Graph Neural Networks: Foundation, Frontiers and Applications

Lingfei Wu*
lwu@email.wm.edu
JD.COM Silicon Valley Research Center
Mountain View, CA, USA

Peng Cui
cuip@tsinghua.edu.cn
Tsinghua University
Beijing, China

Jian Pei
jpei@cs.sfu.ca
Simon Fraser University
British Columbia, Canada

Liang Zhao
liang.zhao@emory.edu
Emory University
Atlantic, GA, USA

Xiaojie Guo
xguo7@gmu.edu
JD.COM Silicon Valley Research Center
Mountain View, CA, USA

ABSTRACT

The field of graph neural networks (GNNs) has seen rapid and incredible strides over the recent years. Graph neural networks, also known as deep learning on graphs, graph representation learning, or geometric deep learning, have become one of the fastest-growing research topics in machine learning, especially deep learning. This wave of research at the intersection of graph theory and deep learning has also influenced other fields of science, including recommendation systems, computer vision, natural language processing, inductive logic programming, program synthesis, software mining, automated planning, cybersecurity, and intelligent transportation. However, as the field rapidly grows, it has been extremely challenging to gain a global perspective of the developments of GNNs. Therefore, we feel the urgency to bridge the above gap and have a comprehensive tutorial on this fast-growing yet challenging topic.

This tutorial of Graph Neural Networks (GNNs): Foundation, Frontiers and Applications will cover a broad range of topics in graph neural networks, by reviewing and introducing the fundamental concepts and algorithms of GNNs, new research frontiers of GNNs, and broad and emerging applications with GNNs. In addition, rich tutorial materials will be included and introduced to help the audience gain a systematic understanding by using our recently published book-*Graph Neural Networks (GNN): Foundation, Frontiers and Applications* [12], which can be easily accessed at https://graph-neural-networks.github.io/index.html.

CCS CONCEPTS

• **Computing methodologies** → **Neural networks**;

KEYWORDS

Graph neural networks; Deep learning

*Corresponding tutor

KDD '22, August 14–18, 2022, Washington, DC, USA

ACM ISBN 978-1-4503-9385-0/22/08.
https://doi.org/10.1145/3534678.3542609

1 TARGET AUDIENCES AND PREREQUISITES

Target Audience:the conference attendees with interest in the deep learning on graphs techniques as well as their important real-world applications.

Prerequisites: The audience is expected to have some basic understanding of data mining, machine learning, linear algebra, and optimization. However, the tutorial will be presented at college junior/senior level and should be comfortably followed by academic researchers and practitioners from the industry.

2 TUTORIAL TOPICS

This tutorial can benefit both the academic and industry community from the following aspects:

- The basic ideas and typical variants of graph neural networks for the graph representation learning [3, 13, 14].
- Foundations of the graph neural networks including node classification, the expressive power, the interpretability and scalability of the graph neural networks[5, 9–11].
- Frontiers problems, including graph generation and transformation, graph matching and graph structural learning [1, 4, 6, 7].
- Emerging Applications of GNN in recommendation system and natural language processing [2, 8].

2.1 Tutors biography

Lingfei Wu is a Principal Scientist at JD.COM Silicon Valley Research Center, leading a team of 30+ machine learning/natural language processing scientists and software engineers to build intelligent e-commerce personalization systems. He earned his Ph.D. degree in computer science from the College of William and Mary in 2016. Previously, he was a research staff member at IBM Thomas J. Watson Research Center and led a 10+ research scientist team for developing novel Graph Neural Networks methods and systems, which leads to the #1 AI Challenge Project in IBM Research and multiple IBM Awards. He was the recipients of the Best Paper Award and Best Student Paper Award of several conferences such as IEEE ICC'19, AAAI workshop on DLGMA'20 and KDD workshop on DLG'19. His research has been featured in numerous media outlets, including NatureNews, YahooNews, Venturebeat, TechTalks, SyncedReview, Leiphone, QbitAI, MIT News, IBM Research News, and SIAM News. He has co-organized 10+ conferences (KDD, AAAI, IEEE BigData) and is the founding co-chair for Workshops of Deep

Learning on Graphs (with AAAI'21, AAAI'20, KDD'21, KDD'20, KDD'19, and IEEE BigData'19).

Peng Cui is an Associate Professor with tenure at Department of Computer Science in Tsinghua University. He obtained his PhD degree from Tsinghua University in 2010. His research interests include data mining, machine learning and multimedia analysis, with expertise on network representation learning, causal inference and stable learning, social dynamics modeling, and user behavior modeling, etc. He is keen to promote the convergence and integration of causal inference and machine learning, addressing the fundamental issues of today's AI technology, including explainability, stability and fairness issues. He is recognized as a Distinguished Scientist of ACM, Distinguished Member of CCF and Senior Member of IEEE. He has published more than 100 papers in prestigious conferences and journals in machine learning and data mining. He is one of the most cited authors in network embedding. His recent research won the IEEE Multimedia Best Department Paper Award, IEEE ICDM 2015 Best Student Paper Award, IEEE ICME 2014 Best Paper Award, ACM MM12 Grand Challenge Multimodal Award, MMM13 Best Paper Award, and were selected into the Best of KDD special issues in 2014 and 2016, respectively. He received ACM China Rising Star Award in 2015, and CCF-IEEE CS Young Scientist Award in 2018.

