

Representation Learning in Continuous-Time Dynamic Signed Networks

Kartik Sharma* Georgia Institute of Technology Atlanta, Georgia, USA ksartik@gatech.edu Mohit Raghavendra[†]* Georgia Institute of Technology Atlanta, Georgia, USA mraghavendra6@gatech.edu Yeon-Chang Lee Georgia Institute of Technology Atlanta, Georgia, USA yeonchang@gatech.edu

Anand Kumar M NIT Karnataka India m_anandkumar@nitk.edu.in Srijan Kumar Georgia Institute of Technology Atlanta, Georgia, USA srijan@gatech.edu

ABSTRACT

Signed networks allow us to model conflicting relationships and interactions, such as friend/enemy and support/oppose. These signed interactions happen in real-time. Modeling such dynamics of signed networks is crucial to understanding the evolution of polarization in the network and enabling effective prediction of the signed structure (i.e., link signs) in the future. However, existing works have modeled either (static) signed networks or dynamic (unsigned) networks but not dynamic signed networks. Since both sign and dynamics inform the graph structure in different ways, it is non-trivial to model how to combine the two features. In this work, we propose a new Graph Neural Network (GNN)-based approach to model dynamic signed networks, named SEMBA: Signed link's Evolution using Memory modules and Balanced Aggregation. Here, the idea is to incorporate the signs of temporal interactions using separate modules guided by balance theory and to evolve the embeddings from a higher-order neighborhood. Experiments on 4 real-world datasets and 3 different tasks demonstrate that SEMBA consistently and significantly outperforms the baselines by up to 80% on the tasks of predicting signs of future links while matching the stateof-the-art performance on predicting existence of these links in the future. We find that this improvement is due specifically to superior performance of SEMBA on the minority negative class. Code is made available at https://github.com/claws-lab/semba.

CCS CONCEPTS

• Computing methodologies → Supervised learning; Neural networks; • Information systems → Data mining.

KEYWORDS

Graph Neural Networks; Signed Networks; Dynamic Graphs

[†]This work was done during MR's internship at Georgia Tech

CIKM '23, October 21–25, 2023, Birmingham, United Kingdom

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0124-5/23/10.

https://doi.org/10.1145/3583780.3615032

ACM Reference Format:

Kartik Sharma, Mohit Raghavendra, Yeon-Chang Lee, Anand Kumar M, and Srijan Kumar. 2023. Representation Learning in Continuous-Time Dynamic Signed Networks. In Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23), October 21–25, 2023, Birmingham, United Kingdom. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3583780.3615032

1 INTRODUCTION

Background. Social interactions are often signed and dynamic, representing the evolving nature of trust/distrust, support/oppose, agreement/disagreement, or positive/negative sentiments among people. Signed networks have been used to measure polarization in social networks [2, 5], performance in financial networks [3], and predicting international conflicts [10]. Dynamic signed networks model conflicting relationships that evolve over time, as shown in Figure 1. For instance, online social media discussions involve people supporting/opposing a claim, which continuously evolve as new users and new replies are made in real time. Dynamic signed networks have several societal applications such as predicting veracity of posts [40], toxicity in conversations [38], conflict across communities [22], and analyzing polarization in debates [2, 26, 27, 36]. Thus, it is important to create predictive models for dynamic signed networks so that one can devise intervention strategies against future unwanted behavior, maintaining a healthy online environment [4, 35, 44].

Structural balance theory models the organization of signed networks [6]. It states that signed relationships tend to form balanced triads such that "a friend of my friend is my friend," "a friend of my enemy is my enemy," "an enemy of my friend is my enemy," and "an enemy of my enemy is my friend." Several Signed Graph Neural Networks (GNNs) have been proposed to learn effective representations of static signed graphs, while incorporating balance theory [9, 13–15, 31, 32, 41].

Challenges. Modeling dynamic signed networks with existing (static) signed GNNs faces two challenges.

(C1) Temporal-Awareness, i.e., modeling the sequential order and evolution of signed interactions to learn node representations. Static signed GNNs only learn the network structure at a given timestamp [9, 13, 15], ignoring the temporal order of the interactions. Additionally, while signed GNNs can incorporate balance

^{*}KS and MR contributed equally to this work.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).



Figure 1: Visualizing Dynamic Signed Networks.

theory in static networks, they are unable to follow balance theory when the network is dynamic [1, 10, 16, 34, 45].

(C2) Staleness, i.e., in the absence of recent direct interactions of a node, its representation becomes stale. Staleness is a pervasive issue since real-world networks are long-tailed and sparse, i.e., most nodes only have a few direct interactions.

Meanwhile, there has been a surge of interest in learning temporal representations of nodes in dynamic networks by developing dynamic GNN models [25, 37, 42]. However, these model architectures are not **(C3) Sign-Aware**. In other words, since these models have been developed specifically for unsigned networks, they treat negative and positive edges the same and thus, fail to incorporate the higher-order organization induced by the link signs (i.e., structural balance theory). Due to their homophily assumptions in learning [9], these models fail to model heterophilic network organizations such as negative links between polarized communities. **Proposed Ideas.** In this work, we aim to <u>effectively learn repre</u>-

sentations of dynamic signed networks while addressing temporal-awareness, staleness, and sign-awareness problems. Towards this goal, we propose a new dynamic signed GNN model, named as **SEMBA**: Signed link Evolution using Memory modules and Balanced Aggregation. SEMBA learns a short-term 'memory' encoding and a long-term 'embedding' per node. SEMBA consists of three key modules:

- (1) *Signed Memories:* SEMBA learns two signed memories (positive and negative) per node to distinctly encode past interactions of a node with its friends and enemies, respectively.
- (2) Balanced Aggregation: when the network evolves by adding a link (u, v), SEMBA updates signed memories of nodes u and v while explicitly incorporating balance theory.
- (3) Long-term Propagation: SEMBA learns a node's embedding by attentively aggregating signed memories over all its neighbors until the current timestep. This overcomes staleness by learning not only from direct interactions at that time but any interactions in any of its past neighbors.

Representations thus formed can be used to solve multiple different tasks on dynamic signed networks: predicting the existence of a future link (i.e., **dynamic link existence prediction**), predicting its sign (i.e., **dynamic link sign prediction**), and predicting both existence and sign in future (i.e., **dynamic link signed existence prediction**). While dynamic link existence prediction has been studied extensively for unsigned dynamic networks in the literature [7, 33, 37, 45, 48], no work has studied it for signed networks. Similarly, while sign prediction has been studied for static signed networks [9, 13, 15, 17, 24], no prior work exists that studies

Kartik Sharma, Mohit Raghavendra, Yeon-Chang Lee, Anand Kumar M, and Srijan Kumar

Table 1: Comparison of existing models with the proposed approach SEMBA based on the challenges that they address.



Figure 2: Balanced and unbalanced triads [6].

the dynamic setting. Due to the inductive nature of the dynamic case, SEMBA generalizes to new nodes and links, which has never been addressed in the signed network literature. We show experimentally that (1) modeling both balance and network evolution is essential to accurately predict sign and existence of a link, and (2) SEMBA consistently and significantly outperforms state-of-the-art baselines for signed networks and for dynamic networks.

Contributions. Our contributions are summarized as follows:

- Novel Problem: We devise a novel problem of representing dynamic signed networks via 3 different prediction tasks.
- **Dynamic Signed GNN:** We propose SEMBA, a GNN-based model to represent dynamic signed networks. To the best of our knowledge, we are the first to devise a model to jointly address temporal-awareness, staleness, and sign-awareness problems.
- **Comprehensive Validation:** We validate the effectiveness of SEMBA by comparing it with 7 representative baselines on 4 realworld datasets under the problems of dynamic link existence and sign prediction. Specifically, SEMBA matches or outperforms the baselines in all tasks across datasets with an improvement of up to 80% in predicting signs.

2 RELATED WORK

In this section, we briefly review existing works for static signed GNNs and dynamic unsigned GNNs. Table 1 compares our model with the existing works. Unlike other models, SEMBA uses an end-to-end GNN-based model to generate node representations of dynamic signed networks.

