Using Temporal Graph Isomorphism to Understand the Expressivity of Temporal Graph Neural Networks

temporal graphs, temporal graph isomorphism, machine learning, graph neural networks

Extended Abstract

Graph neural networks (GNNs) are a cornerstone of deep learning in complex networks. They have recently been generalized to temporal GNNs (TGNNs) that capture patterns in time series data on *temporal graphs*, where edges carry timestamps. An important characteristic of temporal graphs is how the directed *arrow of time* influences their *causal topology*, i.e., which nodes can possibly influence each other causally via *time-respecting paths*:

Definition 1 (Time-respecting path)**.** *A path of length k in a temporal graph $G^\tau = (V, E^\tau)$ is an alternating sequence of nodes and timestamped edges $p = (v_0, e_1, v_1, \dots, e_k, v_k)$ with $e_i = (v_{i-1}, v_i; t_i) \in E^\tau$ for $i \in \{1, \dots, k\}$. For a maximum time difference (or waiting time) $\delta \in \mathbb{N}$, we say that p is* time-respecting *if $1 \leq t_i - t_{i-1} \leq \delta$ for $i \in \{1, \dots, k\}$. We denote the set of time-respecting paths in G^τ as $P^\tau(G^\tau)$.*

Many works in network science have shown how the temporal ordering of edges in temporal graphs influences processes like spreading or diffusion, node centralities or communities [1, 2]. Considering the arrow of time is thus an important prerequisite for causality-aware machine learning in complex networks. However, many TGNNs do not account for time-respecting paths, limiting their expressivity and performance. Limitations of *static* GNNs have recently been investigated in works studying their expressivity based on the Weisfeiler-Leman (WL) algorithm to heuristically distinguish non-isomorphic graphs [3]. These works yield insights into fundamental limitations of GNNs, e.g. regarding which non-isomorphic graphs cannot be distinguished in graph classification tasks. To leverage these results to understand limitations of *temporal* GNNs (TGNNs), we lack a generalization of isomorphism to temporal graphs that captures their *causal topology*. Addressing this gap, we propose the following temporal generalization of graph isomorphism, which preserves time-respecting paths:

Definition 2 (Time-respecting path isomorphism)**.** *Let $G_1^\tau = (V_1, E_1^\tau)$ and $G_2^\tau = (V_2, E_2^\tau)$ be two temporal graphs. We say that G_1^τ and G_2^τ are* time-respecting path isomorphic *if there is a bijective node mapping $\pi_V : V_1 \to V_2$ and a bijective timestamped edge mapping $\pi_E : E_1^\tau \to E_2^\tau$ such that the following holds for all alternating node/edge sequences $(v_0, e_1, v_1, \dots, e_{k-1}, v_k)$ with $k \in \mathbb{N}$:*

$$
\begin{aligned}
& (v_0, e_1, v_1, \dots, e_{k-1}, v_k) \in P^\tau(G_1^\tau) \\
\iff \quad & (\pi_V(v_0), \pi_E(e_1), \pi_V(v_1), \dots, \pi_E(e_{k-1}), \pi_V(v_k)) \in P^\tau(G_2^\tau).
\end{aligned}
$$

In our work, we show that this definition is equivalent to static graph isomorphism on the *augmented event graph*, a static line graph expansion of temporal graphs that (i) captures time-respecting paths, and (ii) is augmented by nodes in the original graph:

Definition 3 (Augmented event graph)**.** *Let $G^\tau = (V, E^\tau)$ be a temporal graph. The temporal event graph is given by $G^{\mathscr{E}} = (E^\tau, \mathscr{E})$ with*

$$\mathscr{E} = \{((u,v;t),(v,w;t')) \mid (u,v;t),(v,w;t') \in E^\tau, 1 \leq t' - t \leq \delta\}.$$

An augmented event graph *is a static, directed, node-labeled graph $G^{aug} = (V^{aug}, E^{aug}, \ell)$ with*

$$V^{aug} = V \cup E^\tau, \qquad E^{aug} = \mathscr{E} \cup E^{out} \cup E^{in},$$

$$\ell(v) = \begin{cases} 0 & \text{if } v \in V, \\ 1 & \text{if } v \in E^\tau, \end{cases} \qquad \begin{aligned} E^{out} &= \{(u, (u,v;t) \mid (u,v;t) \in E^\tau\}, \\ E^{in} &= \{((u,v;t), v) \mid (u,v;t) \in E^\tau\}. \end{aligned}$$

The following theorem (ref. to preprint with proof blinded for review) reduces time-respecting path isomorphism to a static graph isomorphism on two augmented event graphs.

