# Fairness and Sparse Link Prediction through Ranking-Aware Metrics and Attribute Enhancement

João Mattos
jrm28@rice.edu
Rice University
Houston, TX, USA

## Abstract

This research addresses challenges in link prediction, a foundational task in graph machine learning, by investigating the limitations in the evaluation frameworks for sparse graphs and fair link recommendation. Our first project introduces **Gelato**, a novel similarity-based framework that leverages graph learning to integrate node attributes into the graph topology and a topological heuristic optimized via a ranking loss, designed to tackle extreme class imbalance and hard negative sampling. Gelato outperforms state-of-the-art GNNs in accuracy and scalability on multiple real-world datasets. The second project focuses on fairness in link prediction by highlighting the importance of considering the ranking aspect in fairness metrics, borrowing properties from Information Retrieval, and proposing a simple post-processing method to mitigate bias effectively while maintaining high utility. Experiments on diverse datasets demonstrate the efficacy of our fairness approach compared with existing methods. Collectively, this work enhances the robustness and ethical considerations of link prediction models.

**ACM Reference Format:**
João Mattos. 2018. Fairness and Sparse Link Prediction through Ranking-Aware Metrics and Attribute Enhancement. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (KDD '25).* ACM, New York, NY, USA, 3 pages. https://doi.org/XXXXXXX.XXXXXXX

## 1 Introduction

Graphs naturally represent complex relational data in domains such as social networks, recommender systems, natural language processing, and scientific modeling. Link prediction—inferring missing or future edges—plays a critical role both as an independent task and as a fundamental component in graph-based machine learning pipelines. Recently, Graph Neural Networks (GNNs) have achieved remarkable success in node and graph classification, and have become the dominant approach for link prediction [9, 10, 29]. However, our empirical studies reveal that standard evaluation protocols for GNN-based link prediction are overly optimistic, especially on sparse and modular real-world graphs, due to biased negative sampling and balanced testing sets.

Furthermore, fairness of link prediction models remains underexplored, with prior methods[3, 5, 11, 12, 15] often neglecting the influence of ranking on exposure bias [17]. Our research addresses these gaps by proposing novel methodologies combining rigorous analysis, methodological innovations, and practical fairness interventions.

## 2 Gelato — Enhancing Link Prediction in Sparse Graphs via Graph Learning and Topological Heuristics

### 2.1 Motivation and Challenges

Real-world graphs such as citation networks are extremely sparse with strong modular structures, where fewer than 0.2% of node pairs are linked [9]. GNN-based link prediction methods typically frame the task as binary classification, handling extreme class imbalance via uniform negative sampling [10, 29]. This results in balanced training and evaluation, which do not capture the challenges of the real-world scenario where all disconnected pairs (i.e., all negatives) must be considered. Under this realistic evaluation (we refer to as *unbiased sampling*), classical topological heuristics (e.g. CommonNeighbors) often outperform GNNs, showing GNNs struggle to learn topological similarity measures under extreme imbalance[13].

We claim that successful link prediction in realistic settings requires discriminating positive links from hard negatives that are similar in structure and attributes. We demonstrate that the key limitations of existing GNN approaches include reliance on classification losses sensitive to imbalance, biased training mimicking biased evaluation, and relying on topological heuristics that struggle in distinguishing between positive samples and hard negatives.

### 2.2 Methodology

**Gelato** combines an expressive topological heuristic (Autocovariance [4]) with graph learning to jointly leverage topology and node attributes for link prediction. Our approach learns an attribute-enhanced adjacency matrix via a ranking-based N-pair loss [18], distinguishing between easy and hard negative pairs.

The method comprises three components:

- *Negative Sampling*: We sample negative and positive pairs following the distribution found in the original graph. This attempts to simulate the challenging scenario of link prediction found in the real world, for both training and evaluation.
- *Graph Learning*: Starting from the original graph adjacency matrix, Gelato augments edges between node pairs based on untrained attribute similarity, yielding an enhanced graph. A Multi-layer Perceptron (MLP) learns edge weights on this graph to integrate attribute information directly into the topology. The MLP is trained end-to-end with respect to


*KDD '25, Toronto, ON*

ACM ISBN 978-1-4503-XXXX-X/2018/06
https://doi.org/XXXXXXX.XXXXXXX

a ranking loss, enabling the enhanced adjacency matrix to capture similarities in both feature and structure domains.
- *Topological Heuristic*: Gelato applies the Autocovariance (AC) heuristic on the attribute-enhanced adjacency matrix to compute pairwise link scores. This exploits local structural similarity while also leveraging node attributes.
- *N-pair Loss*: To enable distinguishing between positive links and hard negatives and learning in unbalanced scenarios, we model link prediction as a ranking task, adopting the N-pair Loss. Since most pairs are *easy* negatives, this contrastive approach enables the model to capture what sets easy from hard negatives apart, addressing the issue with binary classification that considers all negative pairs the same.