2.1 Static Signed GNN

Machine learning on graphs has seen widespread interest, especially since the advent of message-passing GNNs. Plenty of GNNs have been adapted to perform a variety of graph-related tasks [47]. Signed networks allow us to model polarity of the relationship which becomes particularly important in a variety of applications [18, 23, 24, 28, 46]. Tasks of predicting the sign of the links are often studied in the literature [9, 24, 29]. With the advent of GNNs, several works have developed GNNs for signed networks [9, 13, 15, 17, 19]. They make use of the inductive bias induced by the neighborhood message passing mechanism, but augments it with ideas from social science like balance (as shown in Figure 2) [6] and status [30]. However, most of these signed GNNs ignore the temporal aspect of the signed network and are limited to static networks. As mentioned in Section 1, this leads to temporal-awareness and staleness problems. Since majority of networks in the real world are dynamic and the temporal interactions are rich sources of data, we present a model that exploits network evolution along with the signed information to make prediction on future signed links.

2.2 Dynamic Unsigned GNN

Dynamic networks are popularly modeled as a stream of continuously time-stamped links, and allow us to capture the network evolution in real time. Methods like DyRep [42], JODIE [25] and TGN [37] have successfully modeled dynamic networks in continuous time, enabling them to solve multiple network-based problems. In particular, the problem of predicting whether a link would arrive in future has been studied extensively in the literature. GNNs are effective in this task as it is often the case that link dynamics is affected by the past interactions in the local neighborhood of the involved nodes.

However, these models do not use the signed information, which we pointed out as the sign-awareness problem in Section 1. In particular, signed interactions are shown to reorganize in well-defined higher-order structures such as balanced triads. However, the existing works based on standard message-passing assume homophily and fail to capture the heterophilic nature of signed networks (e.g., polarized social networks). In this work, we incorporate balance theory to effectively learn representations of continuous-time dynamic networks. To the best of our knowledge, there has been only one prior work [8] in this direction, but not using GNNs. It uses random walks to model the network as an evolving time-series which is then fed into a recurrent neural network for sign prediction. However, it does not make use of balance theory, GNNs for processing the graph structure, or attention mechanism for processing temporal data. We shall build on these recent advancements to create a more accurate model.

3 PROBLEM DEFINITION

A directed continuous-time dynamic signed network at time *t* is denoted by $\mathcal{G}(t) = (\mathcal{V}(t), \mathcal{E}^+(t), \mathcal{E}^-(t))$, where $\mathcal{V}(t)$ denotes the set of nodes, and $\mathcal{E}^+(t)$ (resp. $\mathcal{E}^-(t)$) denotes the positive (resp. negative) links at time *t*. It can be modeled using an initial network $\mathcal{G}_0 = (\mathcal{V}_0, \mathcal{E}_0^+, \mathcal{E}_0^-)$ along with a sequence of time-stamped events (ordered by time), $\mathbb{E} := \mathbb{E}[t_0, t_n] = \{(\tau, ev) : t_0 \le \tau \le t_n\}$, such that t_n is the final time step and ev represents a link addition at timestep τ .¹ More precisely, one can construct a network $\mathcal{G}(t)$ at snapshot *t* by augmenting links to \mathcal{G}_0 such that $\mathcal{G}(t_n) = \mathcal{G}_0 \sqcup \{ev :$ $(t, ev) \in \mathbb{E}, t_0 \le t \le t_n\}$. The operator \sqcup sequentially updates the network based on the events $\{ev\}$. For example, if the event ev is addition of the link (u, v, e_{uv}) , i.e., $\text{Add}_{edge}(u, v, e_{uv})$, and the link signed weight $\mathbf{e}_{uv}(t) > 0$, then network $\mathcal{G} = (\mathcal{V}, \mathcal{E}^+, \mathcal{E}^-)$

Table 2: Notations used in this paper

Notation	Description
$\mathcal{G}(t) = (\mathcal{V}(t), \mathcal{E}(t))$	graph snapshot (vertex set, edge set) at t
$\mathcal{E}^+(t), \mathcal{E}^-(t)$	positive and negative edge set at t
$\mathcal{G}_0, \mathcal{G}_n$	initial and final snapshot
$\mathcal{E}_{train}, \mathcal{E}_{val}, \mathcal{E}_{test}$	training, validation, test set of edges
ev	event like $ADD_{edge}(u, v, \mathbf{e}_{uv})$
$\mathbb{E}[t_i, t_f]$	sequence of time-stamped events from t_i till t_f
$\mathbf{e}_{uv}(t)$	signed weight of (u, v) at t
$\mathbf{x}_{u}(t)$	node features of u at t
$\mathcal{N}_{u}^{k}([0,t])$	k-hop neighbor set of u from time 0 to t
$\mathbf{s}_{u}^{+}(t), \mathbf{s}_{u}^{-}(t)$	positive/negative memory of node u at t
$\mathbf{m}_{u}^{+}(t), \mathbf{m}_{u}^{-}(t)$	positive/negative message encoding of u at t
$\mathbf{M}_{u}^{+}(t),\mathbf{M}_{u}^{-}(t)$	aggregated positive/negative messages on \boldsymbol{u} at \boldsymbol{t}
В	Batch size
$T_B(t)$	First time after t with B interactions in between
$T_{B}^{-1}(t)$	Last time from t that was batched
$p(u, v, t; \Theta)$	probability of edge (u, v) for specified task at t

 $\rightarrow (\mathcal{V}, \mathcal{E}^+ \cup \{(u, v, \mathbf{e}_{uv}(t))\}, \mathcal{E}^-)$. For the sake of simplicity, we will use $\mathcal{E}(t) := \mathcal{E}^+(t) \cup \mathcal{E}^-(t)$ and \mathcal{G}_n to denote $\mathcal{G}(t_n)$.

In this work, we study the effectiveness of signed links' representation in time through the following three problems:

PROBLEM 1 (DYNAMIC LINK EXISTENCE PREDICTION). Given G_0 , $\mathbb{E}[t_1, t_j]$, the goal is to predict whether a link exists between u and v in future, i.e., $(u, v, \cdot) \in \mathcal{E}(t_k)$ for $t_k > t_j$.

PROBLEM 2 (DYNAMIC LINK SIGN PREDICTION). Given \mathcal{G}_0 , $\mathbb{E}[t_1, t_j]$, and a link $(u, v, \mathbf{e}_{uv}) \in \mathcal{E}(t_k)$ for $t_k > t_j$, the goal is to predict its sign, i.e., whether $\mathbf{e}_{uv}(t_k) > 0$ or $\mathbf{e}_{uv}(t_k) < 0$.

PROBLEM 3 (**DYNAMIC LINK SIGNED EXISTENCE PREDICTION**). Given \mathcal{G}_0 , $\mathbb{E}[t_1, t_j]$, the goal is to predict whether a positive, negative or no link exists between u and v in future, i.e., whether $(u, v, \cdot) \in \mathcal{E}^+(t_k), (u, v, \cdot) \in \mathcal{E}^-(t_k)$, or $(u, v, \cdot) \notin \mathcal{E}(t_k)$ for $t_k > t_j$.

4 SEMBA: OUR APPROACH

In this section, we propose SEMBA for representation learning in dynamic signed networks, which stands for **S**igned links Evolution using **M**emory modules and **B**alanced **A**ggregation. We learn two representations per node: node memories to encode temporal signed interactions of the node, and node embeddings to incorporate interactions in its long-term neighbors. Finally, we use the learned embeddings to make the prediction.

4.1 Overview

SEMBA is motivated by the following three observations that help to understand signed links evolution in a dynamic signed network:

- Link Evolution: In a dynamic network, new links of a node *u* at time *t* tend to appear depending on the past interactions in the ego-network of *u* [37, 42], which consists of the focal node *u* and its neighborhood to whom *u* had a connection at some path length before *t*.
- (2) **Structural Balance**: Balance theory [6] guides signed network evolution. If $(u, v) \notin \mathcal{E}$ but $(u, w), (w, v) \in \mathcal{E}$, then the sign

¹We only consider link additions in this work since none of the publicly-available dynamic signed datasets have link deletions. Link removals and node-level events will be considered in future work.



Figure 3: Balanced aggregation: Schema representing how link addition events $(t_1, (1, 2, +))$ and $(t_2, (1, 3, -))$ update the corresponding signed memories at the end-nodes.

of (u, w) and (w, v) would determine how information from v reaches u and consequently, the sign of future link (u, v). As a result, u should consider the information coming from a friend of a friend (resp. a friend of an enemy) differently than the information coming from an enemy of a friend (resp. an enemy of an enemy).

