Time-Aware Robustness of Temporal Graph Neural Networks for Link Prediction

Marco Sälzer
School of Electrical Engineering and Computer Science, University of Kassel, Germany

Silvia Beddar-Wiesing
School of Electrical Engineering and Computer Science, University of Kassel, Germany

Abstract

We present a first notion of a time-aware robustness property for Temporal Graph Neural Networks (TGNN), a recently popular framework for computing functions over continuous- or discrete-time graphs, motivated by recent work on time-aware attacks on TGNN used for link prediction tasks. Furthermore, we discuss promising verification approaches for the presented or similar safety properties and possible next steps in this direction of research.

2012 ACM Subject Classification Computing methodologies → Neural networks; Security and privacy → Logic and verification

Keywords and phrases graph neural networks, temporal, verification

Digital Object Identifier 10.4230/LIPIcs.TIME.2023.19

Category Extended Abstract

Introduction

Graph Neural Networks (GNN) provide a framework for computing functions over graphs based on learnable parameters and have gained much attention in recent years [7]. The most popular GNN models, so-called convolutional GNN or message-passing GNN apply a neighborhood aggregation procedure to each node in a graph to compute its output. Usually, such GNNs are used for classification or prediction tasks over static graphs. However, this limits their applicability in contexts like social networks or knowledge graphs, where underlying graphs change stepwise or time-continuously. Temporal Graph Neural Networks[1] (TGNN) [5, 6] try to close this gap. The general idea of TGNN is to generalize the neighborhood aggregation procedure mentioned above to temporal graphs, usually represented as a tuple of a base graph with a series of time-stamped observed changes. In most applications involving Neural Network based models, giving reliable safety certificates is highly desirable but also a significant challenge, especially because of the blackbox nature of neural models. In this extended abstract, we address the topic of verifying TGNN, which is an unexplored area of research. We present a time-aware robustness property for TGNN used for link prediction tasks, which is motivated by recent work on similar time-aware attacks [3]. Additionally, we discuss our ongoing work regarding promising verification approaches for the introduced (or similar) safety property.

Preliminaries

Temporal Graphs: A Continuous-Time Temporal Graph (CTG) is a tuple $(G, \mathcal{O})$ where G is a graph, often called start graph, and $\mathcal{O}$ is a finite set of time-stamped observations, including events like node or edge additions or deletions. We denote by $G^{\mathcal{O}}_{\leq t}$ for some $t \in \mathbb{Q}^{\geq 0}$ the

[1] Equivalently, these models are also called Dynamic Graph Neural Networks (DGNN).

graph constructed by applying the observations from $\mathcal{O}$ with a timestamp $t' \leq t$ to G. Note that each CTG can be seen as a finite sequence of graphs by unfolding $\mathcal{O}$ in a stepwise fashion. For more details on the notions of temporal graphs used in the context of temporal graph learning, see [4]. *Link Prediction for CTG:* Given a CTG $C = (G, \mathcal{O})$, the link prediction task is to predict for a time $t \in \mathbb{Q}^{\geq 0}$ and pair of nodes u, v present in $G^{\mathcal{O}}_{\leq t}$ whether an edge (u, v) will be present in the graph at time t. We denote the output of a TGNN for the above-described link prediction task by $N(C, (u, v), t)$.

Time-aware Robustness in the Context of Link Prediction

Chen et al. [3] present a notion of pointwise (adversarial) attacks on link predicting TGNN, exploiting the time component of temporal graphs. Similarly, we present the following definition of a pointwise time-aware (adversarial) robustness certificate for TGNN.

▶ **Definition 1.** *Let N be a TGNN, $C = (G, \mathcal{O})$ a CTG and $B = (\mathcal{O}_1, \mathcal{O}_2)$ an adversarial budget of two sets O_1, O_2 of observations where $\mathcal{O}_1 \subseteq \mathcal{O}$. We say that N is* robust for nodes u, v and time t under influence of B *if $N(C, (u, v), t) = N(C', (u, v), t)$ where $C' = (G, (\mathcal{O} \setminus \mathcal{O}_1) \cup \mathcal{O}_2)$.*

While this exact robustness certificate is desirable, a computationally feasible and complete verification algorithm is unlikely, as recent results about the decidability and complexity of similar safety properties for GNN [9] indicate. Therefore, we propose two approaches: (A) one focuses on the development of non-complete verification algorithms, similar to [10] for GNN, or (B) one gives up on exact verification and relaxes Def. 1 to a probabilistic certificate, similar to [2] for GNN. The two approaches have advantages and disadvantages: (A) allows for exact verification but most likely depends on the underlying TGNN model, making model-specific verification algorithms necessary. Approach (B) can be model-agnostic but can only give probabilistic certificates.

Outlook

We introduced a first notion of robustness for TGNN in the context of link prediction, inspired by common (adversarial) attack and robustness certificates for Neural Network based models, and discussed possible verification approaches. However, this can only be seen as a first step in developing a well-founded framework for the verification of TGNN. Next to developing efficient verification algorithms, a desirable goal is to combine TGNN verification with well-founded specification languages or temporal logic. Since TGNNs work over finite sequences or traces of graphs, a similar logic to Linear Temporal Logic (LTL) on finite traces [8] could be promising. However, the traces considered here work over infinite domains, making more expressive LTL variants necessary, like in [1].

References

1 Artale et al. First-order temporal logic on finite traces: Semantic properties, decidable fragments, and applications. *CoRR*, abs/2202.00610, 2022. URL: `https://arxiv.org/abs/2202.00610`.

2 Bojchevski et al. Efficient robustness certificates for discrete data: Sparsity-aware randomized smoothing for graphs, images and more. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1003–1013. PMLR, July 2020. URL: `https://proceedings.mlr.press/v119/bojchevski20a.html`.

3 Chen et al. Time-aware gradient attack on dynamic network link prediction. *IEEE Trans. Knowl. Data Eng.*, 35(2):2091–2102, 2023. doi:10.1109/TKDE.2021.3110580.

4 Kazemi et al. Representation learning for dynamic graphs: A survey. *J. Mach. Learn. Res.*, 21:70:1–70:73, 2020. URL: http://jmlr.org/papers/v21/19-447.html.

5 Longa et al. Graph Neural Networks for Temporal Graphs: State of the Art, Open Challenges, and Opportunities. *arXiv preprint arXiv:2302.01018*, 2023.

6 Skarding et al. Foundations and Modeling of Dynamic Networks using Dynamic Graph Neural Networks: A Survey. *IEEE Access*, 9:79143–79168, 2021.

7 Wu et al. A comprehensive survey on graph neural networks. *IEEE Trans. Neural Networks Learn. Syst.*, 32(1):4–24, 2021. doi:10.1109/TNNLS.2020.2978386.

8 Giuseppe De Giacomo and Moshe Y. Vardi. Linear temporal logic and linear dynamic logic on finite traces. In *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*, pages 854–860. IJCAI/AAAI, 2013. URL: http://www.aaai.org/ocs/index.php/IJCAI/IJCAI13/paper/view/6997.

9 Marco Sälzer and Martin Lange. Fundamental limits in formal verification of message-passing neural networks. In *The Eleventh International Conference on Learning Representations*, 2023. URL: https://openreview.net/forum?id=WlbG820mRH-.

10 Daniel Zügner and Stephan Günnemann. Certifiable robustness and robust training for graph convolutional networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4-8, 2019*, pages 246–256. ACM, 2019. doi:10.1145/3292500.3330905.