Towards Efficient Temporal Graph Learning: Algorithms, Frameworks, and Tools

Ruijie Wang
University of Illinois
Urbana-Champaign
Champaign, Illinois, USA
ruijiew2@illinois.edu

Wanyu Zhao
University of Illinois
Urbana-Champaign
Champaign, Illinois, USA
wanyu2@illinois.edu

Dachun Sun
University of Illinois
Urbana-Champaign
Champaign, Illinois, USA
dsun18@illinois.edu

Charith Mendis
University of Illinois
Urbana-Champaign
Champaign, Illinois, USA
charithm@illinois.edu

Tarek Abdelzaher
University of Illinois
Urbana-Champaign
Champaign, Illinois, USA
zaher@illinois.edu

ABSTRACT

Temporal graphs capture dynamic node relations via temporal edges, finding extensive utility in wide domains where time-varying patterns are crucial. Temporal Graph Neural Networks (TGNNs) have gained significant attention for their effectiveness in representing temporal graphs. However, TGNNs still face significant efficiency challenges in real-world low-resource settings. First, from a data-efficiency standpoint, training TGNNs requires sufficient temporal edges and data labels, which is problematic in practical scenarios with limited data collection and annotation. Second, from a resource-efficiency perspective, TGNN training and inference are computationally demanding due to complex encoding operations, especially on large-scale temporal graphs. Minimizing resource consumption while preserving effectiveness is essential. Inspired by these efficiency challenges, this tutorial systematically introduces state-of-the-art *data-efficient* and *resource-efficient* TGNNs, focusing on algorithms, frameworks, and tools, and discusses promising yet under-explored research directions in efficient temporal graph learning. This tutorial aims to benefit researchers and practitioners in data mining, machine learning, and artificial intelligence.

CCS CONCEPTS

• **Mathematics of computing → Graph algorithms**; • **Computing methodologies → Machine learning**; • **Information systems → Data mining**.

KEYWORDS

Temporal Graphs; Graph Neural Networks; Data-Efficient Learning; Resource-Efficient Learning

ACM Reference Format:
Ruijie Wang, Wanyu Zhao, Dachun Sun, Charith Mendis, and Tarek Abdelzaher. 2024. Towards Efficient Temporal Graph Learning: Algorithms, Frameworks, and Tools. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM '24), October 21–25, 2024, Boise, ID, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3627673.3679104

CIKM '24, October 21–25, 2024, Boise, ID, USA

ACM ISBN 979-8-4007-0436-9/24/10
https://doi.org/10.1145/3627673.3679104

1 INTRODUCTION

Many real-world graphs exhibit dynamic changes in their topology and node/edge features. For example, this occurs in knowledge graphs with new events, social networks with new interactions, and sensor networks with evolving signals. Recently, there has been significant interest in Temporal Graph Neural Networks (TGNNs) as advanced techniques for temporal graph learning [3, 9, 14, 15, 19, 24]. These networks are designed to simultaneously capture both the graph structure information (e.g., temporal neighbors and node features) and temporal information (e.g., edge timestamps) into time-dependent node embedding vectors, greatly enhancing their utility in various temporal prediction and classification tasks.

However, using TGNNs in real-world applications is still far from efficient as it requires costly data collection, annotation, and computation resources. **First**, training TGNNs without sufficient supervision (e.g., temporal edges and data labels) is challenging. While some TGNNs can be trained in a self-supervised manner by reconstructing observed temporal edges [3, 14, 24], certain tasks (e.g., dynamic node classification) and specific scenarios (e.g., scarce observation [13], new tasks [11, 18], or new nodes [20, 21, 25]) still impose stringent requirements for data collection and annotation, where task-related labels and temporal edges are often unavailable. **Second**, the training and inference of TGNNs are computationally intensive due to complex encoding operations, which greatly limits their practicality for large-scale temporal graphs [10]. Thus, reducing the resource consumption of TGNN training and inference while maintaining their effectiveness is crucial but challenging.