Jian Pei is a Professor in the School of Computing Science at Simon Fraser University. He is a well-known leading researcher in the general areas of data science, big data, data mining, and database systems. His expertise is on developing effective and efficient data analysis techniques for novel data intensive applications, and transferring his research results to products and business practice. He is recognized as a Fellow of the Royal Society of Canada (Canada's national academy), the Canadian Academy of Engineering, the Association of Computing Machinery (ACM) and the Institute of Electrical and Electronics Engineers (IEEE). He was the editor-in-chief of the IEEE Transactions of Knowledge and Data Engineering (TKDE) in 2013-16, the chair of the Special Interest Group on Knowledge Discovery in Data (SIGKDD) of the Association for Computing Machinery (ACM) in 2017-2021, and a general co-chair or program committee co-chair of many premier conferences. He received many prestigious awards, including the 2017 ACM SIGKDD Innovation Award, the 2015 ACM SIGKDD Service Award, the 2014 IEEE ICDM Research Contributions Award, the British Columbia Innovation Council 2005 Young Innovator Award, an NSERC 2008 Discovery Accelerator Supplements Award (100 awards cross the whole country), an IBM Faculty Award (2006), a KDD Best Application Paper Award (2008), an ICDE Influential Paper Award (2018), a PAKDD Best Paper Award (2014), a PAKDD Most Influential Paper Award (2009), and an IEEE Outstanding Paper Award (2007).

Liang Zhao is an assistant professor at the Department of Compute Science at Emory University. He obtained his PhD degree in 2016 from Virginia Tech in the United States. His research interests include data mining, artificial intelligence, and machine learning, with special interests in spatiotemporal and network data mining, deep learning on graphs, nonconvex optimization, model parallelism, event prediction, and interpretable machine learning. He received AWS Machine Learning Research Award in 2020 from Amazon Company for his research on distributed graph neural networks. He won NSF Career Award in 2020 awarded by National Science Foundation for his research on deep learning for spatial networks, and Jeffress Trust Award in 2019 for his research on deep generative models for biomolecules, awarded by Jeffress Memorial Trust Foundation and Bank of America. He won the Best Paper Award in the 19th IEEE International Conference on Data Mining (ICDM 2019) for the paper of his lab on deep graph transformation. He has also won Best Paper Award Shortlist in WWW 2021 and Best Paper Candidate in ICDM 2021. He was selected as "Top 20 Rising Star in Data Mining" by Microsoft Search in 2016 for his research on spatiotemporal data mining. He is awarded as CIFellow Mentor 2021 by the Computing Community Consortium for his research on deep learning for spatial data.

Xiaojie Guo is an Research Scientist at JD.COM Silicon Valley Research Center. She got her Ph.D. degree from George Mason University. Her research topics include data mining, artificial intelligence, and machine learning, with special interests in deep learning on graphs, graph transformation and generation, and interpretable representation learning. She won the Best Paper Award in ICDM 2019 and has one paper awarded as an ESI Hot and Highly Cited Paper as the first author. She also won the AAAI/IAAI 2022 Award for the applied AI project in JD.com.

REFERENCES

[1] Yu Chen and Lingfei Wu. 2022. Graph Neural Networks: Graph Structure Learning. In *Graph Neural Networks: Foundations, Frontiers, and Applications*. Springer, 297–321.

[2] Yunfei Chu, Jiangchao Yao, Chang Zhou, and Hongxia Yang. 2022. Graph Neural Networks in Modern Recommender Systems. In *Graph Neural Networks: Foundations, Frontiers, and Applications*. Springer, 423–445.

[3] Peng Cui, Lingfei Wu, Jian Pei, Liang Zhao, and Xiao Wang. 2022. Graph Representation Learning. In *Graph Neural Networks: Foundations, Frontiers, and Applications*. Springer, 17–26.

[4] Xiaojie Guo, Shiyu Wang, and Liang Zhao. 2022. Graph Neural Networks: Graph Transformation. In *Graph Neural Networks: Foundations, Frontiers, and Applications*. Springer, 251–275.

[5] Pan Li and Jure Leskovec. 2022. The Expressive Power of Graph Neural Networks. In *Graph Neural Networks: Foundations, Frontiers, and Applications*. Springer, 63–98.

[6] Renjie Liao. 2022. Graph Neural Networks: Graph Generation. In *Graph Neural Networks: Foundations, Frontiers, and Applications*. Springer, 225–250.

[7] Xiang Ling, Lingfei Wu, Chunming Wu, and Shouling Ji. 2022. Graph Neural Networks: Graph Matching. In *Graph Neural Networks: Foundations, Frontiers, and Applications*. Springer, 277–295.

[8] Bang Liu and Lingfei Wu. 2022. Graph Neural Networks in Natural Language Processing. In *Graph Neural Networks: Foundations, Frontiers, and Applications*. Springer, 463–481.

[9] Ninghao Liu, Qizhang Feng, and Xia Hu. 2022. Interpretability in Graph Neural Networks. In *Graph Neural Networks: Foundations, Frontiers, and Applications*. Springer, 121–147.

[10] Hehuan Ma, Yu Rong, and Junzhou Huang. 2022. Graph Neural Networks: Scalability. In *Graph Neural Networks: Foundations, Frontiers, and Applications*. Springer, 99–119.

[11] Jian Tang and Renjie Liao. 2022. Graph Neural Networks for Node Classification. In *Graph Neural Networks: Foundations, Frontiers, and Applications*. Springer, 41–61.

[12] Lingfei Wu, Peng Cui, Jian Pei, and Liang Zhao. 2022. Graph Neural Networks: Foundations, Frontiers, and Applications.

[13] Lingfei Wu, Peng Cui, Jian Pei, Liang Zhao, and Le Song. 2022. Graph Neural Networks. In *Graph Neural Networks: Foundations, Frontiers, and Applications*. Springer, 27–37.

[14] Liang Zhao, Lingfei Wu, Peng Cui, and Jian Pei. 2022. Representation Learning. In *Graph Neural Networks: Foundations, Frontiers, and Applications*. Springer, 3–15.