

JIANG Chenyu

Mobile: (+852) 55711843

Email: jchenyu@connect.hku.hk

Homepage: <https://chenyu-jiang.github.io/>

Address: Rm 412, Chow Yei Ching Building, The University of Hong Kong, Hong Kong SAR

EDUCATION

The University of Hong Kong (HKU)

Sep 2020-Present (Late 2025)

- **Program:** Ph.D. in Computer Science
- **Supervisor:** Prof. Chuan Wu

The Chinese University of Hong Kong (CUHK)

Sep 2016-Aug 2020

- **Program:** B.Sc. in Computer Science
- Graduated with **First Class Honors**

RESEARCH INTEREST

My current research focuses on the optimization of distributed DNN training and inference systems. I am also broadly interested in deep learning compilation technologies.

PUBLICATIONS

- **Chenyu Jiang**, Zhenkun Cai, Ye Tian, Zhen Jia, Yida Wang, Chuan Wu, *DCP: Addressing Input Dynamism in Long-Context Training via Dynamic Context Parallelism*, in **SOSP'25**
- **Chenyu Jiang**, Ye Tian, Zhen Jia, Chuan Wu, Yida Wang, Shuai Zheng, *Lancet: Accelerating Mixture-of-Experts Training by Overlapping Weight Gradient Computation and All-to-All Communication*, in **MLSys'24**
- **Chenyu Jiang**, Zhen Jia, Yida Wang, Shuai Zheng, Chuan Wu, *DynaPipe: Optimizing Multi-task Training through Dynamic Pipelines*, in **EuroSys'24**
- Hanpeng Hu, **Chenyu Jiang**, Yuchen Zhong, Yanghua Peng, Chuan Wu, Yibo Zhu, Haibin Lin, *dPRO: A Generic Performance Diagnosis and Optimization Toolkit for Expediting Distributed DNN Training*, in **MLSys'22**
- Shiyuan Deng, Xiao Yan, Kelvin K.W. Ng, **Chenyu Jiang**, James Cheng, *Pyramid-A General Framework for Distributed Similarity Search*, IEEE Big Data 2019 Short Paper

INTERNSHIPS

Meituan

May 2025-Present

- **Team:** Central Local Commerce - Computing Platform - LLM pretraining

Amazon Web Services

Jun 2021-Feb 2025

- **Position:** Applied Scientist Intern
- **Supervisors:** Yida Wang, Zhen Jia

TALKS

- "Optimizing Distributed Large Language Model Training Under Novel Parallelisms and Dynamic Inputs", Anuttacon, Online, Oct 2025
- "DCP: Addressing Input Dynamism in Long-Context Training via Dynamic Context Parallelism", SOSP'25 Conference, Seoul, South Korea, Oct 2025
-

- “Lancet: Accelerating Mixture-of-Experts Training by Overlapping Weight Gradient Computation and All-to-All Communication”, *MLSys '24 Conference*, San Jose, USA, May 2024
- “Optimizing novel DNN training scenarios-from data to model architecture”, *Bytedance Ltd*, San Jose, USA, May 2024
- “Optimizing novel DNN training scenarios-from data to model architecture”, *Amazon Web Services (AWS)*, San Jose, USA, May 2024
- “DynaPipe: Optimizing Multi-task Training through Dynamic Pipelines”, *EuroSys '24 Conference*, Athens, Greece, April 2024

SELECTED AWARDS

- HKU University Postgraduate Fellowship [2020]
- Professor Charles K. Kao Research Exchange Scholarship [2019]
- ELITE Stream Student Scholarship [2017-2019]
- CSE Outstanding Academic Performance [2018-2019]
- Wu Yee Sun College Academic Excellence Scholarship [2017-2018]

TEACHING EXPERIENCES

- Teaching Assistant for COMP9602 Convex Optimization, HKU
- Teaching Assistant for COMP3322 Modern Technologies on World Wide Web, HKU
- Teaching Assistant for COMP3234 Computer and Communication Networks, HKU

OTHER RESEARCH EXPERIENCES

Final Year Project: Distributed Deep Learning Systems, CUHK Sep 2019-Aug 2020

Instructor: Prof. James Cheng

- Investigated communication optimizations to accelerate distributed DNN training with data parallelism.

Summer Internship: Fairness in Machine Learning, Carnegie Mellon University Jun 2019-Aug 2019

Instructor: Prof. Pulkit Grover

- Investigated methods to alleviate undesired political bias from recommendation systems using tools from information theory.

Research Internship: Representation Learning in NLP, CUHK Jun 2018-Jan 2019

Instructor: Prof. Helen Meng, Dr. Kerney Wu

- Examined methods to create multi-sense dictionary embeddings.