Jiong Zhu Santa Clara, CA | jiongzhu@umich.edu | jiongzhu.net

Research Interests

I am an Applied Scientist at Amazon working on **pretraining** foundational **Small Language Models (SLMs)** with up to 10B parameters. I focus on building tiny but mighty models following state-of-the-art practice in LLM architectural design and pretraining data curation, and conduct research on **data & hyperparameter scaling law**. Besides generative model pretraining, I also have experience in **contrastive fine-tuning** of **text embedding models** for **retrieval** tasks.

During my PhD, my research focused on **Graph Representation Learning** and **Graph Neural Networks (GNNs)**, particularly their limitations, robustness, fairness, and applications in complex, large-scale environments. A key aspect of my work involves advancing GNN performance on **heterophilous graphs** (where connected nodes have dissimilar labels & features).

Education

University of Michigan, Ann Arbor, MI, USA
- **Ph.D. in Computer Science and Engineering** **Sep. 2019 - Jun. 2024**
  Advancing Graph Neural Networks for Complex Data: A Perspective Beyond Homophily
  Advisor: Prof. Danai Koutra
- **M.S. in Electrical and Computer Engineering** **Aug. 2017 - May. 2019**

Xi'an Jiaotong University, Xi'an, Shaanxi, China
- **B.Eng. in Automation (Honors Engineering Program)** **Aug. 2013 - Jun. 2017**
- **Special Class for Gifted Young** **Aug. 2011 - Jun. 2013**
  A honors high school program for nationwide selected middle school graduates.

Work Experiences

Amazon, Palo Alto, CA, USA **Jun. 2024 - Present**
Applied Scientist, Store Foundational AI
- Responsible for pretraining foundational Small Language Models (SLMs) with less than 10B params following state-of-the-art LLM architectural design and pretraining data curation.
- Conducts research on data & hyperparameter scaling law for optimizing SLM pretraining.

Amazon, Palo Alto, CA, USA **Jun. 2023 - Mar. 2024**
Applied Scientist Intern, Search Science & AI
- Developed a retrieval system by fine-tuning text encoders (e.g., T5) in a two-tower approach, tailored to meet the unique requirements and goals of creating a new Amazon search experience.
- Implemented a comprehensive full-stack pipeline, which includes data curation and cleaning, language model fine-tuning with contrastive learning, inferencing, and evaluation.
- Conducted targeted experiments to assess the model's performance and effectiveness.

Amazon, Palo Alto, CA, USA **Jan. 2022 - Feb. 2023**
Applied Scientist Intern, Search Science & AI
- Developed an advanced distributed framework for Graph Neural Network (GNN) training on web-scale datasets. This framework streamlines previous methods and enhances performance and speed. It achieves these improvements by utilizing randomized graph partitions and a novel time-based model aggregation mechanism.
- Published a research paper in ACM Transactions on Knowledge Discovery from Data (TKDD).

Adobe Inc., San Jose, CA, USA **Jun. 2020 - Aug. 2020**
Data Science Research Intern (Remote), Adobe Research
- Designed CPGNN, a GNN framework which adapts to the homophily level in the graphs by incorporating belief propagation with learnable compatibility into GNN message passing.
- Paper published in AAAI 2021 [4]; US and CN patents pending.

Publications

Google Scholar: https://scholar.google.com/citations?user=KjGFQ0QAAAAJ
- **Impact**: 1400+ citations (as of September 2024)

In Peer-reviewed Conference Proceedings

[1] *On the Impact of Feature Heterophily on Link Prediction with Graph Neural Networks*
Jiong Zhu*, Gaotang Li*, Yao-An Yang, Jing Zhu, Xuehao Cui, Danai Koutra
NeurIPS 2024 - Conference on Neural Information Processing Systems. [link]

[2] *On Performance Discrepancies Across Local Homophily Levels in Graph Neural Networks*
Donald Loveland, Jiong Zhu, Mark Heimann, Ben Fish, Michael T Schaub, Danai Koutra
LoG 2023 - Learning on Graphs Conference 2023. [link]
★ Oral presentation at proceedings track.