## 2.3 Results

Experiments on five real-world graph datasets show that Gelato consistently outperforms state-of-the-art GNN-based link prediction methods such as SEAL[29], BUDDY[2], and NCN[21] in terms of hits@1000 and average precision. Ablation studies confirm that combining graph learning, AC heuristic, N-pair ranking loss, and partition-based negative sampling is crucial for superior performance. Sensitivity analyses reveal robustness with respect to thresholding parameters controlling the incorporation of learned vs. attribute similarity.

We also empirically demonstrate the efficacy of the partitioned sampling mechanism. This sampling technique filters out easy negative samples, focusing training on the (more informative) hard negatives, enabling a lower training time. Gelato thus addresses fundamental challenges in link prediction on sparse graphs, providing a scalable alternative to GNN-based approaches.

# 3 Fairness in Link Prediction — Ranking-Aware Metrics and Bias Mitigation

## 3.1 Problem Statement and Motivation

Link prediction increasingly impacts socially sensitive applications, making fairness with respect to protected attributes a critical concern. Prior work often focuses on demographic parity[3, 5, 11, 12, 15], considering only dyadic fairness settings and not accounting for the ranking nature of link prediction, which governs exposure and can induce bias [8, 17, 26]. In addition, previous GNN-based approaches address bias in link prediction by either obtaining fair node representations (node-based approaches) or by forcing predictions across sensitive groups to have similar positive rates. Both approaches raise fairness concerns, since GNNs are not expressive enough for link-level tasks [23, 29], casting doubt on the capacity of these methods to distinguish and ensure fairness between sensitive groups.

We tackle the task of fairness in link prediction on a graph with sensitive attributes by demonstrating the necessity of considering more than two sensitive groups and incorporating ranking notions on fairness evaluation. Our work aims to shed light on hidden biases that arise from the existing dyadic formulation of fairness in link prediction based on demographic parity as a fairness metric and propose a new post-processing method for bias mitigation.

## 3.2 Contributions

- *Dyadic Fairness and Hidden Biases*: We formalize and demonstrate the existence of two main sources of biases that are not addressed by previous literature. The first one is based on the dyadic formulation and evaluation of fairness adopted by previous approaches. The second is based on unawareness of ranking positions, which leaves an open door for exposure biases and feedback loops in which the same group of pairs is overrepresented in recommendations.
- *Fairness Properties:* We propose to adopt two novel properties that evaluation metrics should satisfy: (1) respecting the original graph's edge distribution to prevent hidden bias, and (2) sensitivity to exposure disparity inherent in ranking scenarios. We demonstrate how these properties help mitigate the hidden biases occurring in the current link prediction methods existing in the literature.
- *Bias Mitigation:* We propose a simple post-processing technique that allocates predicted scores from decoupled classifiers to enforce these fairness properties while maintaining high link prediction utility and scalability.

## 3.3 Experimental Validation

Evaluations on six real-world datasets [6] with varying graph structures and sensitive attribute distributions corroborate our claims and reveal biases undetected by demographic parity in previous approaches. Results show our post-processing method significantly reduces statistical unfairness measured via normalized discounted KL divergence (NDKL) and preserves *prec@k*, when compared to pre-processing (EDITS [5]), in-processing (GRAPHAIR [12]), and random-walk (FairWalk [15]) baselines.

# 4 Future Research Directions

Spatial-temporal learning is an important research area with applications in domains such as Traffic Analysis[19, 22, 25, 27], Simulation [16] and Weather Forecasting[14]. Previous approaches are successful in adopting GNNs and Graph Transformers as the main backbone to capture spatial-temporal inductive biases in these domains. One of our future research directions is to explore pair-wise relationships in spatial-temporal tasks, an active area of research[24].

Finally, our second research direction to be explored is the intersection between Graph Neural Networks and Large Language Models (LLMs) ([1, 7, 20]). In particular, recent research sheds light on how LLMs can be interpreted as Graph Neural Networks [20] and on the theoretical properties that arise by doing so. On the opposite direction, adopting the Message Passing mechanism for reasoning tasks or AI agents' communication networks [28] is a promising way of improving reasoning and the general capabilities of these systems that we also aim to explore in the near future.

# References


[1] Federico Barbero, Alvaro Arroyo, Xiangming Gu, Christos Perivolaropoulos, Michael Bronstein, Razvan Pascanu, et al. 2025. Why do LLMs attend to the first token? *arXiv preprint arXiv:2504.02732* (2025).

[2] Benjamin Paul Chamberlain, Sergey Shirobokov, Emanuele Rossi, Fabrizio Frasca, Thomas Markovich, Nils Hammerla, Michael M Bronstein, and Max Hansmire. 2023. Graph Neural Networks for Link Prediction with Subgraph Sketching. In *ICLR*.