(3) Long-tailed Distribution: It is well known that real-world networks tend to have a power-law degree distribution, which indicates only a few nodes in the network have a high degree, while most of the other nodes have a very low degree. In a dynamic network, it makes the representation of long-tailed nodes that have no *recent* interactions more challenging. In other words, their representations can easily become stale.

Following these observations, our model has three key architectural components — signed memories, balanced aggregation, and long-term propagation.

Signed Memories. A node's memory serves the purpose of representing the node's recent interactions with its immediate positive and negative neighbors, which is useful for predicting new links according to Observation 1. To be able to distinguish between positive versus negative neighbors as per Observation 2, our architecture uses two different memory modules, one positive and one negative, to accumulate the past interactions according to their sign.

Balanced Aggregation. Upon addition of a signed link (u, v), both nodes' positive and negative memories must be updated based on Observation 2:

- If link (u, v) is positive, memories of the same polarity influence each other. Thus, the positive memory of u is updated using the positive memory of v, while u's negative memory is updated using the negative memory of v.
- (2) If link (u, v) is negative, memories of *opposite polarity* provide influence. Thus, the positive memory of u is updated using the negative memory of v, while u's negative memory is updated using the positive memory of v.

Figure 3 visualizes the balanced aggregation mechanism to update a node's memory due to incident positive and negative links.

Furthermore, it should be noted that *balanced aggregation naturally follows higher-order structural rules of balance theory.* Figure 4 shows how updating node memories using neighbor-to-neighbor

Kartik Sharma, Mohit Raghavendra, Yeon-Chang Lee, Anand Kumar M, and Srijan Kumar



Figure 4: Schema showing how updating memories based on balanced aggregation enforces balance theory in triadic relations 2 - 1 - 3 and 2 - 1 - 4 when $t_2 < t_1$, $t_2 < t_3$.

balanced aggregation results in complying with higher-order balance theory relationships. At the first time step t_1 , node 2's positive memory updates node 1's positive memory, which then updates node 3's negative memory at the next time step t_2 . Hence, one can note that positive memory from a friend of an enemy (i.e., node 2 is a friend of an enemy of node 3) updates the negative memory, which stores the information from enemies, thus, in line with balance theory. Similarly, other rules of balance theory can also be realized by local updates in future time steps. For example, node 2's positive memory updates node 4's positive memory via node 1's positive memory by the 'friend of a friend is a friend' rule.

Long-Term Propagation. A node's embedding serves the purpose of propagating signed memories of a node to the neighbors with no active interaction, which is useful for addressing staleness for long-tailed nodes according to Observation 3. To generate a node's embedding, we attentively aggregate signed memories of a node over all of its neighbors in the past and not just the neighbors in the current batch.

4.2 SEMBA Model Pipeline

Figure 5 demonstrates the pipeline of our model. It consists of the following four steps:

- (1) Encode the events as signed messages following balance theory
- (2) Aggregate signed messages in a batch (for efficient processing)
- (3) Update node memories with the aggregated messages
- (4) Generate node embeddings using memories and other features

Temporal Batching. Since networks evolve rapidly, it is computationally inefficient to update node embeddings after each change in the network. Thus, we combine a fixed number *B* of sequential events happening after a time *t* into a temporal batch. Node embeddings are then updated based on all the events included in such a batch. This also allows us to parallelize the computation of events within the batches and make the training process efficient. We create uniform batches of size *B* starting from *t*₀. The first batch stores these *B* interactions $\mathbb{E}[t_0, T_B(t_0)]$ such that $T_B(x)$ is Representation Learning in Continuous-Time Dynamic Signed Networks



Figure 5: Pipeline of our method, SEMBA. The encircled numbers dictate the sequence in which each step happens. The model first adds the new link to the network, followed by generating messages as dictated by balance theory, in step 1. It then aggregates the messages and updates its memory using these new messages in steps 2 and 3. A graph attention based neural network is then used to update the embeddings of the nodes involved in step 4, which are used for downstream prediction tasks.

the smallest timestep after x such that there are exactly B events in $\mathbb{E}[x, T_B(x)]$. Consequently, the kth batch consists of the events $\mathbb{E}\left[T_B^{(k-1)}(t_0), T_B^{(k)}(t_0)\right]$ with $T_B^{(j)}(x) = T_B \circ T_B^{(j-1)}(x)$.

Now, we describe the details of each step. To this end, we consider the updates due to an event of the form $ev = (t, ADD(u, v, e_{uv}))$. Let us also consider $T_B^{-1}(t)$ to be the latest batch time until t, i.e., $T_B^{-1}(t) = \max_k T_B^{(k)}(t_0)$ such that $T_B^{(k)} \leq t$.

Step 1: Message Generation

Given a new interaction between u and v by $ev = (t, ADD(u, v, e_{uv}))$, we generate signed messages from v to u, which contain v's information passed to u, according to balance theory. The signed messages from u to v are also generated similarly. Thus, we here explain how to generate the signed messages from v to u.

Specifically, if e_{uv} is positive (resp. negative), then the message from positive (resp. negative) memories of node v will be used by node u to update u's positive (resp. negative) memory. More formally, positive and negative messages from v to u at time t, $\mathbf{m}_{u \leftarrow v}^+(t)$ and $\mathbf{m}_{u \leftarrow v}^-(t)$, are generated as follows²:

$$\mathbf{m}_{u \leftarrow v}^{+}(t) = \begin{cases} \operatorname{msg}(\mathbf{s}_{u}^{+}(t - \Delta t_{u}) \oplus \mathbf{s}_{v}^{+}(t - \Delta t_{v}) \oplus \Delta t_{u} \oplus |\mathbf{e}_{uv}(t)|), \mathbf{e}_{uv}(t) > 0, \\ \operatorname{msg}(\mathbf{s}_{u}^{+}(t - \Delta t_{u}) \oplus \mathbf{s}_{v}^{-}(t - \Delta t_{v}) \oplus \Delta t_{u} \oplus |\mathbf{e}_{uv}(t)|), \mathbf{e}_{uv}(t) < 0, \end{cases}$$
(1)

$$\begin{split} \mathbf{m}_{u\leftarrow v}^{-}(t) &= \\ \begin{cases} \mathsf{msg}(\mathbf{s}_{u}^{-}(t-\Delta t_{u})\oplus\mathbf{s}_{v}^{-}(t-\Delta t_{v})\oplus\Delta t_{u}\oplus|\mathbf{e}_{uv}(t)|), \mathbf{e}_{uv}(t)>0, \\ \mathsf{msg}(\mathbf{s}_{u}^{-}(t-\Delta t_{u})\oplus\mathbf{s}_{v}^{+}(t-\Delta t_{v})\oplus\Delta t_{u}\oplus|\mathbf{e}_{uv}(t)|), \mathbf{e}_{uv}(t)<0, \end{cases} \end{split}$$

where \oplus indicates the concatenation operator and \mathbf{s}_u^+ (resp. \mathbf{s}_u^-) denotes positive (resp. negative) memory of *u*. Also, $t - \Delta t_u$ denotes the time when the signed memories of node *u* were last updated (i.e., the last time when there was a direct interaction involving the node *u*). In addition, $|\mathbf{e}_{uv}(t)|$ indicates the link signed weight between *u* and *v*. Lastly, we employ a standard multi-layer perception (MLP) to implement the message function msg.

Intuitively, $\mathbf{m}_{u \leftarrow v}^{\dagger}(t)$ and $\mathbf{m}_{u \leftarrow v}^{-}(t)$ preserve the information received through recent positive and negative memories of both u

and v while taking into consideration of the balance theory. Furthermore, it considers the strength of the interaction between u and v. By doing so, SEMBA can (1) differentiate the impact of recent versus older interactions, (2) treat signed memories differently based on the link sign, and (3) differentiate the influence per link. That is, signed messages are designed so that SEMBA can jointly address the temporal-awareness (C1) and sign-awareness (C3) problems.

