Responsible AI in Graph Learning and More

Yushun Dong
Department of Computer Science
Florida State University
yd24f@fsu.edu

References

[1] Rong Ding, Yushun Dong, Daniel P Aldrich, Jundong Li, Kelsey Pieper, and Qi Ryan Wang. Post-disaster private well water contamination with geosocial network: A case study of post-hurricane harvey. In *Computing in Civil Engineering 2023*, pages 194–201.

[2] Yushun Dong. Explanatory essay requirements.

[3] Yushun Dong. *Algorithmic Fairness in Graph Machine Learning: Explanation, Optimization, and Certification*. PhD thesis, University of Virginia, 2024.

[4] Yushun Dong, Kaize Ding, Brian Jalaian, Shuiwang Ji, and Jundong Li. Adagnn: Graph neural networks with adaptive frequency response filter. In *Proceedings of the 30th ACM international conference on information & knowledge management*, pages 392–401, 2021.

[5] Yushun Dong, Jian Kang, Hanghang Tong, and Jundong Li. Individual fairness for graph neural networks: A ranking based approach. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 300–310, 2021.

[6] Yushun Dong, Oyku Deniz Kose, Yanning Shen, and Jundong Li. Fairness in graph machine learning: Recent advances and future prospectives. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 5794–5795, 2023.

[7] Yushun Dong, Zhenyu Lei, Zaiyi Zheng, Song Wang, Jing Ma, Alex Jing Huang, Chen Chen, and Jundong Li. Pygdebias: A python library for debiasing in graph learning. In *Companion Proceedings of the ACM on Web Conference 2024*, pages 1019–1022, 2024.

[8] Yushun Dong, Jundong Li, and Tobias Schnabel. When newer is not better: Does deep learning really benefit recommendation from implicit feedback? In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 942–952, 2023.

[9] Yushun Dong, Ninghao Liu, Brian Jalaian, and Jundong Li. Edits: Modeling and mitigating data bias for graph neural networks. In *Proceedings of the ACM web conference 2022*, pages 1259–1269, 2022.

[10] Yushun Dong, Jing Ma, Song Wang, Chen Chen, and Jundong Li. Fairness in graph mining: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 35(10):10583–10602, 2023.

[11] Yushun Dong, Jing Ma, Song Wang, Chen Chen, Jundong Li, Y Dong, J Ma, S Wang, and C Chen. No known security issues. *IEEE Transactions on Knowledge and Data Engineering*, 2023.

[12] Yushun Dong, Yingxia Shao, Xiaotong Li, Sili Li, Lei Quan, Wei Zhang, and Junping Du. Forecasting pavement performance with a feature fusion lstm-bpnn model. In *Proceedings of the 28th ACM international conference on information and knowledge management*, pages 1953–1962, 2019.

[13] Yushun Dong, William Shiao, Yozen Liu, Jundong Li, Neil Shah, and Tong Zhao. Seesaw: Do graph neural networks improve node representation learning for all?

[14] Yushun Dong, Patrick Soga, Yinhan He, Song Wang, and Jundong Li. Graph neural networks are more than filters: Revisiting and benchmarking from a spectral perspective. *arXiv preprint arXiv:2412.07188*, 2024.

[15] Yushun Dong, Song Wang, Zhenyu Lei, Zaiyi Zheng, Jing Ma, Chen Chen, and Jundong Li. A benchmark for fairness-aware graph learning. *arXiv preprint arXiv:2407.12112*, 2024.

[16] Yushun Dong, Song Wang, Jing Ma, Ninghao Liu, and Jundong Li. Interpreting unfairness in graph neural networks via training node attribution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 7441–7449, 2023.

[17] Yushun Dong, Song Wang, Yu Wang, Tyler Derr, and Jundong Li. On structural explanation of bias in graph neural networks. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 316–326, 2022.

[18] Yushun Dong, Binchi Zhang, Zhenyu Lei, Na Zou, and Jundong Li. Idea: A flexible framework of certified unlearning for graph neural networks. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 621–630, 2024.

[19] Yushun Dong, Binchi Zhang, Hanghang Tong, and Jundong Li. Elegant: Certified defense on the fairness of graph neural networks. *arXiv preprint arXiv:2311.02757*, 2023.