In response to the aforementioned challenges, extensive research has been carried out to enhance both *data efficiency* and *resource efficiency* of TGNNs. To be concrete, **Data-efficient** TGNNs aim to mitigate the costs and time associated with data collection and annotation. According to their reliance on temporal edges and data labels, existing works are divided into three categories: self-supervised learning [1–3, 12, 14, 16, 24], weakly-supervised learning [13, 27], and few-shot learning [11, 21, 25]. **Resource-efficient** TGNNs aim to accelerate model training and inference without compromising accuracy. This is accomplished by improving computation resource utilization (e.g., parallel temporal neighbor sampling), reducing data transfer (e.g., feature caching), and minimizing computation (e.g., eliminating redundancy). Depending on the types of temporal graphs, there are two main categories: optimizations for discrete-time graphs (a sequence of snapshots) [5, 8], and optimizations for continuous-time graphs (event streams) [6, 7, 23, 28]. Furthermore, distributed frameworks utilizing multiple GPUs are proposed to facilitate efficient learning on large-scale real-world graphs [29].

In this tutorial, we will systematically review efficient TGNN algorithms, frameworks, and tools designed to facilitate temporal graph learning in real-world low-resource settings, emphasizing both data efficiency and resource efficiency. To this end, the tutorial is organized into four parts. In Part I, we provide a background introduction of temporal graph learning, and discuss the motivations and problem definitions of TGNNs. In part II, we first discuss the challenges in designing data-efficient TGNNs and then introduce three categories of methods to reduce costs in data collection and annotation. Part III focuses on the challenges of designing resource-efficient TGNNs, and introduces the major efficiency bottleneck and representative methods for discrete-time and continuous-time TGNNs respectively. We also introduce state-of-the-art distributed frameworks for large-scale training. Finally, in Part IV, we discuss potential research opportunities and future directions for enhancing the data efficiency and resource efficiency of TGNNs.

2 TUTORIAL DETAILS

2.1 Outline of the Tutorial

The tutorial will last for 3 hours, and the outline is as follows:

- Part I: Introduction
  - Background and Motivations
  - Problem Definitions and Settings
- Part II: Data-Efficient Temporal Graph Learning
  - Key Challenges of Data-Efficient TGNNs
  - Self-Supervised Temporal Graph Learning
  - Weakly-Supervised Temporal Graph Learning
  - Few-Shot Temporal Graph Learning
- Part III: Resource-Efficient TGNNs
  - Key Challenges of Resource-Efficient TGNNs
  - Efficient Discrete-Time TGNN Frameworks
  - Efficient Continuous-Time TGNN Frameworks
  - Efficient Distributed TGNN Training Frameworks
- Part IV: Open Questions and Challenges
  - Generative Pre-training on Temporal Graphs
  - Distributed Training on Temporal Graphs

The details for Parts II-IV are introduced as follows.

2.2 Part II: Data-Efficient TGNNs

Data-Efficient TGNNs aim to mitigate the costs and time associated with data collection and annotation. In particular, three fundamental problems have been studied to achieve the goal: self-supervised learning, weakly-supervised learning, and few-shot learning. They differ in their reliance on temporal edges and task-related labels. *Self-supervised learning* only leverages temporal edges information to pretrain task-agnostic models without relying on external data labels [2, 14, 16, 24]. In this direction, we will introduce two major techniques: (1) generative/predictive approaches that train models by reconstructing temporal edges [14, 24], and (2) contrastive approaches that learns parameters by performing discrimination between positive and negative pairs [1, 2, 12, 16]. *Weakly-supervised learning* uses incomplete, indirect, or inaccurate supervision signals to develop effective TGNNs. We will review recent advances in TGNNs with various weak supervision types, including sparse temporal edges [13] and limited task-related labels [27]. *Few-shot learning* aims to improve TGNNs' generalization to new tasks with limited edges and labels. We will explore how few-shot learning techniques, such as meta-learning, can enhance TGNNs to tackle different unseen temporal graph learning tasks, like new nodes [20, 21, 25], new data [18] and new classes [11].

2.3 Part III: Resource-Efficient TGNNs

Introducing the temporal dimension makes TGNN training more computationally complex and resource-intensive compared to static GNNs. The unique data access and computation patterns of TGNN models prevent them from directly utilizing optimized static GNN frameworks like PyG and DGL [4, 17]. However, several frameworks have recently been proposed to optimize TGNN training. In this part, we discuss the challenges of training TGNNs resource-efficiently, such as reducing high I/O overhead [6, 7, 22], increasing resource utilization [5, 6], and maintaining large-batch training effectiveness [6, 28]. We categorize related works based on the types of temporal graphs they address and introduce state-of-the-art TGNN frameworks and optimizations for discrete-time dynamic graphs (DTDGs) [5, 8] and continuous-time dynamic graphs (CTDGs) [6, 7, 23, 28]. Moreover, we review distributed training frameworks for large-scale real-world graphs with hundreds of millions of vertices and billions of edges [29].