Theorem 1. *Let G_1^τ and G_2^τ be two temporal graphs with corresponding augmented event graphs G_1^{aug} and G_2^{aug}. Then the following statements are equivalent:*
(i) G_1^τ and G_2^τ are time-respecting path isomorphic.
(ii) G_1^{aug} and G_2^{aug} are isomorphic.

We use our insights to derive a new TGNN architecture generating representations that allow to distinguish non-isomorphic temporal graphs. We further prove that the resulting TGNN has the same expressive power as the WL algorithm on the augmented event graph[1]. We experimentally evaluate our TGNN in synthetic temporal graphs constructed such that all classes share the same time-aggregated static graph, differing only in terms of time-respecting paths.

Experiment A: Starting from random graphs, we generate temporal edge sequences and **shuffle timestamps** of a fraction α of edges. Increasing α destroys more causal dependencies without changing the time-aggregated graph. Our TGNN reliably separates original graphs from their shuffled counterparts, achieving near-perfect accuracy for $\alpha > 0.2$ (Fig. 1, left).

Experiment B: We generate temporal graphs with two communities and **vary the likelihood of time-respecting paths between communities** using a parameter σ. For $\sigma < 0$ cross-community paths are suppressed, while for $\sigma > 0$ they are overrepresented. We assign graphs with $\sigma = 0$ to one class and graphs with $\sigma \neq 0$ to the other. Our TGNN detects these differences with high accuracy, while accuracy peaks as $|\sigma|$ increases (Fig. 1, middle/right).

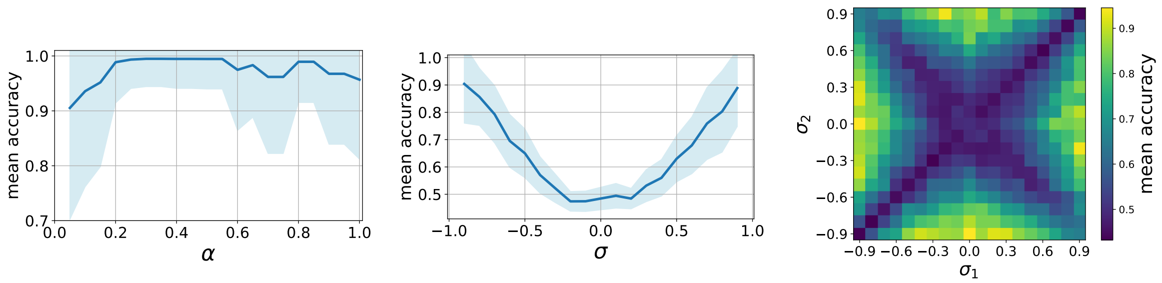

Figure 1: Results of classification experiment A (left) and B (middle), averaged over 100 runs (hull curve = std. dev.). Right panel: mean classification accuracy for temporal graphs generated with σ_1 vs. σ_2 (25 runs per pair).

Our work contributes to the theoretical foundation of temporal graph learning, providing a basis for the design and analysis of TGNN architectures that consider how the arrow of time shapes the causal topology of temporal graphs. It also shows how concepts from network science can help to improve our understanding of deep graph learning models.

References

[1] Ingo Scholtes et al. "Causality-driven slow-down and speed-up of diffusion in non-Markovian temporal networks". In: *Nature Communications* 5 (2014), p. 5024.
[2] Arash Badie-Modiri et al. "Directed percolation in temporal networks". In: *Physical Review Research* 4.2 (2022), p. L022047.
[3] Christopher Morris et al. "Weisfeiler and leman go neural: Higher-order graph neural networks". In: *Proceedings of AAAI'19*. Honolulu, Hawaii, USA, 2019.

[1] details on TGNN message passing scheme and proof in preprint, ref. blinded for review