[20] Yushun Dong, Binchi Zhang, Yiling Yuan, Na Zou, Qi Wang, and Jundong Li. Reliant: Fair knowledge distillation for graph neural networks. In *Proceedings of the 2023 SIAM International Conference on Data Mining (SDM)*, pages 154–162. Society for Industrial and Applied Mathematics, 2023.

[21] Yushun Dong, Binichi Zhang, Yiling Yuan, Na Zou, Qi Wang, and Jundong Li. Abstract graph neural networks (gnns) have shown satisfying performance on various graph learning tasks. to achieve better fitting capability, most gnns are with a large number of parameters, which makes these gnns computationally expensive. therefore, it is difficult to deploy them onto edge devices with scarce computational resources, eg, mobile phones and wearable smart devices. knowledge distillation (kd) is a common solution to compress gnns, where a light-weighted model (ie, the student model) is encouraged to mimic the behavior of a computationally expensive gnn (ie, the teacher gnn model). nevertheless, most existing gnn-based kd methods lack fairness consideration. as a consequence, the student model usually inherits and even exaggerates the bias from the teacher gnn. to handle such a problem, we take initial steps towards fair knowledge distillation for gnns. specifically, we first formulate a novel problem of fair knowledge distillation for gnn-based teacher-student frameworks. then we propose a principled framework named reliant to mitigate the bias exhibited by the student model. notably the design of reliant is decoupled from any specific teacher and student model structures, and thus can be easily adapted to various gnn-based kd frameworks. we perform extensive experiments on multiple real-world datasets, which corroborates that reliant achieves less biased gnn knowledge distillation while maintaining high prediction utility. open-source code can be found at https://github. com/yushundong/reliant. In *Proceedings of the 2023 SIAM International Conference on Data Mining (SDM)*, 2023.

[22] Xingbo Fu, Chen Chen, Yushun Dong, Anil Vullikanti, Eili Klein, Gregory Madden, and Jundong Li. Spatial-temporal networks for antibiogram pattern prediction. In *2023 IEEE 11th International Conference on Healthcare Informatics (ICHI)*, pages 225–234. IEEE, 2023.

[23] Xingbo Fu, Song Wang, Yushun Dong, Binchi Zhang, Chen Chen, and Jundong Li. Federated graph learning with graphless clients. *arXiv preprint arXiv:2411.08374*, 2024.

[24] Xingbo Fu, Binchi Zhang, Yushun Dong, Chen Chen, and Jundong Li. Federated graph machine learning: A survey of concepts, techniques, and applications. *ACM SIGKDD Explorations Newsletter*, 24(2):32–47, 2022.

[25] Yinhan He, Zaiyi Zheng, Patrick Soga, Yaochen Zhu, Yushun Dong, and Jundong Li. Explaining graph neural networks with large language models: A counterfactual perspective on molecule graphs. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 7079–7096, 2024.

[26] Zheng Huang, Jing Ma, Yushun Dong, Natasha Zhang Foutz, and Jundong Li. Empowering next poi recommendation with multi-relational modeling. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2034–2038, 2022.

[27] Lincan Li, Jiaqi Li, Catherine Chen, Fred Gui, Hongjia Yang, Chenxiao Yu, Zhengguang Wang, Jianing Cai, Junlong Aaron Zhou, Bolin Shen, et al. Political-llm: Large language models in political science. *arXiv preprint arXiv:2412.06864*, 2024.

[28] Zhixun Li, Yushun Dong, Qiang Liu, and Jeffrey Xu Yu. Rethinking fair graph neural networks from re-balancing. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1736–1745, 2024.

[29] Haochen Liu, Song Wang, Yaochen Zhu, Yushun Dong, and Jundong Li. Knowledge graph-enhanced large language models via path selection. *arXiv preprint arXiv:2406.13862*, 2024.

[30] Jing Ma, Yushun Dong, Zheng Huang, Daniel Mietchen, and Jundong Li. Assessing the causal impact of covid-19 related policies on outbreak dynamics: A case study in the us. In *Proceedings of the ACM Web Conference 2022*, pages 2678–2686, 2022.