Step 2: Message Aggregation

Given a set of positive and negative messages, i.e., $\mathbf{m}_{u\leftarrow(\cdot)}^+(t)$ and $\mathbf{m}_{u\leftarrow(\cdot)}^-(t)$, received by each node u in the batch, we aggregate them separately as follows:

$$\mathbf{M}_{u}^{+}(t) = \operatorname{agg}(\{\mathbf{m}_{u \leftarrow i}^{+}: (\tau, \operatorname{Add}(i, u, \mathbf{e}_{iu})) \in \mathbb{E}[T_{B}^{-1}(t), t]\}), \\
\mathbf{M}_{u}^{-}(t) = \operatorname{agg}(\{\mathbf{m}_{u \leftarrow i}^{-}: (\tau, \operatorname{Add}(i, u, \mathbf{e}_{iu})) \in \mathbb{E}[T_{B}^{-1}(t), t]\}).$$
(3)

To implement the aggregation function agg, we can employ any methods used in [37], such as mean aggregation, temporal aggregation using long short-term memories (LSTMs), and a function that returns the most recent message from its input set of messages. We found no significant difference among these and used the function returning the most recent message.

Step 3: Memory Update

Now, we have signed messages, i.e., $\mathbf{M}_{u}^{+}(t)$ and $\mathbf{M}_{u}^{-}(t)$, which encode new knowledge from the events involving each node u in the batch from $\mathbb{E}[T_{B}^{-1}(t), t]$. Given $\mathbf{M}_{u}^{+}(t)$ and $\mathbf{M}_{u}^{-}(t)$, we update u's positive and negative memories, $\mathbf{s}_{u}^{+}(t)$ and $\mathbf{s}_{u}^{-}(t)$, at time t and then form a joint memory $\mathbf{s}_{u}(t)$ for the node u as follows:

$$s_{u}^{+}(t) = \operatorname{mem}(\mathbf{M}_{u}^{+}(t), \mathbf{s}_{u}^{+}(t - \Delta t)),$$

$$s_{u}^{-}(t) = \operatorname{mem}(\mathbf{M}_{u}^{-}(t), \mathbf{s}_{u}^{-}(t - \Delta t)),$$

$$s_{u}(t) = \mathbf{s}^{+}(t) \oplus \mathbf{s}^{-}(t),$$

(4)

where mem indicates a memory update function, which we employ as an LSTM in this paper.

Step 4: Embedding Generation

Given a joint memory $\mathbf{s}_u(t)$ of u and memories $\mathbf{s}_i(t)$ of all of its neighbors i until that time t, i.e., $i \in \mathcal{N}_u([0, t])$, we find each node

 $^{{}^{2}\}mathbf{m}_{v\leftarrow u}^{+}(t)$ and $\mathbf{m}_{v\leftarrow u}^{-}(t)$ are also formed similarly.

u's embedding $\mathbf{z}_u(t)$ at time *t* using Graph Transformer [39] as:

$$\mathbf{z}_{u}(t) = \mathbf{W}_{1}\mathbf{s}_{u}(t) + \sum_{\substack{(\tau, \mathbf{e}v) \in \mathbb{E}[0, t]\\ \mathbf{e}v = ADD(i, u, \mathbf{e}_{iu})}} \alpha_{i,u}(\mathbf{W}_{2}\mathbf{s}_{i}(t) + \mathbf{W}_{4}\tau + \mathbf{W}_{4}|\mathbf{e}_{iu}|), \quad (5)$$

where the attention coefficients $\alpha_{i,u}$ are found using multi-head dot product attention, and \mathbf{W}_j denote learnable weights. If there are node features $\mathbf{x}_u(t)$, we concatenate them with the node memories $\mathbf{s}_u(t)$ and pass the concatenated value in place of $\mathbf{s}_u(t)$ in Eq. (5) to obtain the embeddings $\mathbf{z}_u(t)$.

Therefore, one can note that u's embedding z_u will be updated at time t even if there are no direct interactions with u in the batch. This is due to the fact that we attentively aggregate the memories of all of its neighbors formed until now and not just in this batch. Thus, if any of these nodes have direct interactions in this batch, their memories would be updated, subsequently updating u's embedding. In this way, SEMBA can generate robust node embeddings even in the lack of direct interactions during a time period, thereby successfully mitigating the staleness problem (C2).

4.3 Learning and Inference

To make effective predictions about events, the node embeddings are first updated using events until the previous temporal batch. The predictions are made using these latest embeddings, which ensures that the embeddings account for the immediate history. The node embeddings are used to make inference on a node pair (u, v) at a time *t* by concatenating the corresponding embeddings $\mathbf{z}_u(t)$ and $\mathbf{z}_v(t)$ and then passing it through an MLP, i.e.,

$$\widehat{y}(u, v, t; \Theta) = \mathrm{MLP}(\mathbf{z}_u(t) \oplus \mathbf{z}_v(t)), \tag{6}$$

where Θ denotes the set of learnable model parameters.

We treat Dynamic Link Existence and Dynamic Link Sign Prediction (Problems 1, 2) as binary classification tasks. For Problem 1, negative sampling was employed to find an equal number of nonexisting edges. To train the models for these tasks, we minimize the Binary Cross Entropy (BCE) loss in sequential batches as:

$$\mathcal{L}_{BCE}(\mathcal{E}_{tr}; \Theta) = -\sum_{(u,v,t)\in\mathcal{E}_{tr}} \log p_y(u, v, T_B^{-1}(t); \Theta)), \tag{7}$$

where \mathcal{E}_{tr} denotes the training set and $p_y := \sigma(\widehat{y})$ if $y_{uv}(t) = 1$ and $1 - \sigma(\widehat{y})$ otherwise. Note that $T_B^{-1}(t)$ is the latest batch time until *t* and $y_{uv}(t)$ is the actual label of the edge, which is 1 for link existence in Problem 1 and 1 for positive link in Problem 2.

Dynamic Signed Link Prediction (Problem 3) is treated as a multiclass classification task. Here, we minimize the Cross-Entropy Loss similar to Eq. (7) with the true labels are in {0, 1, 2} denoting positive, negative, and non-existing edges, respectively. Corresponding p_y is found after applying softmax of \hat{y} .

4.4 Discussion

Here, we summarize the unique architectural contribution of SEMBA as compared to closely related TGN [37]. TGN is a temporal graph learning module specifically designed for homophilic unsigned networks. On the other hand, the use of two different memories and their update rules inspired by balance theory, SEMBA makes the learning appropriate to heterophilic signed networks. As we shall

Table 3: Summary of the datasets. $f_+ = |\mathcal{E}_f^+|/|\mathcal{E}_f|$ denotes the fraction of positive links and $f_{ub} = \triangle_{ub}(\mathcal{G}_f)/\triangle(\mathcal{G}_f)$ denotes the fraction of unbalanced triangles in the signed network. Weight denotes the presence of link weight in a dataset.

Dataset	#nodes	#links	<i>f</i> +	f_{ub}	#days	Weight
ВТС-Отс	5.9K	21.5K	0.89	0.13	1,905	1
BTC-Alpha	3.8K	14.1K	0.85	0.14	1,902	1
WikiRfA	12.8K	170.5K	0.77	0.24	3,758	
Epinions	131.9K	871.4K	0.85	0.09	945	

see in the experiments, such simple modification can be very effective in learning representations of links in a continuous-time dynamic signed network, by jointly addressing the temporal awareness, staleness, and sign-awareness problems.

5 EXPERIMENTS

We designed our experiments, aiming at answering the following key evaluation questions (EQs):

- EQ1: How does SEMBA perform compared to its competitors on the sign and link prediction tasks (Problems 1, 2, and 3)?
- EQ2: How does the performance vary in transductive (nodes seen during training) and inductive (unseen nodes) settings?
- EQ3: How important is each component of SEMBA?
- EQ4: Is SEMBA time and memory efficient?

5.1 Experimental Setup

5.1.1 Datasets. We used four publicly available³ real-world signed networks with temporal links: BTC-OTC, BTC-ALPHA, WIKIRFA, and EPINIONS. Table 3 shows some statistics of these datasets.

- **BTC-OTC** and **BTC-ALPHA** [23, 24] are two who-trusts-whom networks of people trading Bitcoin on OTC and Alpha platforms, respectively. In addition to the sign and timestamp, these links are also weighted on a scale from -10 to +10.
- WIKIRFA [46] is a voting network by fellow members for a Wikipedia member's request of adminship.
- **EPINIONS** is a trust network between members of an online consumer review site.

From Table 3, we note that all datasets are skewed towards positive links (see f_+) and follow balance theory (see f_{ub}). Note that WIKIRFA originally has links with null signs that denote abstention, which we filtered out to make it into a binary signed network. Furthermore, we also removed any links with no given timestamps or no given weights from all the datasets. Since no datasets contain node features, we use a zeroed vector of dimension 8 as the features.

