Towards Trustworthy, Efficient, and Scalable Machine Learning

Yezi Liu *

Department of Electrical Engineering and Computer Science
University of California Irvine, Irvine, CA, United States
yezil3@uci.edu

1 Work Completed by September 2024

Throughout the development of machine learning, researchers have increasingly focused on the challenges of trustworthiness, efficiency, and scalability. Our research specifically addresses these critical aspects. I am the primary contributor to all studies in this thesis summary, including experiments and writing, except for (Xie, Liu, and Shen 2022), which was conducted collaboratively.

1.1 Trustworthy Machine Learning

Studying trustworthiness in machine learning is essential due to the potentially harmful consequences of its applications. We will introduce three works that correspond to three key aspects of Trustworthy Machine Learning.

Address Bias. Existing studies primarily concentrate on creating fairness-aware algorithms for graph neural networks (GNNs). However, these approaches often modify the model architecture, making it difficult to apply a general GNN for fair decision-making. To address this limitation, we propose a data-centric approach (Liu 2023) that enables the learning of a fair graph. Consequently, a general GNN can be trained on this fair graph without requiring debiasing techniques, allowing it to produce fair predictions.

Protect Data Privacy. In recent years, there has been significant progress in developing legal regulations to protect data privacy. To address these requirements in the graph data, graph unlearning has emerged, allowing algorithms to effectively handle such deletion requests. In this study, titled 'FGU: Enabling Group Fairness in Graph Unlearning via Global Alignment', we address the challenge of ensuring fairness and privacy in GNNs. Our proposed framework, FGU, utilizes two primary strategies: shard debiasing and global alignment. Shard debiasing focuses on reducing discrimination in predictions across sensitive groups, while global alignment promotes fairness in the overall distribution by adjusting model weights through back-propagation.

Improve Explainability. Deep learning models frequently operate as "black boxes", obscuring their decision-making processes. To address this issue, we propose the DGExplainer (Xie, Liu, and Shen 2022), a method designed to identify the key subset of node features that influence predictions in dynamic GNNs. DGExplainer employs Layer-wise Relevance Propagation to clarify how these dynamic GNNs arrive at their predictions.

1.2 Scalable Machine Learning

Scalable Graph Neural Networks for Large Graphs. In this work (Liu and Shen 2024), we present TinyGraph, a novel joint graph condensation framework. Unlike traditional methods that focus solely on condensing nodes, TinyGraph is designed to condense both nodes and features within large-scale graphs. To optimize the trainable condensed graph, we utilize a gradient matching strategy, complemented by a structure-aware dimensionality function that maintains the integrity of the graph structure. This dual condensation approach enables TinyGraph to achieve high test accuracies across various datasets, showcasing its efficiency.

Scalable Image Classifier for Large Image Data. In this study, we present a novel approach, TinyData, that tackles the challenge of training deep neural networks on large-scale image datasets. Our method, known as joint condensation, simultaneously reduces both feature and sample size. Specifically, TinyData integrates a trainable function into the training process, which minimizes a gradient matching loss. This allows for effective data condensation while retaining essential information from the original dataset. The results not only validate the effectiveness of our approach but also showcase its applicability across various model architectures.

2 Planned Projects Before February 2025

In the near future, my research will focus on efficient machine learning, as efficiency and scalability are essential for addressing the time-consuming and resource-intensive nature of training machine learning models. I also plan to explore trustworthiness across various machine learning applications, including link prediction and text-to-image models.

*Thesis supervisor: Mohsen Imani, Donald Bren School of Information and Computer Sciences, University of California, Irvine, Irvine, CA, United States, m.imani@uci.edu.

