Feng Zhu

Department of Electrical and Computer Engineering, NC State University, Raleigh, NC 27606
919-274-7431 fzhu5@ncsu.edu Website Google Scholar GitHub

Education

North Carolina State University **Aug 2023 - Present**
PhD candidate *in Electrical and Computer Engineering, GPA: 4.00/4.00* *Raleigh, NC, USA*

- *Advisors:* Co-advised by Prof. Aritra Mitra and Prof. Robert W. Heath
- *Relevant Coursework:* Theoretical Foundations of Large-Scale Machine Learning and Optimization (A+), Convex Optimization in Data Science (A−), Detection and Estimation Theory (A+), Real Analysis (A+), Linear Transformation & Matrix Theory (A+), Systems Control Engineering (A+)

Fudan University **Sep 2020 - Jun 2023**
Master *of Engineering in Communications and Information Systems, GPA: 3.79/4.00* *Shanghai, China*

- *Advisor:* Prof. Xin Wang
- *Relevant Coursework:* Artificial Intelligence (A), Optimization Theory and Algorithms (A), Fundamentals of Digital Communication (A)

Fudan University **Sep 2016 - Jun 2020**
Bachelor *of Engineering in Communications Engineering, GPA: 3.47/4.00* *Shanghai, China*

- *Relevant Coursework:* Principles of Communications (A)

University of Texas at Austin **Aug 2018 - Dec 2018**
Exchange student *in Cockrell School of Engineering, GPA: 4.00/4.00* *Austin, TX, USA*

- *Relevant Coursework:* Wireless Communication Lab (A)

Research Interests

- **Federated Reinforcement Learning (PhD Focus)** – Developing sample-efficient federated algorithms for decision-making problems in multi-agent dynamic environments; published work at **IEEE CDC**, **ACC**, **IEEE ASILOMAR**, **TMLR**.
- **Federated and Distributed Learning** – Designing communication-efficient and robust learning algorithms for large-scale federated systems; published work at **IEEE TSP**, **IEEE ISIT**, etc.
- **Machine Learning for Wireless Systems** – Experienced with ML-based channel prediction and GPU-accelerated simulation during internship at **Qualcomm**.

Experience

Qualcomm Inc. **May 2024 – August 2024**
Engineer Intern – Transformer-Based Channel-Prediction Project *San Diego, CA, USA*

- Conducted a comprehensive literature review on **machine learning–based channel prediction** and presented key findings to the research team.
- Developed and optimized **Python (PyTorch)** and **MATLAB** scripts for large-scale simulations on **GPU servers**, benchmarking model performance across diverse channel conditions.
- Investigated **model design trade-offs** between configurations using multiple smaller neural networks and fewer larger ones; achieved a comparable overall parameter count and accuracy with an approximately **60× reduction in individual model size**, improving efficiency and scalability.
- Performed **comparative experiments** between **Transformer**-based and **autoregressive (AR)** models, demonstrating that the ML approach achieved **1.7 dB performance gain** under equivalent testing conditions.

Publications and Preprints

- **Zhu, F.**, Heath, R.W., & Mitra, A. (2026). A short and unified convergence analysis of the SAG, SAGA, and IAG algorithms. *ArXiv preprint arXiv: 2602.05304. Submitted to ICML.*
- **Zhu, F.**, Heath, R.W., & Mitra, A. (2026). Variance-reduced federated Q-Learning with constant communication. *Submitted to TAC.*
- **Zhu, F.**, Heath, R.W., & Mitra, A. (2026). One-shot clustering for personalized federated policy evaluation. *Submitted to TSP.*
- Maity, S.*, **Zhu, F.***, Mitra, A., & Heath, R.W. (2026). Variance-reduced Q-Learning over static and time-varying networks. *(*Equal contribution.) Accepted by **ACC**.*
- **Zhu, F.**, Heath, R.W., & Mitra, A. (2026). Achieving fast finite-time rates for heterogeneous federated stochastic approximation under Markovian sampling. ***Transactions on Machine Learning Research (TMLR)**.*
- **Zhu, F.**, Heath, R.W., & Mitra, A. (2025). Distributed stochastic approximation with constant communication. *Accepted at **IEEE ASILOMAR**.*
- **Zhu, F.**, Heath, R.W., & Mitra, A. (2024). Towards fast rates for federated and multi-task reinforcement learning. In *2024 **IEEE Conference on Decision and Control (CDC)*** (pp. 2658-2663).
- **Zhu, F.**, Zhang, J., & Wang, X. (2024). STSyn: Speeding up local SGD with straggler-tolerant synchronization. ***IEEE Transactions on Signal Processing (TSP)**.*
- **Zhu, F.**, Zhang, J., Liu, S., & Wang, X. (2023). DRAG: Divergence-based adaptive aggregation in federated learning on non-iid Data. *ArXiv preprint arXiv: 2309.01779. Submitted to IEEE TIFS.*
- **Zhu, F.**, Zhang, J., & Wang, X. (2023). Communication-efficient local SGD with age-based worker selection. *The Journal of Supercomputing*, 1-23.
- **Zhu, F.**, Zhang, J., Simeone, O., & Wang, X. (2022). Adaptive worker grouping for communication-efficient and straggler-tolerant distributed SGD. In *2022 **IEEE International Symposium on Information Theory (ISIT)*** (pp. 2996-3000).

Awards

Graduate Merit Award **2023**
North Carolina State University *Raleigh, NC*

Excellent Academic Scholarship for Master Students - First Prize **2021**
Fudan University *Shanghai, China*

Excellent Academic Scholarship for Master Students - Excellence Award **2020**
Fudan University *Shanghai, China*

Excellent Undergraduate Scholarship - Second Prize **2017, 2018, 2020**
Fudan University *Shanghai, China*

Technical Skills

Languages: Python, Matlab, Latex

Frameworks & Libraries: Pytorch, Numpy, Scipy, Transformers

ML/RL Tooling: Hugging Face Transformers, Open AI Gym

Expertise: Reinforcement Learning (policy gradient, Q-learning, multi-agent RL), Federated Learning, Distributed Optimization