[3] Sean Current, Yuntian He, Saket Gurukar, and Srinivasan Parthasarathy. 2022. FairEGM: Fair Link Prediction and Recommendation via Emulated Graph Modification. In *Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–14. doi:10.1145/3551624.3555287 arXiv:2201.11596 [cs].
[4] J-C Delvenne, Sophia N Yaliraki, and Mauricio Barahona. 2010. Stability of graph communities across time scales. *PNAS* 107, 29 (2010), 12755–12760.
[5] Yushun Dong, Ninghao Liu, Brian Jalaian, and Jundong Li. 2022. EDITS: Modeling and Mitigating Data Bias for Graph Neural Networks. In *Proceedings of the ACM Web Conference 2022*. 1259–1269. doi:10.1145/3485447.3512173 arXiv:2108.05233 [cs].
[6] Yushun Dong, Jing Ma, Song Wang, Chen Chen, and Jundong Li. 2023. Fairness in graph mining: A survey. *IEEE Transactions on Knowledge and Data Engineering* (2023).
[7] Bahare Fatemi, Jonathan Halcrow, and Bryan Perozzi. 2023. Talk like a graph: Encoding graphs for large language models. *arXiv preprint arXiv:2310.04560* (2023).
[8] Shantanu Gupta, Hao Wang, Zachary Lipton, and Yuyang Wang. 2021. Correcting exposure bias for link recommendation. In *International Conference on Machine Learning*. PMLR, 3953–3963.
[9] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. 2020. Open graph benchmark: Datasets for machine learning on graphs. In *NeurIPS*.
[10] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
[11] Peizhao Li, Yifei Wang, Han Zhao, Pengyu Hong, and Hongfu Liu. 2021. On dyadic fairness: Exploring and mitigating bias in graph connections. In *International conference on learning representations*.
[12] Hongyi Ling, Zhimeng Jiang, Youzhi Luo, Shuiwang Ji, and Na Zou. 2023. LEARNING FAIR GRAPH REPRESENTATIONS VIA AUTOMATED DATA AUGMENTATIONS. (2023).
[13] João Mattos, Zexi Huang, Mert Kosan, Ambuj Singh, and Arlei Silva. 2025. Attribute-Enhanced Similarity Ranking for Sparse Link Prediction. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*. To appear.
[14] Ilan Price, Alvaro Sanchez-Gonzalez, Ferran Alet, Tom R Andersson, Andrew El-Kadi, Dominic Masters, Timo Ewalds, Jacklynn Stott, Shakir Mohamed, Peter Battaglia, et al. 2025. Probabilistic weather forecasting with machine learning. *Nature* 637, 8044 (2025), 84–90.
[15] Tahleen Rahman, Bartlomiej Surma, Michael Backes, and Yang Zhang. 2019. Fairwalk: Towards Fair Graph Embedding. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, Macao, China, 3289–3295. doi:10.24963/ijcai.2019/456
[16] Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. 2020. Learning to simulate complex physics with graph networks. In *International conference on machine learning*. PMLR, 8459–8468.
[17] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of Exposure in Rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2219–2228. doi:10.1145/3219819.3220088 arXiv:1802.07281 [cs].
[18] Kihyuk Sohn. 2016. Improved deep metric learning with multi-class n-pair loss objective. In *NeurIPS*.
[19] Yan Tian, Kaili Zhang, Jianyuan Li, Xianxuan Lin, and Bailin Yang. 2018. LSTM-based traffic flow prediction with missing data. *Neurocomputing* 318 (2018), 297–305.
[20] Alex Vitvitskyi, Joao GM Araújo, Marc Lackenby, and Petar Veličković. 2025. What makes a good feedforward computational graph? *arXiv preprint arXiv:2502.06751* (2025).
[21] Xiyuan Wang, Haotong Yang, and Muhan Zhang. 2023. Neural Common Neighbor with Completion for Link Prediction. *arXiv preprint arXiv:2302.00890* (2023).
[22] Jianli Xiao. 2019. SVM and KNN ensemble learning for traffic incident detection. *Physica A: Statistical Mechanics and its Applications* 517 (2019), 29–35.
[23] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2018. How Powerful are Graph Neural Networks?. In *ICLR*.
[24] Lu Yi, Jie Peng, Yanping Zheng, Fengran Mo, Zhewei Wei, Yuhang Ye, Yue Zixuan, and Zengfeng Huang. 2025. TGB-Seq Benchmark: Challenging Temporal GNNs with Complex Sequential Dynamics. *arXiv preprint arXiv:2502.02975* (2025).
[25] Mehmet Yildirimoglu and Nikolas Geroliminis. 2013. Experienced travel time prediction for congested freeways. *Transportation Research Part B: Methodological* 53 (2013), 45–63.
[26] Meike Zehlike and Carlos Castillo. 2020. Reducing Disparate Exposure in Ranking: A Learning To Rank Approach. In *Companion Proceedings of the Web Conference 2020*. 172–175. doi:10.1145/3366424.3383534 arXiv:1805.08716 [cs].
[27] Anqin Zhang, Qizheng Liu, and Ting Zhang. 2022. Spatial–temporal attention fusion for traffic speed prediction. *Soft Computing* 26, 2 (2022), 695–707.
[28] Guibin Zhang, Yanwei Yue, Zhixun Li, Sukwon Yun, Guancheng Wan, Kun Wang, Dawei Cheng, Jeffrey Xu Yu, and Tianlong Chen. 2024. Cut the crap: An economical communication pipeline for llm-based multi-agent systems. *arXiv preprint arXiv:2410.02506* (2024).
[29] Muhan Zhang and Yixin Chen. 2018. Link prediction based on graph neural networks. In *NeurIPS*.