

Zhaocheng Zhu

2788 San Tomas Expy, Santa Clara, CA 95051, United States
(+1)650-450-6584 ◊ <https://kiddozhu.github.io> ◊ zhaochengz@nvidia.com

WORK EXPERIENCE

- NVIDIA**, Santa Clara, United States *Feb. 2025 - Now*
Senior Research Scientist
· Nemotron Post-training Team
- Google**, Sunnyvale, United States *Sep. 2023 - Apr. 2024*
Student Researcher
· Tool Use, Retrieval-Augmented Generation, Multi-Step Reasoning
Mentor: Hanjun Dai, Yuan Xue
- Google**, Remote from Canada *Apr. 2023 - June 2023*
Student Researcher
· Factuality, Hallucination, Rule Learning
Mentor: Hanjun Dai, Yuan Xue
- Microsoft Research Asia**, Beijing, China *Sep. 2017 - May 2018*
Research Intern
· Object Detection, Video Object Detection, Keypoint Detection
Mentor: Jifeng Dai
- Mitsubishi Information Technology R&D Center**, Kamakura, Japan *July 2016 - Aug. 2016*
Research Intern
· Dialog Systems, Natural Language Understanding
Mentor: Yusuke Koji

EDUCATION

- Mila - Québec AI Institute / Université de Montréal**, Canada *Sep. 2019 - Sep. 2024*
Ph.D. in Computer Science
· Reasoning, Knowledge Graphs, Machine Learning Systems
Advisor: Jian Tang
- Mila - Québec AI Institute / Université de Montréal**, Canada *Sep. 2018 - Aug. 2019*
M.Sc. in Computer Science, transferred to Ph.D.
· Graph Representation learning, Machine Learning Systems
Advisor: Jian Tang
- Peking University**, China *Sep. 2014 - July 2018*
B.S. in Computer Science (with honors)
· Natural Language Processing, Unsupervised Representation Learning, Word Semantics
Advisor: Junfeng Hu
· Computer Vision, Object Detection
Thesis Advisor: Yizhou Wang, Jifeng Dai (Microsoft Research Asia)
- Carnegie Mellon University**, Pittsburgh, United States *July 2017 - Sep. 2017*
Visiting Student
· Interpretable Machine Learning
Advisor: Pradeep Ravikumar

TUTORIALS

Reasoning on Knowledge Graphs: Symbolic or Neural?

Meng Qu, **Zhaocheng Zhu**, Jian Tang. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022.

PUBLICATIONS

Landscape of Thoughts: Visualizing the Reasoning Process of Large Language Models

Zhanke Zhou*, **Zhaocheng Zhu***, Xuan Li*, Mikhail Galkin, Xiao Feng, Sanmi Koyejo, Jian Tang, Bo Han. In *International Conference on Learning Representations*, 2026.

Enhancing Link Prediction in Biomedical Knowledge Graphs with BioPathNet

Emy Yue Hu, Svitlana Oleshko, Samuele Firmani, Hui Cheng, **Zhaocheng Zhu**, Maria Ulmer, Matthias Arnold, Maria Colom-Tatch, Jian Tang, Sophie Xhonneux, Annalisa Marsico. In *Nature Biomedical Engineering*, 2026.

Overcoming Long-Context Limitations of State-Space Models via Context-Dependent Sparse Attention

Zhihao Zhan*, Jianan Zhao*, **Zhaocheng Zhu**, Jian Tang. In *Conference on Neural Information Processing Systems*, 2025.

From Passive to Active Reasoning: Can Large Language Models Ask the Right Questions under Incomplete Information?

Zhanke Zhou*, Xiao Feng*, **Zhaocheng Zhu**, Jiangchao Yao, Sanmi Koyejo, Bo Han. In *International Conference on Machine Learning*, 2025.

Learning Equivalence Classes of Bayesian Network Structures with GFlowNet

Michelle Liu, **Zhaocheng Zhu**, Olexa Bilaniuk, Emmanuel Bengio. In *Transactions on Machine Learning Research*, 2025.

Fully-inductive Node Classification on Arbitrary Graphs

Jianan Zhao*, **Zhaocheng Zhu***, Mikhail Galkin, Hesham Mostafa, Michael Bronstein, Jian Tang. In *International Conference on Learning Representations*, 2025.

A Foundation Model for Zero-shot Logical Query Reasoning

Mikhail Galkin, Jincheng Zhou, Bruno Ribeiro, Jian Tang, **Zhaocheng Zhu**. In *Conference on Neural Information Processing Systems*, 2024.

Neural Graph Reasoning: Complex Logical Query Answering Meets Graph Databases

Hongyu Ren*, Mikhail Galkin*, Michael Cochez, **Zhaocheng Zhu**, Jure Leskovec. In *Transactions on Machine Learning Research*, 2024.