2.4 Part IV: Open Questions and Challenges

To mitigate the dependency on data annotation and computational resources in downstream applications, the pre-training plus lightweight prompting/fine-tuning paradigm has become a promising approach [26]. In this paradigm, pre-training extracts general knowledge from vast amounts of unlabeled data, while prompting or fine-tuning bridges the gap between pretext and downstream tasks. This approach introduces new challenges for both data efficiency and resource efficiency. In this open direction, we will discuss (1) how to design general pre-training objectives for temporal graphs to improve data efficiency and (2) how to develop large-scale distributed pre-training frameworks considering various temporal graph characteristics to improve resource-efficiency.

3 COVERED WORKS IN THE TUTOTIAL

For space, we only list part of representative works here. This is not an exhaustive list of papers that are relevant to the topic.

- Part II: Data-Efficient Temporal Graph Learning
  - Self-Supervised Learning of TGNNs
    - Da Xu, Chuanwei Ruan, Evren Korpeoglu, Sushant Kumar, and Kannan Achan. Inductive representation learning on temporal graphs. ICLR 2020.

    * Emanuele Rossi, Ben Chamberlain, Fabrizio Frasca, Davide Eynard, Federico Monti, and Michael Bronstein. Temporal graph networks for deep learning on dynamic graphs. ICML Workshop on Graph Representation Learning, 2020.
    * Sheng Tian, Ruofan Wu, Leilei Shi, Liang Zhu, and Tao Xiong. Self-supervised representation learning on dynamic graphs. CIKM 2021.
    * Mohammad Alomrani, Mahdi Biparva, Yingxue Zhang, and Mark Coates. Dyg2vec: Efficient representation learning for dynamic graphs. TMLR 2024.
    * Lingwen Liu, Guangqi Wen,Peng Cao,Jinzhu Yang,Weiping Li,and Osmar R. Zaiane. Capturing temporal node evolution via self-supervised learning: A new perspective on dynamic graph learning. WSDM 2024.
  - Weakly-Supervised Learning of TGNNs
    * Linhao Luo, Gholamreza Haffari, and Shirui Pan. Graph sequential neural ode process for link prediction on dynamic and sparse graphs. WSDM 2023.
    * Guolin Zhang, Zehui Hu, Guoqiu Wen, Junbo Ma, Xiaofeng Zhu. Dynamic graph convolutional networks by semi-supervised contrastive learning. Pattern Recogn 2023.
  - Few-Shot Learning of TGNNs
    * Cheng Yang, Chunchen Wang, Yuanfu Lu, Xumeng Gong, Chuan Shi, Wei Wang, and Xu Zhang. Few-shot link prediction in dynamic networks. WSDM 2022.
    * Ruijie Wang, Zheng li, Dachun Sun, Shengzhong Liu, Jinning Li, Bing Yin, and Tarek Abdelzaher. Learning to sample and aggregate: Few-shot reasoning over temporal knowledge graphs. NeurIPS 2022.
    * Tiancheng Huang, Feng Zhao, and Donglin Wang. Lgp: Few-shot class- evolutionary learning on dynamic graphs. CIKM 2022.
- Part III: Resource-Efficient Temporal Graph Learning
  - Efficient Discrete-Time Frameworks
    * Mingyu Guan, Anand Padmanabha Iyer, Taesoo Kim. DynaGraph: Dynamic Graph Neural Networks at Scale. GRADES-NDA 2022.
    * Kaihua Fu, Quan Chen, Yuzhuo Yang, Jiuchen Shi, Chao Li, Minyi Guo. BLAD: Adaptive Load Balanced Scheduling and Operator Overlap Pipeline for Accelerating the Dynamic GNN Training. SC 2023.
  - Efficient Continuous-Time Frameworks
    * Hongkuan Zhou, Da Zheng, Israt Nisa, Vasileios Ioannidis, Xiang Song, George Karypis. TGL: A General Framework for Temporal GNN Training on Billion-Scale Graphs. VLDB 2022.
    * Shihong Gao, Yiming Li, Xin Zhang, Yanyan Shen, Yingxia Shao, Lei Chen. SIMPLE: Efficient Temporal Graph Neural Network Training at Scale with Dynamic Data Placement. SIGMOD 2024.
    * Shihong Gao, Yiming Li, Yanyan Shen, Yingxia Shao, Lei Chen. ETC: Efficient Training of Temporal Graph Neural Networks over Large-Scale Dynamic Graphs. VLDB 2024.
    * Yufeng Wang and Charith Mendis. TGLite: A Lightweight Programming Framework for Continuous-Time Temporal Graph Neural Networks. ASPLOS 2024.
  - Efficient Distributed Training Frameworks
    * Hongkuan Zhou, Da Zheng, Xiang Song, George Karypis, Viktor Prasanna. DistTGL: Distributed Memory-Based Temporal Graph Neural Network Training. SC 2023.