5.1.2 Baselines. We study a novel problem on dynamic signed networks that have been studied in only one prior work, i.e., RW-LSTM [8]; we made its variant **SGC-LSTM**, which can be used in the inductive setting for a fair comparison. In particular, we employ an inductive SGCN module in place of the random walks to encode graphs in each batch that are then passed through an LSTM sequence to obtain the representations. We also use seven

³https://snap.stanford.edu/data/index.html#signnets

other state-of-the-art GNNs as additional baselines: two static unsigned GNNs, i.e., **GCN** [21], **GAT** [43]; two static signed GNNs, i.e., **SGCN** [9] and **SIGAT**⁴ [13]; and three dynamic unsigned GNNs, i.e., **TGAT**⁵ [48], **NAT**⁶ [33], and **TGN** [37].

We used two encoder layers with ReLU activation for each model, following torchgeometric's implementation unless otherwise mentioned. Since the datasets are not attributed, we used a zero vector of dimension 8 as the initial node features. One-hot and random features gave a similar performance. The embedding (and hidden) dimension was fixed to be 64 for all the models along with a dropout rate of 0.5. Furthermore, for TGAT and NAT, we used a 2-hop neighbor finder with 16 neighbors at each hop and an 8 dimensional dictionary vector and a 100-dimensional self-representation for NAT. For TGN, we use an identity message module, a mean message aggregator, a single-layer MLP time encoder, and a graph transformer (with 8 attention heads) for the embedding module. The dimensions of the memory representation and time encoding here are fixed to be 64 and 32 respectively.

To train each model for a given dataset, we did a grid search to find the best number of epochs $\in [5, 10, 15, 20]$ and initial learning rate $\in [0.01, 0.001, 0.0001]$.

5.1.3 Evaluation Metrics. We use appropriate metrics to evaluate the performance on the 3 different tasks, specified in Section 3.

For the binary classification tasks (*Dynamic Link Sign Prediction* and *Dynamic Link Existence Prediction*), we use the F1 and AUROC scores between the true and predicted labels/probabilities. F1 measures the harmonic mean of the precision and recall of the positive class, while AUROC computes the Area Under the Receiver Operating Characteristic (Precision vs Recall) curve drawn at different thresholds. While F1 handles the imbalance of classes in the data, AUROC ignores the class imbalance. Higher F1 and AUROC values are desired from a good classifier.

For multiclass classification, i.e., the *Dynamic Link Signed Existence Prediction*, we aggregate the F1 metric over different classes as an unweighted mean or macro averaging (F1mac), an average weighted by the support (F1wt), and a global calculation of precision/recall (F1mic or equivalently, accuracy Acc). F1wt accounts for imbalance in classes by weighing based on the number of samples while F1mac gives equal weight to all the classes. On the other hand, Acc may perform well by just predicting the majority class.

5.1.4 Implementation Details. To evaluate prediction quality in the future, we order all the links according to their timestamps and then spilt them into training (70%), validation (15%), and test (15%) sets. All models are trained using a common framework described in Section 4.3 with the Adam optimizer [20]. The predictions are made in a temporal batch, i.e., $t_k = T_B(t_j)$, where the batch size *B* is fixed to be 1,000 interactions for BITCOIN-OTC, BITCOIN-ALPHA, and WIKIRFA, and 16,000 interactions for the larger EPINIONS dataset. Note that the evaluation is also done in sequential batches, i.e., to predict a link at time *t*, all events (including those in the test set) before $T_B^{-1}(t)$ are included in the model's context for prediction. In other words, while the parameters remain fixed during the test

Table 4: Performance on dynamic link existence prediction (rounded to 2 decimal digits). Bold and dark-gray values are the best performing and light-gray values are the second best. 'OOM' denotes out of GPU memory.

-	n	-							
Method	BIC	-Отс	BIC-A	ALPHA	WIK	IRFA	EPINIONS		
	F1 ↑	AUROC ↑	F1 ↑	AUROC ↑	F1 ↑	AUROC ↑	F1 ↑	AUROC ↑	
GCN	0.00 ± 0.00	0.50 ± 0.01	0.00 ± 0.00	0.51 ± 0.01	0.00 ± 0.00	0.50 ± 0.01	0.00 ± 0.00	0.49 ± 0.00	
GAT	0.66 ± 0.00	0.51 ± 0.00	0.66 ± 0.00	0.50 ± 0.00	0.66 ± 0.00	0.50 ± 0.01	0.66 ± 0.00	0.50 ± 0.00	
SGCN	0.75 ± 0.00	0.82 ± 0.01	0.80 ± 0.00	0.83 ± 0.01	0.64 ± 0.00	0.75 ± 0.00	0.81 ± 0.00	0.89 ± 0.00	
SIGAT	0.71 ± 0.05	0.87 ± 0.00	0.85 ± 0.00	0.93 ± 0.00	0.86 ± 0.00	0.93 ± 0.00	-	-	
TGAT	0.72 ± 0.00	0.92 ± 0.00	0.69 ± 0.00	0.91 ± 0.01	0.66 ± 0.00	0.55 ± 0.01	0.71 ± 0.01	0.88 ± 0.00	
NAT	0.76 ± 0.01	0.91 ± 0.00	0.92 ± 0.00	0.95 ± 0.01	0.92 ± 0.01	0.98 ± 0.01	0.91 ± 0.01	0.96 ± 0.00	
TGN	0.97 ± 0.00	0.99 ± 0.00	0.98 ± 0.00	1.00 ± 0.00	0.99 ± 0.00	1.00 ± 0.00	0.92 ± 0.01	0.98 ± 0.00	
SGC-LSTM	0.25 ± 0.00	0.57 ± 0.00	0.31 ± 0.00	0.59 ± 0.00	0.50 ± 0.00	0.33 ± 0.00	-		
SEMBA	0.97 ± 0.00	0.99 ± 0.00	0.98 ± 0.00	1.00 ± 0.00	0.99 ± 0.00	1.00 ± 0.00	0.93 ± 0.01	0.98 ± 0.01	

phase, embeddings are updated in an online manner to incorporate past events to predict the future. For static models, i.e., GCN, SGCN, GAT, and SiGAT, we thus pass the graph formed by all the interactions until the current batch, i.e., $\mathcal{G}(t)$ for prediction. For all unsigned models, signs of the links are used as edge attributes if the model can handle heterogeneous interactions (i.e., for NAT, TGN, TGAT) otherwise the signs are ignored (i.e., for GCN, GAT).

We carefully tune the hyperparameters such as the learning rate and the number of epochs of baselines for each method and dataset via grid search on the validation set. The embedding dimension is fixed to be 64 and the transformer has 8 attention heads.

Reproducibility: We conduct the experiments on NVIDIA Tesla K80 GPUs with 12GB memory, and the models were implemented using the PyTorch Geometric framework.⁷

5.2 Prediction Performance

To test the effectiveness of representations learned by SEMBA, we evaluate its performance against representative baselines on the three tasks outlined earlier in Section 3. We ran for 5 different random seeds and present the mean and standard deviation.

5.2.1 Dynamic Link Existence Prediction. Table 4 shows the performance of different methods on the task of predicting whether a link exists between two nodes at a time step t given the graph until the previous batch $T_B^{-1}(t)$. This is a standard task in dynamic graphs literature [37, 48]. We find that SEMBA closely matches the performance of state-of-the-art TGN in this task for signed networks with a maximum difference of 0.01 in the F1 and AUROC metrics. This shows that SEMBA can effectively handle the evolution of link format in dynamic signed networks. NAT is the second closest baseline followed by SIGAT. Relatively low performance of NAT and TGAT compared to TGN can be attributed to their ineffectiveness in handling heterogeneous (signed) edges. Thus, we conclude that dynamic GNNs (such as TGN, NAT and SEMBA) are essential to model link existence in dynamic signed networks as they show an average of 14.2% and a minimum of 6.4% improvement over their static counterparts. We also find that SGC-LSTM is highly ineffective in modeling link formation dynamics and only shows a slight improvement over static unsigned GCN model, which produces random results. We know from the literature that time-and-graph models are theoretically less expressive than time-then-graph models [12]. Our results complement this theory by showing that even

⁴https://github.com/huangjunjie-cs/SiGAT

 $[\]label{eq:stars} $$^{thtps://github.com/StatsDLMathsRecomSys/Inductive-representation-learning-on-temporal-graphs} $$$