[3] *How does Heterophily Impact the Robustness of Graph Neural Networks? Theoretical Connections and Practical Implications*
Jiong Zhu, Junchen Jin, Donald Loveland, Michael T Schaub, and Danai Koutra
KDD 2022 - ACM SIGKDD Conference on Knowledge Discovery and Data Mining. [link]
★ Research track with oral presentation, acceptance rate: 15.0%.

[4] *Graph Neural Networks with Heterophily*
Jiong Zhu, Ryan A Rossi, Anup Rao, Tung Mai, Nedim Lipka, Nesreen K Ahmed, and Danai Koutra
AAAI 2021 - AAAI Conference on Artificial Intelligence. [link]

[5] *Beyond Homophily in Graph Neural Networks: Current Limitations and Effective Designs*
Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra
NeurIPS 2020 - Conference on Neural Information Processing Systems. [link]
★ Receives more than 970 citations as of September 2024, and was taught in Northeastern University (Spring 2021, PHYS 7332: Network Data Science 2).

[6] *GroupINN: Grouping-based Interpretable Neural Network for Classification of Limited, Noisy Brain Data*
Yujun Yan, Jiong Zhu, Marlena Duda, Eric Solarz, Chandra Sripada, and Danai Koutra
KDD 2019 - ACM SIGKDD Conference on Knowledge Discovery and Data Mining. [link]
★ Oral presentation at research track, acceptance rate: 9.2%.

In Peer-reviewed Journal

[7] *Simplifying Distributed Neural Network Training on Massive Graphs: Randomized Partitions Improve Model Aggregation*
Jiong Zhu, Aishwarya Reganti, Edward Huang, Charles Dickens, Nikhil Rao, Karthik Subbian, Danai Koutra
TKDD 2024 - ACM Transactions on Knowledge Discovery from Data. [link]

In Peer-reviewed Workshops

[8] *Graph Coarsening via Convolution Matching for Scalable Graph Neural Network Training*
Charles Dickens, Edward Huang, Aishwarya Reganti, Jiong Zhu, Karthik Subbian, Danai Koutra
DCAI-WebConf 2024 - Workshop on Data-centric Artificial Intelligence. [link]

[9] *Simplifying Distributed Neural Network Training on Massive Graphs: Randomized Partitions Improve Model Aggregation*
Jiong Zhu, Aishwarya Reganti, Edward Huang, Charles Dickens, Nikhil Rao, Karthik Subbian, Danai Koutra
LLW-ICML 2023 - Workshop on Localized Learning. [link]

[10] *On graph neural network fairness in the presence of heterophilous neighborhoods*
Donald Loveland, Jiong Zhu, Mark Heimann, Ben Fish, Michael T Schaub, Danai Koutra.
DLG-KDD 2022 - Workshop on Deep Learning on Graphs. [link]

Invited Contribution

[11] *Heterophily and Graph Neural Networks: Past, Present and Future*
Jiong Zhu, Yujun Yan, Mark Heimann, Lingxiao Zhao, Leman Akoglu, and Danai Koutra
IEEE Data Engineering Bulletin (2023). [link]

Technical Skills

Programming: Python, C++, MATLAB, C#; **Large-scale Data Processing**: PySpark; **ML Toolkits**: TensorFlow, PyTorch, NeMo, Hugging Face, DGL, scikit-learn, pandas, Faiss; **Visualization**: Matplotlib, Plotly; **Miscellaneous**: LaTeX, Git, Docker, AWS

Selected Research Projects

How does Heterophily Impact the Robustness of Graph Neural Networks? Theoretical Connections and Practical Implications
Published at KDD 2022 [3]

- Goal: To explore the relation between heterophily and robustness in GNNs and its implications.
- Showed theoretically and empirically that effective structural attacks on homophilous graphs increase heterophily, and extended the analysis of change in homophily to heterophilous graphs.
- Demonstrated that a key design principle which improves GNN performance under heterophily can also inherently offer increased empirical and certifiable robustness, with up to 32.92% gain in performance against state-of-the-art attacks and 3.4x gain in certifiable robustness metrics.