4 RELATED TUTORIALS

- **Natural and Artificial Dynamics in GNNs: A Tutorial**
  - **Presenters**: Dongqi Fu, Zhe Xu, Hanghang Tong, Jingrui He
  - **Conference**: WSDM '23, February, 2023, Singapore
  - **Connection**: This tutorial presents GNNs designed for dynamic and imperfect graphs. The discussion on GNNs handling natural dynamics sets the foundation for ours on efficient temporal graph neural networks (TGNNs).
  - **Differences**: Our primary focus is on efficient TGNNs in real-world low-resource environments. We explore how related works are designed to improve both data- and resource-efficiency. Additionally, we discuss relevant works centered on both algorithms and frameworks/libraries.
- **Toward Graph Minimally-Supervised Learning**
  - **Presenters**: Kaize Ding, Chuxu Zhang, Jie Tang, Nitesh Chawla, and Huan Liu
  - **Conference**: KDD '22, August 14–18, 2022, DC
  - **Connection**: This tutorial introduces data-efficient learning on static graphs, including few-shot, weakly-supervised, and self-supervised learning. The covered works motivate recent studies of data-efficient learning on temporal graphs.
  - **Differences**: This tutorial emphasizes data-efficient learning on static graphs. We systematically introduce recent works on temporal graphs. We also cover resource-efficient learning, which is crucial for low-resource setting.
- **Efficient Machine Learning On Large-Scale Graphs**
  - **Presenters**: Parker Erickson, Victor E. Lee, Feng Shi, and Jiliang Tang
  - **Conference**: KDD '22, August 14–18, 2022, DC
  - **Connection**: This tutorial introduces resource-efficient GNN algorithms and frameworks on static graphs.
  - **Differences**: This tutorial presents several widely-used efficient frameworks for static graphs, such as PyG and DGL. Ours covers emerging algorithms and frameworks specifically for designing resource-efficient TGNNs.

5 PRESENTER BIOGRAPHY

Ruijie Wang is a Postdoctoral Research Associate at the Department of Computer Science, the University of Illinois at Urbana-Champaign. He received his Ph.D. in Computer Science from the University of Illinois at Urbana-Champaign. His research interests lie in deep graph learning algorithms for real-world graphs at scale to understand the underlying dynamic patterns and predict future knowledge. He is also generally interested in machine learning and deep learning on graphs, natural language, and time-series data, with applications on social network analysis, knowledge graph, and dynamic systems. He has published more than 30 papers in refereed international conferences and journals including NeurIPS, WWW, ACL, SIGIR, AAAI, CIKM, SenSys, etc.

Wanyu Zhao is a first-year Ph.D. student in Computer Science at the University of Illinois Urbana-Champaign (UIUC). Her current research focuses on developing efficient and scalable systems for temporal graph learning. With an interest in the intersection of systems and machine learning, she aims to explore novel techniques in constructing efficient AI systems through algorithmic insights and comprehensive systems understanding. She is also interested in the application of machine learning to enhance system performance.

Dachun Sun is a senior Ph.D. student in Computer Science at the University of Illinois at Urbana-Champaign (UIUC). His research focuses on computational social analysis with deep graph learning and large language models. Main topics include social network data mining and multimodal embedding for social data. Additionally, his academic interests extend to natural language processing, knowledge graphs, and diffusion-based methods on graphs. He has a dozen of published papers at renowned international conferences and journals including TPAMI, NeurIPS, WWW, AAAI, SIGIR, etc.

Charith Mendis is an Assistant Professor at the University of Illinois at Urbana-Champaign. Previously, he was a visiting faculty researcher at Google and was instrumental in designing and developing the learned TPU cost model used in production. His research interests are in automating compiler construction and in building high-performance ML systems. He received his Ph.D. and Masters from the Massachusetts Institute of Technology and his B.Sc. from the University of Moratuwa. He is the recipient of an NSF CAREER Award, an IEEE Micro Top Picks honorable mention, the William A. Martin outstanding master's thesis award at MIT, a best student paper award, a best paper award, and the university gold medal for his B.Sc. He has published work at both top programming languages venues such as PLDI and ASPLOS as well as at top machine learning venues such as ICML and NeurIPS.