⁶https://github.com/Graph-COM/Neighborhood-Aware-Temporal-Network

⁷https://github.com/pyg-team/pytorch_geometric

Mathad	ВТС-Отс			BTC-Alpha			WIKIRFA			Epinions		
Method	F1wt↑	F1mac↑	Acc↑	F1wt↑	F1mac↑	Acc↑	F1wt↑	F1mac↑	Acc↑	F1wt↑	F1mac↑	Acc↑
GCN	0.37	0.30	0.34	0.37	0.29	0.33	0.39	0.31	0.37	0.35	0.30	0.33
GAT	0.23	0.21	0.39	0.23	0.21	0.32	0.22	0.21	0.31	0.21	0.22	0.31
SGCN	0.65	0.53	0.63	0.70	0.57	0.69	0.57	0.48	0.59	0.71	0.55	0.74
SIGAT	0.67	0.55	0.62	0.70	0.57	0.65	0.65	0.55	0.61	OOM	OOM	OOM
TGAT	0.11	0.14	0.16	0.13	0.15	0.16	0.22	0.23	0.34	0.17	0.18	0.22
NAT	0.39	0.34	0.35	0.55	0.44	0.50	0.69	0.56	0.66	0.63	0.49	0.65
TGN	0.88	0.76	0.87	0.86	0.72	0.84	0.85	0.71	0.83	0.79	0.68	0.76
SGC-LSTM	0.40	0.32	0.51	0.41	0.35	0.52	0.33	0.22	0.49	OOM	OOM	OOM
SEMBA	0.88	0.76	0.87	0.86	0.73	0.85	0.91	0.80	0.90	0.82	0.74	0.81

Table 5: Mean Performance on the 3-class dynamic link signed existence prediction (rounded to 2 digits). Higher F1wt, F1mac, and Acc are desired. Bold and dark-gray values are the best performing and light-gray values are the second best. 'OOM' denotes out of GPU memory. Standard deviations are found to be less than 0.05 on all metrics.

Table 6: Performance on dynamic link sign prediction (rounded to 2 decimal digits). Bold and dark-gray values are the best performing and light-gray values are the second best. 'OOM' denotes out of GPU memory.

M. d 1	BTC	-Отс	BTC-A	Alpha	WIK	IRFA	Epinions		
Method	F1↑	AUROC ↑	F1 ↑	AUROC ↑	F1 ↑	AUROC ↑	F1 ↑	AUROC ↑	
GCN	0.01 ± 0.01	0.51 ± 0.01	0.04 ± 0.02	0.50 ± 0.04	0.00 ± 0.00	0.51 ± 0.01	0.00 ± 0.00	0.49 ± 0.01	
GAT	0.65 ± 0.01	0.53 ± 0.01	0.63 ± 0.00	0.45 ± 0.01	0.62 ± 0.01	0.50 ± 0.00	0.60 ± 0.00	0.49 ± 0.00	
SGCN	0.68 ± 0.00	0.65 ± 0.01	0.73 ± 0.00	0.62 ± 0.02	0.42 ± 0.00	0.65 ± 0.00	0.22 ± 0.00	0.65 ± 0.01	
SIGAT	0.40 ± 0.00	0.69 ± 0.02	0.50 ± 0.00	0.71 ± 0.00	0.55 ± 0.00	0.68 ± 0.00	OOM	OOM	
TGAT	0.45 ± 0.01	0.57 ± 0.00	0.47 ± 0.00	0.55 ± 0.01	0.61 ± 0.00	0.51 ± 0.01	0.28 ± 0.00	0.52 ± 0.00	
NAT	0.28 ± 0.00	0.65 ± 0.00	0.26 ± 0.00	0.63 ± 0.01	0.00 ± 0.00	0.51 ± 0.01	0.59 ± 0.00	0.60 ± 0.01	
TGN	0.74 ± 0.02	0.82 ± 0.03	0.69 ± 0.02	0.70 ± 0.01	0.62 ± 0.19	0.64 ± 0.02	0.44 ± 0.12	0.59 ± 0.10	
SGC-LSTM	0.12 ± 0.00	0.54 ± 0.00	0.12 ± 0.00	0.56 ± 0.00	0.00 ± 0.00	0.05 ± 0.00	OOM	OOM	
SEMBA	0.81 ± 0.04	0.79 ± 0.02	0.75 ± 0.01	0.70 ± 0.02	0.80 ± 0.04	0.85 ± 0.01	0.68 ± 0.15	0.80 ± 0.01	

a simple *time-then-graph* model that aggregates all the temporal interactions in a single graph and pass the aggregated graph to a static signed GNN (as is done for SGCN and SIGAT baselines) is more effective than a *time-and-graph* model such as SGC-LSTM. More sophisticated *time-then-graph* models like TGN and SEMBA are found to be even more effective.

5.2.2 Dynamic Link Sign Prediction. Table 6 shows the performance of different methods on the task of predicting whether the sign of a future link at t is positive or negative given the graph until $T_B^{-1}(t)$. We find that SEMBA outperforms the baselines across all the datasets with an improvement of up to 80% in the F1 measure and up to 33% in the AUROC measure. Note that due to the scarcity of negative signs in the datasets, there is a class imbalance (see Table 3). Thus, F1 is a better measure to assess the quality of the classifier. TGN is the second best performing method, closely followed by the static signed SIGAT and SGCN models. This shows that sign classification requires effective modeling both temporal (recall C1) and sign (C3) features, which is achieved jointly by SEMBA. As shown by the suboptimality of SIGAT and TGN models through our results, neither balance theory nor temporal modeling alone can predict the signs of links in the future. Here too, we see poor performance of SGC-LSTM which can be attributed, as earlier, to its time-and-graph modeling.

5.2.3 Dynamic Link Signed Existence Prediction. Here, we combine the previous two tasks to do a 3-class classification, i.e., to predict whether a link exists between two nodes in the future and

if it exists, predict its sign. Table 5 compares the performance of different methods for this task. SEMBA outperforms or matches the baselines across all the datasets with an improvement of up to 7% in the weighted F1 (F1wt), 12.6% in the macro F1 (F1mac), and 8.4% in the micro F1 (or accuracy Acc) scores. TGN is the closest baseline and matches SEMBA's performance on BTC-OTC and BTC-ALPHA datasets. However, SEMBA shows significant improvements over TGN for the larger WIKIRFA and EPINIONS datasets. Note that F1mac takes a mean by giving equal weights to all three classes while F1wt weighs each class by the number of samples it has. More improvement in F1mac than F1wt would then mean more improvement in the minority (i.e., negative signed links) class. Thus, we can note that SEMBA specifically improves over TGN by improving on classifying node pairs with a negative link, owing to the implicit prior of balance theory in its architecture. Here too, we note the poor performance of SGC-LSTM for similar reasons as mentioned above. Furthermore, we find that static signed methods (SGCN and SIGAT) are highly suboptimal for this task as the temporal methods (TGN and SEMBA) outperform them by a margin of at least 15.5% in F1wt, at least 28.1% in F1mac, and 9.5% in Acc measures. This shows that modeling temporal evolution (recall C1) is essential for predicting signed structure of graphs in future.

5.3 Transductive and Inductive Settings

In this section, we compare the performance of the methods in transductive and inductive settings. A transductive setting is one where all the links in the test set have nodes that were seen earlier during the training phase. On the other hand, in an inductive setting, all the links in the test set have nodes that were not seen earlier during training. We sample the original test set to form a transductive and an inductive test set for each dataset and test our previously trained models on each of these sets. Table 7 shows the performance on the dynamic link signed existence prediction on these two settings. We find that SEMBA outperforms the baselines in all but two cases with an improvement of up to 16% in F1wt and 22% in F1mac for the transductive setting and 17% and 21%, respectively, for the inductive setting. TGN is the closest baseline on average but interestingly, we find that it is often outperformed by static signed methods (SGCN and SIGAT) specifically in the inductive setting. In particular, we note that SGCN beats SEMBA in the inductive setting of Epinions by a margin of 15% and 17% in F1wt and F1mac values

Table 7: Mean Performance (rounded to 2 decimal digits) on the 3-class dynamic link signed existence prediction for transductive (no new nodes) and inductive (both new nodes) settings. Bold and dark-gray values are the best performing and light-gray values are the second best. 'OOM' denotes out of GPU memory. Standard deviations are found to be less than 0.05 on all metrics.

