2002.
- [6] C. Donnat, M. Zitnik, D. Hallac, and J. Leskovec. Learning structural node embeddings via diffusion wavelets. In *KDD*, pages 1320–1329. ACM, 2018.
- [7] P. Goyal, S. R. Chhetri, and A. Canedo. dyngraph2vec: Capturing network dynamics using dynamic graph representation learning. *Knowledge-Based Systems*, page 104816, 2019.
- [8] P. Goyal, N. Kamra, X. He, and Y. Liu. Dyngem: Deep embedding method for dynamic graphs. *arXiv preprint arXiv:1805.11273*, 2018.
- [9] A. Grover and J. Leskovec. node2vec: Scalable feature learning for networks. In *KDD*, pages 855–864, 2016.
- [10] R. Hisano. Semi-supervised graph embedding approach to dynamic link prediction, 2016.
- [11] D. Jin, M. Heimann, R. Rossi, and D. Koutra. node2bits: Compact time-and attribute-aware node representations for user stitching. In *ECML PKDD*, page 22, 2019.
- [12] D. Jin, R. A. Rossi, E. Koh, S. Kim, A. Rao, and D. Koutra. Latent network summarization: Bridging network embedding and summarization. In *KDD*. ACM, 2019.
- [13] S. Kumar, X. Zhang, and J. Leskovec. Predicting dynamic embedding trajectory in temporal interaction networks. In *KDD*, pages 1269–1278. ACM, 2019.
- [14] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graphs over time: Densification laws, shrinking diameters and possible explanations. In *KDD*, 2005.
- [15] J. Leskovec and A. Krevl. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>, June 2014.
- [16] J. Li, H. Dani, X. Hu, J. Tang, Y. Chang, and H. Liu. Attributed network embedding for learning in a dynamic environment. In *CIKM*, pages 387–396. ACM, 2017.
- [17] G. H. Nguyen, J. B. Lee, R. A. Rossi, N. K. Ahmed, E. Koh, and S. Kim. Continuous-time dynamic network embeddings. In *WWW BigNet*, 2018.
- [18] A. Pareja, G. Domeniconi, J. Chen, T. Ma, T. Suzumura, H. Kanezashi, T. Kaler, T. Schardl, and C. Leiserson. Evolvegnn: Evolving graph convolutional networks for dynamic graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5363–5370, 2020.
- [19] L. F. Ribeiro, P. H. Saverese, and D. R. Figueiredo. Struc2vec: Learning node representations from structural identity. In *KDD*, 2017.
- [20] R. A. Rossi and N. K. Ahmed. The network data repository with interactive graph analytics and visualization. In *AAAI*, 2015.
- [21] R. A. Rossi, N. K. Ahmed, H. Eldardiry, and R. Zhou. Similarity-based multi-label learning. In *IJCNN*, pages 1–8, 2018.
- [22] R. A. Rossi, D. Jin, S. Kim, N. K. Ahmed, D. Koutra, and J. B. Lee. From community to role-based graph embeddings. In *arXiv:1908.08572*, 2019.
- [23] R. A. Rossi and J. Neville. Time-evolving relational classification and ensemble methods. In *PAKDD*, volume 7301, pages 1–13. Springer, 2012.
- [24] A. Sankar, Y. Wu, L. Gou, W. Zhang, and H. Yang. Dysat: Deep neural representation learning on dynamic graphs via self-attention networks. In *WSDM*, pages 519–527, 2020.
- [25] U. Singer, I. Guy, and K. Radinsky. Node embedding over temporal graphs. In *IJCAI*, pages 4605–4612, 7 2019.
- [26] J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan, and Q. Mei. LINE: Large-scale Information Network Embedding. In *WWW*, pages 1067–1077, 2015.
- [27] H. Xue, L. Yang, W. Jiang, Y. Wei, Y. Hu, and Y. Lin. Modeling dynamic heterogeneous network for link prediction using hierarchical attention with temporal rnn. In *Proceedings of the 2020 European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD)*, 2020.
- [28] Z. Zhang, P. Cui, J. Pei, X. Wang, and W. Zhu. Timers: Error-bounded svd restart on dynamic networks. In *AAAI*, 2018.
- [29] L. Zhou, Y. Yang, X. Ren, F. Wu, and Y. Zhuang. Dynamic network embedding by modeling triadic closure process. In *AAAI*, 2018.
- [30] Y. Zuo, G. Liu, H. Lin, J. Guo, X. Hu, and J. Wu. Embedding temporal network via neighborhood formation. In *KDD*, pages 2857–2866, 2018.