Tarek Abdelzaher is a Professor and Willett Faculty Scholar at the Department of Computer Science, the University of Illinois at Urbana-Champaign. He received his Ph.D. in Computer Science from the University of Michigan in 1999. He has authored/coauthored more than 300 refereed publications in real-time computing, CPS/IoT, distributed systems, intelligent networked sensing, machine learning, and control. He served as Editor-in-Chief of the Journal of Real-Time Systems for 20 years, and as Associate Editor of the IEEE Transactions on Mobile Computing, IEEE Transactions on Parallel and Distributed Systems, IEEE Embedded Systems Letters, the ACM Transaction on Sensor Networks, ACM Transactions on Internet Technology, ACM Transactions on Internet of Things, and the Ad Hoc Networks Journal. He chaired (as Program or General Chair) several conferences in his area including RTAS, RTSS, IPSN, Sensys, ICDCS, Infocom, and ICAC. He is a recipient of the IEEE Outstanding Technical Achievement and Leadership Award in Real-time Systems (2012), the Xerox Award for Faculty Research (2011), and several best paper awards. He is a fellow of ACM and IEEE.

ACKNOWLEDGEMENTS

This work was sponsored by DARPA award HR001121C0165, DARPA award HR00112290105, DoD Basic Research Office award HQ00342110002, the Army Research Laboratory under Cooperative Agreement W911NF-17-20196, and NSF award CCF-2316233. It was also supported in part by ACE, one of the seven centers in JUMP 2.0, a Semiconductor Research Corporation (SRC) program sponsored by DARPA.