Towards Foundation Models for Knowledge Graph Reasoning

Mikhail Galkin, Xinyu Yuan, Hesham Mostafa, Jian Tang, **Zhaocheng Zhu**. In *International Conference on Learning Representations*, 2024.

Towards Foundation Models for Molecular Learning on Large-Scale Multi-Task Datasets

Dominique Beaini, Shenyang Huang, Joao Alex Cunha, Gabriela Moisescu-Pareja, Oleksandr Dymov, Samuel Maddrell-Mander, Callum McLean, Frederik Wenkel, Luis Mller, Jama Hussein Mohamud, Ali Parviz, Michael Craig, Micha Koziarski, Jiarui Lu, **Zhaocheng Zhu**, Cristian Gabellini, Kerstin Klaser, Josef Dean, Cas Wognum, Maciej Sypetkowski, Guillaume Rabusseau, Reihaneh Rabbany, Jian Tang, Christopher Morris, Ioannis Koutis, Mirco Ravanelli, Guy Wolf, Prudencio Tossou, Hadrien Mary, Therence Bois, Andrew Fitzgibbon, Baej Banaszewski, Chad Martin, Dominic Masters. In *International Conference on Learning Representations*, 2024.

A*Net: A Scalable Path-based Reasoning Approach for Knowledge Graphs

Zhaocheng Zhu*, Xinyu Yuan*, Mikhail Galkin, Sophie Xhonneux, Ming Zhang, Maxime Gazeau,

Jian Tang. In *Conference on Neural Information Processing Systems*, 2023.

Large Language Models can Learn Rules

Zhaocheng Zhu, Yuan Xue, Xinyun Chen, Denny Zhou, Jian Tang, Dale Schuurmans, Hanjun Dai. *arXiv preprint arXiv:2310.07064*, 2023.

GraphText: Graph Reasoning in Text Space

Jianan Zhao, Le Zhuo, Yikang Shen, Meng Qu, Kai Liu, Michael Bronstein, **Zhaocheng Zhu**, Jian Tang. *arXiv preprint arXiv:2310.01089*, 2023.

Inductive Logical Query Answering in Knowledge Graphs

Mikhail Galkin, **Zhaocheng Zhu**, Hongyu Ren, Jian Tang. In *Conference on Neural Information Processing Systems*, 2022.

PEER: A Comprehensive and Multi-Task Benchmark for Protein Sequence Understanding

Minghao Xu, Zuobai Zhang, Jiarui Lu, **Zhaocheng Zhu**, Yangtian Zhang, Chang Ma, Runcheng Liu, Jian Tang. In *Conference on Neural Information Processing Systems (Datasets and Benchmarks Track)*, 2022.

Neural-Symbolic Models for Logical Queries on Knowledge Graphs

Zhaocheng Zhu, Mikhail Galkin, Zuobai Zhang, Jian Tang. In *International Conference on Machine Learning*, 2022.

TorchDrug: A Powerful and Flexible Machine Learning Platform for Drug Discovery

Zhaocheng Zhu, Chence Shi, Zuobai Zhang, Shengchao Liu, Minghao Xu, Xinyu Yuan, Yangtian Zhang, Junkun Chen, Huiyu Cai, Jiarui Lu, Chang Ma, Runcheng Liu, Louis-Pascal Xhonneux, Meng Qu, Jian Tang. *arXiv preprint arXiv:2202.08320*, 2022.

Neural Bellman-Ford Networks: A General Graph Neural Network Framework for Link Prediction

Zhaocheng Zhu, Zuobai Zhang, Louis-Pascal Xhonneux, Jian Tang. In *Conference on Neural Information Processing Systems*, 2021. Rank 12/39 in the link prediction task of OGB-LSC.

KEPLER: A Unified Model for Knowledge Embedding and Pre-trained Language Representation

Xiaozhi Wang, Tianyu Gao, **Zhaocheng Zhu**, Zhengyan Zhang, Zhiyuan Liu, Juanzi Li, Jian Tang. In *Transactions of the Association for Computational Linguistics*, 2021.

GraphAF: A Flow-based Autoregressive Model for Molecular Graph Generation

Chence Shi^{*}, Minkai Xu^{*}, **Zhaocheng Zhu**, Weinan Zhang, Ming Zhang, Jian Tang. In *International Conference on Learning Representations*, 2020

Self-Adaptive Network Pruning

Jinting Chen, **Zhaocheng Zhu**, Cheng Li, Yuming Zhao. In *International Conference on Neural Information Processing*, 2019. (Best Student Paper Finalist)

GraphVite: A High-Performance CPU-GPU Hybrid System for Node Embedding

Zhaocheng Zhu, Shizhen Xu, Meng Qu and Jian Tang. In *The World Wide Web Conference*, 2019.

