Learning on Temporal Graphs

Abstract

Graphs are fundamental data structures for modeling real-world systems, including social, transactional, biological, and information networks. While substantial progress has been made in graph modeling via Graph Neural Networks (GNNs), their applicability remains limited in real-world systems that are inherently temporal, such as communication, financial, and transportation networks. These systems have evolving interactions and attributes over time, and collapsing such systems into static graphs discards critical temporal information, limiting the effectiveness of the trained model on downstream tasks such as future link prediction, event time prediction, and dynamic node classification. These limitations have led to the emergence of Temporal Graph Neural Networks (TGNNs), which model temporal graphs directly as sequences of time-stamped interactions.

Recent advances in TGNNs have enabled inductive and dynamic representation learning by aggregating information from past temporal neighbors. However, several challenges in temporal graph learning remain unexplored. These include temporal graph generation, dynamic features, and efficient training and transfer learning in data-scarce settings.

This thesis addresses several of these core problems and makes the following contributions. First, it studies temporal graph generative modeling and proposes a framework, TIGGER, that learns the joint distribution of node interactions and interaction times using temporal random walks. Second, it studies learning on graphs with dynamic feature sets, where features can change over time, and proposes the GraFNNE framework to model such dynamics. This thesis further explores the transfer learning of trained TGNNs in data-scarce conditions and develops MINTT, a framework that enables effective knowledge transfer across temporal graphs. Fourth, based on insights from feature dynamics, this thesis motivates a scalable graph transformer architecture for large attributed graphs that supports all-pair attention without quadratic complexity. Finally, it examines limitations in the current evaluation and training setup for TGNNs and proposes deployment-relevant evaluation metrics, along with an importance-based negative sampling strategy that leads to a more effective training process and improved predictive performance.

Overall, this thesis advances temporal graph learning by addressing multiple key challenges. The proposed systems improve the scalability and inductive generalization capabilities of temporal graph learning systems, making them applicable to real-world settings. Moreover, this thesis identifies potential frontiers that require further investigation to advance the field.