REFERENCES

[1] Mohammad Alomrani, Mahdi Biparva, Yingxue Zhang, and Mark Coates. Dyg2vec: Efficient representation learning for dynamic graphs. *TMLR '24*, 2024.
[2] Ke-Jia Chen, Linsong Liu, Linpu Jiang, and Jingqiang Chen. Self-supervised dynamic graph representation learning via temporal subgraph contrast. *ACM Trans. Knowl. Discov. Data*, 2023.
[3] Weilin Cong, Si Zhang, Jian Kang, Baichuan Yuan, Hao Wu, Xin Zhou, Hanghang Tong, and Mehrdad Mahdavi. Do we really need complicated model architectures for temporal networks? In *ICLR '23*, 2023.
[4] Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In *WICLR on Graphs and Manifolds*, 2019.
[5] Kaihua Fu, Quan Chen, Yuzhuo Yang, Jiuchen Shi, Chao Li, and Minyi Guo. Blad: Adaptive load balanced scheduling and operator overlap pipeline for accelerating the dynamic gnn training. In *SC '23*, 2023.
[6] Shihong Gao, Yiming Li, Yanyan Shen, Yingxia Shao, and Lei Chen. Etc: Efficient training of temporal graph neural networks over large-scale dynamic graphs. In *VLDB '24*, 2024.
[7] Shihong Gao, Yiming Li, Xin Zhang, Yanyan Shen, Yingxia Shao, and Lei Chen. Simple: Efficient temporal graph neural network training at scale with dynamic data placement. In *SIGMOD '24*, 2024.
[8] Mingyu Guan, Anand Padmanabha Iyer, and Taesoo Kim. Dynagraph: Dynamic graph neural networks at scale. In *GRADES-NDA '22*, 2022.
[9] Ehsan Hajiramezanali, Arman Hasanzadeh, Krishna Narayanan, Nick Duffield, Mingyuan Zhou, and Xiaoning Qian. Variational graph recurrent neural networks. In *NeurIPS '19*, 2019.
[10] Shenyang Huang, Farimah Poursafaei, Jacob Danovitch, Matthias Fey, Weihua Hu, Emanuele Rossi, Jure Leskovec, Michael Bronstein, Guillaume Rabusseau, and Reihaneh Rabbany. Temporal graph benchmark for machine learning on temporal graphs. In *NeurIPS '23*, 2023.
[11] Tiancheng Huang, Feng Zhao, and Donglin Wang. Lgp: Few-shot class-evolutionary learning on dynamic graphs. In *CIKM '22*, 2022.
[12] Lingwen Liu, Guangqi Wen, Peng Cao, Jinzhu Yang, Weiping Li, and Osmar R. Zaiane. Capturing temporal node evolution via self-supervised learning: A new perspective on dynamic graph learning. In *WSDM '24*, 2024.
[13] Linhao Luo, Gholamreza Haffari, and Shirui Pan. Graph sequential neural ode process for link prediction on dynamic and sparse graphs. In *WSDM '23*, 2023.
[14] Emanuele Rossi, Ben Chamberlain, Fabrizio Frasca, Davide Eynard, Federico Monti, and Michael Bronstein. Temporal graph networks for deep learning on dynamic graphs. In *ICML 2020 Workshop on Graph Representation Learning*, 2020.
[15] Aravind Sankar, Yanhong Wu, Liang Gou, Wei Zhang, and Hao Yang. Dysat: Deep neural representation learning on dynamic graphs via self-attention networks. In *WSDM '20*, 2020.
[16] Sheng Tian, Ruofan Wu, Leilei Shi, Liang Zhu, and Tao Xiong. Self-supervised representation learning on dynamic graphs. In *CIKM '21*, 2021.
[17] Minjie Wang, Da Zheng, Zihao Ye, Quan Gan, Mufei Li, Xiang Song, Jinjing Zhou, Chao Ma, Lingfan Yu, Yu Gai, Tianjun Xiao, Tong He, George Karypis, Jinyang Li, and Zheng Zhang. Deep graph library: A graph-centric, highly-performant package for graph neural networks. *CoRR*, 2019.
[18] Ruijie Wang, Jingyuan Huang, Yutong Zhang, Jinyang Li, Yufeng Wang, Wanyu Zhao, Shengzhong Liu, Charith Mendis, and Tarek Abdelzaher. Tgonline: Enhancing temporal graph learning with adaptive online meta-learning. In *SIGIR '24*, 2024.
[19] Ruijie Wang, Zheng Li, Jingfeng Yang, Tianyu Cao, Chao Zhang, Bing Yin, and Tarek Abdelzaher. Mutually-paced knowledge distillation for cross-lingual temporal knowledge graph reasoning. In *The Web Conference*, 2023.
[20] Ruijie Wang, Yutong Zhang, Jinyang Li, Shengzhong Liu, Dachun Sun, Tianchen Wang, Tianshi Wang, Yizhuo Chen, Denizhan Kara, and Tarek Abdelzaher. Metahkg: Meta hyperbolic learning for few-shot temporal reasoning. In *SIGIR '24*, 2024.
[21] Ruijie Wang, zheng li, Dachun Sun, Shengzhong Liu, Jinning Li, Bing Yin, and Tarek Abdelzaher. Learning to sample and aggregate: Few-shot reasoning over temporal knowledge graphs. In *NeurIPS '22*, 2022.
[22] Yufeng Wang and Charith Mendis. Tgopt: Redundancy-aware optimizations for temporal graph attention networks. In *PPoPP '23*, 2023.
[23] Yufeng Wang and Charith Mendis. Tglite: A lightweight programming framework for continuous-time temporal graph neural networks. In *ASPLOS '24*, 2024.
[24] Da Xu, Chuanwei Ruan, Evren Korpeoglu, Sushant Kumar, and Kannan Achan. Inductive representation learning on temporal graphs. In *ICLR '20*, 2020.
[25] Cheng Yang, Chunchen Wang, Yuanfu Lu, Xumeng Gong, Chuan Shi, Wei Wang, and Xu Zhang. Few-shot link prediction in dynamic networks. In *WSDM '22*.
[26] Xingtong Yu, Yuan Fang, Zemin Liu, Yuxia Wu, Zhihao Wen, Jianyuan Bo, Xinming Zhang, and Steven C. H. Hoi. Few-shot learning on graphs: from meta-learning to pre-training and prompting. *CoRR*, 2024.
[27] Guolin Zhang, Zehui Hu, Guoqiu Wen, Junbo Ma, and Xiaofeng Zhu. Dynamic graph convolutional networks by semi-supervised contrastive learning. *Pattern Recogn.*, 2023.
[28] Hongkuan Zhou, Da Zheng, Israt Nisa, Vassilis N. Ioannidis, Xiang Song, and George Karypis. TGL: A general framework for temporal GNN training on billion-scale graphs. In *VLDB '22*, 2022.
[29] Hongkuan Zhou, Da Zheng, Xiang Song, George Karypis, and Viktor Prasanna. Disttgl: Distributed memory-based temporal graph neural network training. In *SC '23*, 2023.