Saliency Supervision: An Intuitive and Effective Approach for Pain Intensity Regression

Conghui Li, **Zhaocheng Zhu** and Yuming Zhao. In *International Conference on Neural Information Processing*, 2018.

Context Aware Document Embedding

Zhaocheng Zhu and Junfeng Hu. *arXiv preprint arXiv:1707.01521*, 2017.

Dialog State Tracking with Attention-based Sequence-to-Sequence Learning

Takaaki Hori, Hai Wang, Chiori Hori, Shinji Watanabe, Bret Harsham, Jonathan Le Roux, John R

Hershey, Yusuke Koji, Yi Jing, **Zhaocheng Zhu** and Takeyuki Aikawa. In *IEEE Spoken Language Technology Workshop (SLT)*, 2016. (Runner up at Dialog State Tracking Challenge 5)

SELECTED PROJECTS

TorchDrug: A Powerful and Flexible Machine Learning Platform for Drug Discovery
(leader of TorchDrug team)

Machine learning development platform for drug discovery in PyTorch. Support 6 tasks, more than 25 models. Over 1,500 stars and 77,000 downloads.

Featured in *PyTorch ecosystem*. Supported by NVIDIA Applied Research Accelerator Program.

<https://torchdrug.ai> <https://github.com/DeepGraphLearning/torchdrug>

GraphVite: A General and High-Performance Graph Embedding System for Various Applications (leader of GraphVite team)

General and high-performance graph embedding system. Support 3 applications, 10 models and more than 40 baseline benchmarks. Over 1,200 stars and 102,000 downloads.

<https://graphvite.io> <https://github.com/DeepGraphLearning/graphvite>

Literature of Deep Learning for Graphs (with Meng Qu and Weiping Song)

Comprehensive paper list of deep learning for graphs. Over 3,000 stars.

<https://github.com/DeepGraphLearning/LiteratureDL4Graph>

HONORS AND AWARDS

NVIDIA Applied Research Accelerator Program	2022
Tuition Fee Exemption Scholarships, Université de Montréal	2019 - 2021
Outstanding Graduate Student, Peking University	2018
Top Talent Class of EECS, Peking University	2016 - 2018
Outstanding Research Award, Peking University	2016
Kwang-Hua Scholarship, Peking University	2016
Honorable Mention, Mathematical Contest in Modeling (MCM)	2016
Merit Student, Peking University	2015
Tung OCCL Scholarship, Peking University	2015
Honorable Mention, Mathematical Contest in Modeling (MCM)	2015
Group Champion, Peking University Debate Competition for Freshman	2014
Second Prize (ranked 86th in China), National Olympiad in Informatics (NOI)	2013

SERVICE

Academic Reviewer

- ICLR, ICML 2026
- ICLR, ICML, JMLR, COLM, NeurIPS 2025
- ICML, NeurIPS 2024
- WWW, ICML, KDD, NeurIPS, LoG 2023
- TKDE, ICML, NeurIPS 2022
- GNNSys Workshop (MLSys) 2021
- DLG Workshop (KDD) 2020
- GRL Workshop (NeurIPS) 2019

Mentoring & Teaching

- Mentor, Buddy Program for New Students, Mila 2019 - 2021
- Reviewer, PhD / MSc Recruitment, Mila 2019 - 2021
- Teaching Assistant, Machine Learning II, HEC Montréal Winter 2020

Social & Leadership

- Lab representatives¹ (11 out of 1,000+ students), Mila *Nov. 2022 - Nov. 2023*
- Vice President, Society of Photography, Peking University *Sep. 2017 - June 2018*
- Leader of story portrait group, Society of Photography, Peking University *Apr. 2015 - Aug. 2017*
- Member of organizers & news team, HackPKU, Peking University *Apr. 2016*
- Member of debate team, School of EECS, Peking University *Sep. 2014 - June 2015*

SKILLS

Programming Languages:

- Proficient: C/C++, Python, CUDA, Pascal/Object Pascal
- Capable: MATLAB, SQL, Bash, Assembly, HTML/CSS

Software Development: OOP design, refactorization (> 20k lines), parallel programming

Deep Learning: PyTorch, MXNet, Keras, TensorFlow

Toolchains: Git, Slurm, L^AT_EX, GDB, CMake, Conda(build), Pip(build), Photoshop

Languages: Mandarin Chinese(native), English(proficient), French(beginner)

Open-Source Contributions: PyTorch-Geometric, Gensim, PyKEEN, OpenFold, Graphium, NeMo-RL

¹A committee that serves students and represents students in meetings with professors and staffs.