Beyond Homophily in Graph Neural Networks: Current Limitations and Effective Designs
Published at NeurIPS 2020 [5]

- Goal: To reveal and address the limitations of GNN models under heterophilous graphs.
- Identified three effective designs backed by theoretical analysis that boost learning from the graph structure under heterophily, and proposed a new model, H_2GCN, based on the designs.
- Extensive empirical analysis shows up to 40% gain in accuracy by the designs in heterophily.

Students Mentored

- **Junchen Jin** (B.S.E. in University of Michigan, 2020-2022). Published paper at KDD 2022 on heterophily and robostness in GNN [3]. Currently Machine Learning Scientist at Paypal.
- **Gaotang Li** (B.S. in University of Michigan, 2021-2024). Published paper at NeurIPS 2024 as a co-first author. Currently PhD student at the University of Illinois Urbana-Champaign.
- **Yao-An Yang** (B.S. in University of Michigan, 2022-). Published paper at NeurIPS 2024.
- **Xuehao Cui** (B.S. in University of Michigan, 2022-2024). Published paper at NeurIPS 2024.

Teaching Experience

University of Michigan, Ann Arbor, MI, USA

- *Graduate Student Instructor*, EECS 598-008: Advanced Data Mining **Winter 2019**
  Teaching assistant for the graduate-level course with 45 students. Duties include leading discussion sessions, designing and grading assignments, and answering questions.
- *Lecturer*, Big Data Summer Institute **Summer 2021**
  Gave asynchronous lectures on data mining to undergraduates with a discussion session.

Academic Service

Workshop Organizer:

- GLB 2021-2023: Workshop on Graph Learning Benchmarks at WebConf and KDD

Program Committee: CIKM 2022 (🏅 best PC distinction), CIKM 2021, AAAI 2021
Reviewer: NeurIPS 2024, WebConf 2021, 2023, LoG 2022-2023 (🏅 top reviewer distinction in 2022), JMLR, TNNLS, TKDE, TKDD, JMLC, GroundedML 2022

Invited Presentations

Beyond Homophily in Graph Neural Networks: Current Limitations and Effective Designs [5]

- *Poster*, Google Workshop on Scalable Algorithms for Semi-supervised and Unsupervised Learning, CA, USA (Virtual) **Oct. 2021**
- *Poster*, Graph Exploitation Symposium (GraphEx 2021), MA, USA (Virtual) **May. 2021**

On the Relationship between Heterophily and Robustness of GNNs [3]

- *Poster*, Graph Exploitation Symposium (GraphEx 2022), MA, USA (Virtual) **May. 2022**

Beyond Homophily in GNNs: Current Limitations, Effective Designs, and Impacts on Robustness [3]-[5]

- *Invited Talk*, DataFun GNN Summit, China (Virtual, 4000+ live audience) **Jun. 2022**

Outreach

AI4ALL, instructor of AI introduction lectures to Michigan high-school students. **2021**

Discover Engineering, volunteer of hands-on coding sessions for 8th-10th grade kids. **2022**

Honors and Awards

LoG 2022 Top Reviewer, awarded to the top 20 reviewers served in the 2022 Learning on Graphs Conference. **2022**

CIKM 2022 Best PC, awarded to the top 5.5% of the 523 program committee members who are responsible for paper reviews. **2022**

KDD 2019 Student Travel Award, awarded by ACM SIGKDD. **2019**

National Second Prize in CUMCM (Contemporary Undergraduate Mathematical Contest in Modeling), awarded to the top 5.7% of the 22233 teams. **2014**

Pengkang Scholarship, awarded to top 2% students by XJTU. **2014**