

# Ziyi Zhang

zzhang765@wisc.edu  
github.com/541736690  
RobeZH.com/about/  
linkedin.com/in/ziyi-zhang-19315b16a/

## EDUCATION

---

### University of Wisconsin, Madison

B.S. in Computer Sciences & B.S. in Mathematics

Sep 2017 - May 2020 & Sep 2021 - May 2022

Advisors: Shivaram Venkataraman, Dieter van Melkebeek

- GPA Overall: 3.95 / 4.0, with Core GPA (CS/Math Courses): 3.98 / 4.0
- Selected course projects: benchmarking different DBMS tuning methods, accelerating the einsum operator in PyTorch, compiler for a subset of C language, database index manager using B+ tree, adding functionality to xv6 operating system

## RESEARCH INTEREST

---

My research interests lie in **distributed systems** and **machine learning systems**. The projects I worked on include low-level DNN operator accelerations, high-level ML system pipeline optimizations, and video processing system development. I am particularly interested in designing **high-performance ML/Big Data Systems**. I am actively looking for PhD positions that start in Fall 2022.

## PUBLICATION

---

[1] [Conference Paper] Bagpipe: Accelerating Deep Recommendation Model Training. Saurabh Agarwal, **Ziyi Zhang**, Shivaram Venkataraman. 2021. *In Submission*

## RESEARCH EXPERIENCE

---

### System Group, University of Wisconsin - Madison

Madison, WI, United States

Research Assistant

Sep 2019 - Sep 2020

Advisor: Shivaram Venkataraman

- Developed a flexible framework for processing video DNN tasks with ~2500 lines of C++
- Used the framework to perform tasks in video DNN pipelines including object detection and object tracking

### ByteDance AI Lab

Beijing, China

Research Intern

Sep 2020 - Feb 2021

Advisor: Bairen Yi, Yibo Zhu

- Designed several strategies for overlapping computation and communication across operators
- Developed high performance DNN operators for model parallelism using CUDA/CUTLASS/NCCL
- Saved up to 30% of the communication time for a single operator in the end-to-end benchmark

### ByteDance AML (Applied Machine Learning)

Beijing, China / Madison, WI, United States

Research & Engineering Intern

Feb 2021 - Nov 2021

Advisor: Yibo Zhu, Shivaram Venkataraman

- Research
  - Explored ways to accelerate distributed recommendation system training
  - Formulated the training process as a graph problem and used Metis library to find better parameter placement in asynchronous training
  - Investigated data access improvements such as caching to accelerate synchronous training
- Engineering
  - Incorporated new DNN operators into the recommendation system training framework
  - Built tools for quickly checking correctness of gradients during the training process

## MENTORSHIP/TEACHING EXPERIENCE

---

### Programming Contest Team, University of Wisconsin - Madison

Madison, WI, United States

Coach, Discussion Leader

Sep 2018 - Present

- Led discussions on problem sets and topics including graph theory and data structure

- Discussed with different teams about contest strategies
- Set up 3 contests with other coaches and wrote test generators and checkers for the problems
- Attended ICPC World Finals 2019 as a coach

**Computer Sciences Dept., University of Wisconsin - Madison** Madison, WI, United States  
*Undergraduate Teaching Assistant in Course "Introduction to Algorithms"* Jan 2019 - May 2019

- Designed and led group discussions on topics including greedy and dynamic programming
- Held individual help sessions for students to answer their questions on course material and algorithm design

**Mentorship Program at Chinese Undergrad Student Assoc.** Madison, WI, United States  
*Mentor* Sep 2020 - Present

- Held one-to-one sessions to share academic information/experience
- Helped students explore academic opportunities inside and outside school

## AWARDS

---

- 2021 UW Madison CS Dept. Scholarship in Excellence in Competitive Programming
- 2020 ICPC World Finalist (Ranked 17th out of 118 teams worldwide)
- 2019 ICPC North Central North American Regional Champion (Ranked 1st out of 180 teams regionwide)
- 2019 UW Madison CS Dept. Golden Brick Awards
- Dean's List on Fall 2018 and Fall 2019

## SELECTED COURSEWORK

---

- Graduate CS: Big Data Systems, High-Performance Computing, Advanced Algorithms
- Undergraduate CS: Operating Systems, Networks, Databases, Compilers, Algorithms, Artificial Intelligence, Security
- Undergraduate Math: Stochastic Process, Combinatorics, Numerical Analysis, Abstract Algebra, Analysis
- Coursera Online Course: Neural Networks and Deep Learning, Convolutional Neural Networks, Recurrent Neural Networks

## SKILLS

---

### Engineering

- **Programming Languages** - C/C++, Python, Java, SQL
- **Framework** - CUDA, PyTorch, TensorFlow, MPI, Spark
- **Development** - CMake, Docker, Git, Linux

### General

- **Language Skills** - Mandarin (Primary Language), English (Fluent)