[31] Xihao Piao, Zheng Chen, Yushun Dong, Yasuko Matsubara, and Yasushi Sakurai. Frednormer: Frequency domain normalization for non-stationary time series forecasting. *arXiv preprint arXiv:2410.01860*, 2024.

[32] Yucheng Shi, Yushun Dong, Qiaoyu Tan, Jundong Li, and Ninghao Liu. Gigamae: Generalizable graph masked autoencoder via collaborative latent space reconstruction. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 2259–2269, 2023.

[33] Weihao Song, Yushun Dong, Ninghao Liu, and Jundong Li. Guide: Group equality informed individual fairness in graph neural networks. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1625–1634, 2022.

[34] Jihong Wang, Minnan Luo, Jundong Li, Yun Lin, Yushun Dong, Jin Song Dong, and Qinghua Zheng. Empower post-hoc graph explanations with information bottleneck: A pre-training and fine-tuning perspective. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 2349–2360, 2023.

[35] Song Wang, Yushun Dong, Kaize Ding, Chen Chen, and Jundong Li. Few-shot node classification with extremely weak supervision. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, pages 276–284, 2023.

[36] Song Wang, Yushun Dong, Xiao Huang, Chen Chen, and Jundong Li. Faith: Few-shot graph classification with hierarchical task graphs. *arXiv preprint arXiv:2205.02435*, 2022.

[37] Song Wang, Yushun Dong, Xiao Huang, Chen Chen, and Jundong Li. Learning hierarchical task structures for few-shot graph classification. *ACM Transactions on Knowledge Discovery from Data*, 18(3):1–20, 2024.

[38] Song Wang, Yushun Dong, Binchi Zhang, Zihan Chen, Xingbo Fu, Yinhan He, Cong Shen, Chuxu Zhang, Nitesh V Chawla, and Jundong Li. Safety in graph machine learning: Threats and safeguards. *arXiv preprint arXiv:2405.11034*, 2024.

[39] Song Wang, Peng Wang, Yushun Dong, Tong Zhou, Lu Cheng, Yangfeng Ji, and Jundong Li. On demonstration selection for improving fairness in language models. In *Workshop on Socially Responsible Language Modelling Research*.

[40] Song Wang, Peng Wang, Tong Zhou, Yushun Dong, Zhen Tan, and Jundong Li. Ceb: Compositional evaluation benchmark for fairness in large language models. *arXiv preprint arXiv:2407.02408*, 2024.

[41] Xiao Wang, Quan Yuan, Zhihan Liu, Yushun Dong, Xiaojuan Wei, and Jinglin Li. Learning route planning from experienced drivers using generalized value iteration network. In *Internet of Vehicles. Technologies and Services Toward Smart Cities: 6th International Conference, IOV 2019, Kaohsiung, Taiwan, November 18–21, 2019, Proceedings 6*, pages 88–100. Springer International Publishing, 2020.

[42] Yu Wang, Yuying Zhao, Yushun Dong, Huiyuan Chen, Jundong Li, and Tyler Derr. Improving fairness in graph neural networks via mitigating sensitive attribute leakage. In *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, pages 1938–1948, 2022.

[43] Zhiming Xu, Xiao Huang, Yue Zhao, Yushun Dong, and Jundong Li. Contrastive attributed network anomaly detection with data augmentation. In *Pacific-Asia conference on knowledge discovery and data mining*, pages 444–457. Springer International Publishing Cham, 2022.

[44] Dong Yushun, Ma Jing, Chen Chen, and Li Jundong. Fairness in graph mining: A survey. *arXiv preprint arXiv*, 2204, 2022.

[45] Binchi Zhang, Yushun Dong, Chen Chen, Yada Zhu, Minnan Luo, and Jundong Li. Adversarial attacks on fairness of graph neural networks. *arXiv preprint arXiv:2310.13822*, 2023.

[46] Binchi Zhang, Yushun Dong, Tianhao Wang, and Jundong Li. Towards certified unlearning for deep neural networks. *arXiv preprint arXiv:2408.00920*, 2024.

[47] Xianren Zhang, Jing Ma, Yushun Dong, Chen Chen, Min Gao, and Jundong Li. Sd-attack: Targeted spectral attacks on graphs. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 352–363. Springer Nature Singapore Singapore, 2024.