	ВТС-Отс			BTC-Alpha			WIKIRFA				Epinions					
Method	Trans	ductive	Indu	uctive	Trans	ductive	Indu	uctive	Trans	ductive	Indu	ictive	Trans	ductive	Indu	uctive
	F1wt↑	F1mac↑	F1wt↑	F1mac↑	F1wt↑	F1mac↑	F1wt↑	F1mac↑	F1wt↑	F1mac↑	F1wt↑	F1mac↑	F1wt↑	F1mac↑	F1wt↑	F1mac↑
GCN	0.38	0.31	0.36	0.28	0.34	0.28	0.37	0.29	0.33	0.29	0.42	0.34	0.35	0.30	0.36	0.29
GAT	0.22	0.19	0.22	0.19	0.23	0.26	0.25	0.22	0.18	0.24	0.19	0.19	0.21	0.22	0.23	0.22
SGCN	0.69	0.59	0.77	0.65	0.75	0.69	0.79	0.67	0.60	0.50	0.67	0.57	0.75	0.57	0.84	0.76
SIGAT	0.68	0.58	0.76	0.63	0.70	0.65	0.74	0.61	0.60	0.52	0.71	0.62	OOM	OOM	OOM	OOM
TGAT	0.06	0.10	0.16	0.16	0.11	0.15	0.10	0.16	0.19	0.26	0.12	0.16	0.16	0.17	0.21	0.23
NAT	0.56	0.43	0.51	0.39	0.54	0.47	0.56	0.46	0.65	0.57	0.74	0.57	0.71	0.54	0.58	0.44
TGN	0.82	0.68	0.84	0.68	0.75	0.64	0.83	0.67	0.63	0.53	0.56	0.46	0.82	0.66	0.59	0.50
SGC-LSTM	0.46	0.39	0.32	0.21	0.50	0.44	0.33	0.26	0.33	0.22	0.33	0.22	OOM	OOM	OOM	OOM
SEMBA	0.85	0.73	0.85	0.72	0.83	0.75	0.85	0.70	0.75	0.68	0.83	0.75	0.84	0.76	0.73	0.65

Table 8: Ablation study of SEMBA on the multiclass dynamic link signed existence prediction. Bold and dark-gray values are the best and light-gray values are the second best.

Method	BTC	с-отс	BTC-	Alpha	Wir	IRFA	Epinions		
	F1wt↑	F1mac↑	F1wt↑	F1mac↑	F1wt↑	F1mac↑	F1wt↑	F1mac↑	
SEMBA	0.88	0.76	0.86	0.73	0.91	0.80	0.82	0.74	
- BA	0.88	0.76	0.86	0.72	0.85	0.71	0.79	0.68	
- emb	0.71	0.57	0.72	0.59	0.79	0.67	0.78	0.71	
- mem	0.67	0.54	0.71	0.57	0.63	0.46	0.70	0.54	

respectively. One reason for this observation could be the sparsity of the inductive nodes in the EPINIONS dataset. Around 75% of the inductive nodes have a degree equal to 1, i.e., they interact only once. SEMBA is designed specifically to model nodes from their interactions. Node features are used only in the final propagation step and may not affect the node embeddings too much. Future research can study more powerful mechanisms for learning node embeddings from memories to handle highly inductive cases.

5.4 Ablation Study

In this section, we conduct an ablation study to examine SEMBA's sensitivity to each component. In particular, we consider three variants of SEMBA. (1) - **BA**, which removes the balanced aggregation (BA) of memories. This is equivalent to the original TGN model. (2) - **emb**, where we remove the long-term propagation using a Graph Transformer to form embeddings (i.e., Step 4). The concatenated memories are used as embeddings. (3) - **mem**, where we remove the memories and message generation steps. We directly propagate the node features through the aggregated graph until that time via a transformer to find embeddings.

Table 8 shows the ablation results for the 3-class dynamic link signed existence prediction. We find that each component is essential for the superior performance of SEMBA as it consistently outperforms its ablation counterparts across datasets. Our results further show that memories are the most important component of the architecture, followed by long-term embedding propagation, and finally balanced aggregation. In fact, removing balanced aggregation can result in a performance comparable to removing the long-term embedding propagation (see EPINIONS). Thus, balanced aggregation is an essential advance over the original TGN model for representation learning for dynamic signed networks.

5.5 Efficiency

We compare the efficiency of SEMBA against baselines in terms of training time and memory usage. This is quantified by the per epoch training time in seconds and the allocated GPU memory in bytes for the largest EPINIONS dataset. First, we note that both SIGAT and SGC-LSTM fail to run on this dataset as they ran out of GPU memory while allocating space during training. SGCN takes 11.9 s on average to train one epoch while allocating 862 MB. On the other hand, more sophisticated TGN and SEMBA take more time and space resources to train. In particular, TGN takes 151.4 s per epoch and requires 11.9 GB, while SEMBA takes 335.2 s per epoch with 12.9 GB of GPU memory. Thus, the superior performance of SEMBA comes at a cost of more training time and GPU resources.

6 CONCLUSION

Our work provides a novel framework to learn and test representations of dynamic signed networks. Our proposed model SEMBA is shown to outperform the baselines across different prediction tasks over dynamic links. We hope that our model opens new future avenues to study signed dynamic GNNs. Extensions of our work can focus on enforcing other theories of signed network configuration such as status theory [11, 30]. This can also inspire designs of other dynamic models influenced by sociological theories of organization. Another direction could focus on predicting the signed weights of the edges. Regression tasks on graphs are known to be notorious and are largely understudied for dynamic graphs. Future works can develop effective models for the difficult task of predicting the signed weights of dynamic links.

ACKNOWLEDGEMENTS

This research/material is based upon work supported in part by NSF grants CNS-2154118, IIS-2027689, ITE-2137724, ITE-2230692, CNS-2239879, Defense Advanced Research Projects Agency (DARPA) under Agreement No. HR00112290102 (subcontract No. PO70745), and funding from Microsoft, Google, and The Home Depot. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the position or policy of DARPA, DoD, SRI International, NSF and no official endorsement should be inferred. We thank the reviewers for their comments.

Kartik Sharma, Mohit Raghavendra, Yeon-Chang Lee, Anand Kumar M, and Srijan Kumar