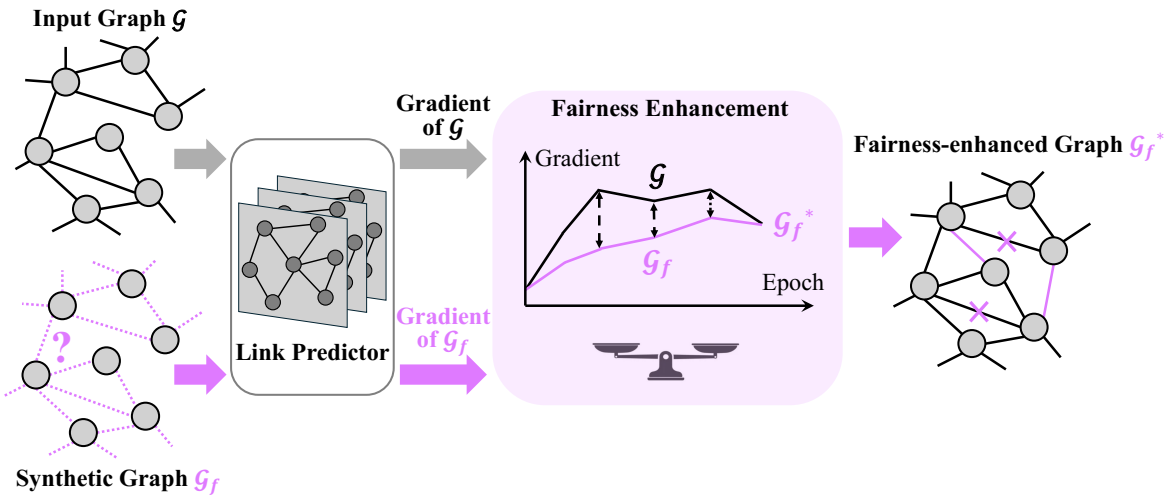

Figure 1: FairGraph aims to learn a fairness-enhanced graph that promotes fairness while preserving utility.

Dataset **Acc.** (%)	Data Cond. Ratio (r_n)	Full Feature	Condensed Feature	
		Coarsening	DC-PCA	FairGraph
MNIST 99.6±0.0	0.017%	88.4 ± 0.3	86.7 ± 1.2	$\mathbf{88.3 \pm 0.7}$
	0.17%	81.1 ± 0.6	91.3 ± 0.4	$\mathbf{95.7 \pm 1.1}$
	0.83%	96.4 ± 0.5	94.1 ± 0.7	$\mathbf{97.8 \pm 1.6}$
CIFAR10 84.7±0.1	0.02%	19.3 ± 0.6	24.6 ± 0.7	$\mathbf{27.6 \pm 0.8}$
	0.2%	13.4 ± 0.5	40.4 ± 1.1	$\mathbf{43.6 \pm 1.0}$
	1%	27.3 ± 0.4	48.8 ± 0.9	$\mathbf{52.4 \pm 0.7}$

Table 1: The performance comparison of baselines. The condensed number of features for joint condensation is 16×16 for MNIST and CIFAR10.

2.1 Fairness-aware Graph Hyperdimensional Computing (FairGHDC)

Hyperdimensional computing (HDC) seeks to emulate brain circuits more robustly and efficiently than traditional neural networks by representing information as points in a high-dimensional space, known as hypervectors. These hypervectors are typically binary or bipolar vectors with ten thousand dimensions. In this study, we examine the fairness problem within the HDC-based node classification framework. This issue is particularly challenging due to the complexities involved in graph hyperdimensional computing, which incorporates various training settings alongside numerous existing machine learning models. Specifically, the design of effective debiasing techniques during the training process for Node Hypervector Encoding, Edge Hypervector Encoding, and state updates has not yet been thoroughly explored.

2.2 Promoting Fairness in Link Prediction with Graph Enhancement

In this work (Liu, Chen, and Imani 2024), we introduce FairLink, a novel method designed to learn a fairness-enhanced graph that eliminates the need for debiasing during link predictor training. FairLink preserves link prediction accuracy by ensuring that the training trajectory of the enhanced graph closely mirrors that of the original input graph. At the same time, it promotes fairness by minimizing the absolute difference in link probabilities between node pairs belonging to the same sensitive group, compared to those between node pairs from different sensitive groups.

Acknowledgements

FGU, FairLink, and the FairGHDC were supported in part by the DARPA Young Faculty Award, the National Science Foundation (NSF) under Grants #2127780, #2319198, #2321840, #2312517, and #2235472, the Semiconductor Research Corporation (SRC), the Office of Naval Research through the Young Investigator Program Award, and Grants #N00014-21-1-2225 and #N00014-22-1-2067, Army Research Office Grant #W911NF2410360. Additionally, support was provided by the Air Force Office of Scientific Research under Award #FA9550-22-1-0253, along with generous gifts from Xilinx and Cisco.

References

Liu, Y. 2023. FairGraph: Automated Graph Debiasing with Gradient Matching. In *CIKM*, 4135–4139.

Liu, Y.; Chen, H.; and Imani, M. 2024. Promoting Fairness in Link Prediction with Graph Enhancement. *arXiv preprint arXiv:2409.08658*.

Liu, Y.; and Shen, Y. 2024. TinyGraph: Joint Feature and Node Condensation for Graph Neural Networks. *arXiv preprint arXiv:2407.08064*.

Xie, J.; Liu, Y.; and Shen, Y. 2022. Explaining dynamic graph neural networks via relevance back-propagation. *arXiv preprint arXiv:2207.11175*.