REFERENCES

- Peter Abell and Mark Ludwig. 2009. Structural balance: a dynamic perspective. Journal of Mathematical Sociology 33, 2 (2009), 129–155.
- [2] Teresa Alsinet, Josep Argelich, Ramón Béjar, and Santi Martínez. 2021. Measuring Polarization in Online Debates. *Applied Sciences* 11, 24 (2021), 11879.
- [3] Omid Askarisichani, Jacqueline Ng Lane, Francesco Bullo, Noah E Friedkin, Ambuj K Singh, and Brian Uzzi. 2019. Structural balance emerges and explains performance in risky decision-making. *Nature communications* 10, 1 (2019), 1–10.
- [4] Robert Axelrod, Joshua J Daymude, and Stephanie Forrest. 2021. Preventing extreme polarization of political attitudes. *Proceedings of the National Academy* of Sciences 118, 50 (2021), e2102139118.
- [5] Francesco Bonchi, Edoardo Galimberti, Aristides Gionis, Bruno Ordozgoiti, and Giancarlo Ruffo. 2019. Discovering polarized communities in signed networks. In Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM). 961–970.
- [6] Dorwin Cartwright and Frank Harary. 1956. Structural balance: a generalization of Heider's theory. Psychological review 63, 5 (1956), 277.
- [7] Jinyin Chen, Xueke Wang, and Xuanheng Xu. 2022. GC-LSTM: Graph convolution embedded LSTM for dynamic network link prediction. *Applied Intelligence* (2022), 1–16.
- [8] Quang-Vinh Dang and Claudia-Lavinia Ignat. 2018. Link-sign prediction in dynamic signed directed networks. In 2018 IEEE 4th International Conference on Collaboration and Internet Computing (CIC). IEEE, 36–45.
- [9] Tyler Derr, Yao Ma, and Jiliang Tang. 2018. Signed graph convolutional networks. In Proceedings of IEEE International Conference on Data Mining (ICDM). IEEE, 929–934.
- [10] Patrick Doreian and Andrej Mrvar. 2019. Structural balance and signed international relations. Journal of Social Structure 16, 1 (2019).
- [11] M Hamit Fişek, Joseph Berger, and Robert Z Norman. 1991. Participation in heterogeneous and homogeneous groups: A theoretical integration. Amer. J. Sociology 97, 1 (1991), 114–142.
- [12] Jianfei Gao and Bruno Ribeiro. 2022. On the equivalence between temporal and static graph representations for observational predictions. In Proceedings of the 39th International Conference on Machine Learning (ICML).
- [13] Junjie Huang, Huawei Shen, Liang Hou, and Xueqi Cheng. 2019. Signed graph attention networks. In Proceedings of International Conference on Artificial Neural Networks. Springer, 566–577.
- [14] Junjie Huang, Huawei Shen, Liang Hou, and Xueqi Cheng. 2021. SDGNN: Learning Node Representation for Signed Directed Networks. In *Thirty-Fifth AAAI* Conference on Artificial Intelligence. 196–203.
- [15] Zexi Huang, Arlei Silva, and Ambuj Singh. 2021. POLE: Polarized Embedding for Signed Networks. In Proceedings of the 15th ACM International Conference on Web Search and Data Mining (WSDM).
- [16] Norman P Hummon and Patrick Doreian. 2003. Some dynamics of social balance processes: bringing Heider back into balance theory. *Social Networks* 25, 1 (2003), 17–49.
- [17] Amin Javari, Tyler Derr, Pouya Esmailian, Jiliang Tang, and Kevin Chen-Chuan Chang. 2020. Rose: Role-based signed network embedding. In Proceedings of The Web Conference 2020. 2782–2788.
- [18] Yoonsuk Kang, Woncheol Lee, Yeon-Chang Lee, Kyungsik Han, and Sang-Wook Kim. 2021. Adversarial Learning of Balanced Triangles for Accurate Community Detection on Signed Networks. In *Proceedings of IEEE International Conference* on Data Mining (ICDM). IEEE, 1150–1155.
- [19] Min-Jeong Kim, Yeon-Chang Lee, and Sang-Wook Kim. 2023. TrustSGCN: Learning Trustworthiness on Edge Signs for Effective Signed Graph Convolutional Networks. In Proceedings of International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR). ACM, 2451–2455.
- [20] Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In Proceedings of the 3rd International Conference on Learning Representations (ICLR).
- [21] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. In Proceedings of the 5th International Conference on Learning Representations (ICLR).
- [22] Srijan Kumar, William L Hamilton, Jure Leskovec, and Dan Jurafsky. 2018. Community interaction and conflict on the web. In Proceedings of The Web Conference. 933–943.
- [23] Srijan Kumar, Bryan Hooi, Disha Makhija, Mohit Kumar, Christos Faloutsos, and VS Subrahmanian. 2018. Rev2: Fraudulent user prediction in rating platforms. In Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining (WSDM). ACM, 333–341.
- [24] Srijan Kumar, Francesca Spezzano, VS Subrahmanian, and Christos Faloutsos. 2016. Edge weight prediction in weighted signed networks. In Data Mining (ICDM), 2016 IEEE 16th International Conference on. IEEE, 221–230.
- [25] Srijan Kumar, Xikun Zhang, and Jure Leskovec. 2019. Predicting dynamic embedding trajectory in temporal interaction networks. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining

(KDD). 1269-1278.

- [26] Mirko Lai, Viviana Patti, Giancarlo Ruffo, and Paolo Rosso. 2018. Stance evolution and twitter interactions in an italian political debate. In International Conference on Applications of Natural Language to Information Systems. Springer, 15–27.
- [27] Mirko Lai, Marcella Tambuscio, Viviana Patti, Giancarlo Ruffo, and Paolo Rosso. 2019. Stance polarity in political debates: A diachronic perspective of network homophily and conversations on Twitter. *Data & Knowledge Engineering* 124 (2019), 101738.
- [28] Wonchang Lee, Yeon-Chang Lee, Dongwon Lee, and Sang-Wook Kim. 2021. Look Before You Leap: Confirming Edge Signs in Random Walk with Restart for Personalized Node Ranking in Signed Networks. In Proceedings of International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR). ACM, 143–152.
- [29] Yeon-Chang Lee, Nayoun Seo, Kyungsik Han, and Sang-Wook Kim. 2020. ASiNE: Adversarial Signed Network Embedding. In Proceedings of International ACM SIGIR conference on research and development in Information Retrieval (SIGIR). ACM, 609–618.
- [30] Jure Leskovec, Daniel Huttenlocher, and Jon Kleinberg. 2010. Signed networks in social media. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI). 1361–1370.
- [31] Yu Li, Yuan Tian, Jiawei Zhang, and Yi Chang. 2020. Learning Signed Network Embedding via Graph Attention. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*. 4772–4779.
- [32] Haoxin Liu, Ziwei Zhang, Peng Cui, Yafeng Zhang, Qiang Cui, Jiashuo Liu, and Wenwu Zhu. 2021. Signed Graph Neural Network with Latent Groups. In KDD '21: The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. 1066–1075.
- [33] Yuhong Luo and Pan Li. 2022. Neighborhood-aware Scalable Temporal Network Representation Learning. In Proceedings of the Learning on Graphs Conference (LoG).
- [34] Seth A Marvel, Jon Kleinberg, Robert D Kleinberg, and Steven H Strogatz. 2011. Continuous-time model of structural balance. *Proceedings of the National Academy of Sciences* 108, 5 (2011), 1771–1776.
- [35] Ana Peleteiro, Juan C Burguillo, and Siang Yew Chong. 2014. Exploring indirect reciprocity in complex networks using coalitions and rewiring. In AAMAS, Vol. 14. 669–676.
- [36] Simonetta Primario, Dario Borrelli, Luca Iandoli, Giuseppe Zollo, and Carlo Lipizzi. 2017. Measuring polarization in Twitter enabled in online political conversation: The case of 2016 US presidential election. In Proceedings of IEEE International Conference on Information Reuse and Integration (IRI). IEEE, 607–613.
- [37] Emanuele Rossi, Ben Chamberlain, Fabrizio Frasca, Davide Eynard, Federico Monti, and Michael Bronstein. 2020. Temporal graph networks for deep learning on dynamic graphs. arXiv preprint arXiv:2006.10637 (2020).
- [38] Martin Saveski, Brandon Roy, and Deb Roy. 2021. The structure of toxic conversations on Twitter. In Proceedings of the Web Conference 2021. 1086–1097.
- [39] Yunsheng Shi, Zhengjie Huang, Shikun Feng, Hui Zhong, Wenjin Wang, and Yu Sun. 2020. Masked label prediction: Unified message passing model for semisupervised classification. In Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI).
- [40] Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 2017. Fake news detection on social media: A data mining perspective. ACM SIGKDD explorations newsletter 19, 1 (2017), 22–36.
- [41] Lin Shu, Erxin Du, Yaomin Chang, Chuan Chen, Zibin Zheng, Xingxing Xing, and Shaofeng Shen. 2021. SGCL: Contrastive Representation Learning for Signed Graphs. In CIKM '21: The 30th ACM International Conference on Information and Knowledge Management. 1671–1680.
- [42] Rakshit Trivedi, Mehrdad Farajtabar, Prasenjeet Biswal, and Hongyuan Zha. 2019. Dyrep: Learning representations over dynamic graphs. In Proceedings of the International Conference on Learning Representations (ICLR).
- [43] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. arXiv preprint arXiv:1710.10903 (2017).
- [44] Nathan Walter, John J Brooks, Camille J Saucier, and Sapna Suresh. 2021. Evaluating the impact of attempts to correct health misinformation on social media: A meta-analysis. *Health Communication* 36, 13 (2021), 1776–1784.
- [45] Yanbang Wang, Yen-Yu Chang, Yunyu Liu, Jure Leskovec, and Pan Li. 2021. Inductive Representation Learning in Temporal Networks via Causal Anonymous Walks. In Proceedings of the International Conference on Learning Representations (ICLR).
- [46] Robert West, Hristo S Paskov, Jure Leskovec, and Christopher Potts. 2014. Exploiting social network structure for person-to-person sentiment analysis. *Transactions of the Association for Computational Linguistics* 2 (2014), 297–310.
- [47] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems* 32, 1 (2020), 4–24.
- [48] Da Xu, Chuanwei Ruan, Evren Korpeoglu, Sushant Kumar, and Kannan Achan. 2020. Inductive representation learning on temporal graphs. (2020).