Gauri Joshi

http://andrew.cmu.edu/user/gaurij
Google Scholar

Employment

Carnegie Mellon University Pittsburgh PA
Associate Professor of Electrical and Computer Engineering — Jul 2022 – present
Assistant Professor of Electrical and Computer Engineering — Sept 2017 – Jul 2022

IBM T. J. Watson Research Center — Yorktown Heights NY
Research Staff Member — July 2016 – Aug 2017

Education

Massachusetts Institute of Technology — Sept 2012– Jun 2016

Ph.D in Electrical Engineering & Computer Science, GPA 5.0/5.0 (minor in Finance)
Thesis: Using Redundancy to Reduce Latency in Cloud Systems
Committee: Gregory Wornell (advisor), Emina Soljanin, Devavrat Shah

Massachusetts Institute of Technology — Sept 2010– Jun 2012

S.M. in Electrical Engineering & Computer Science, GPA 5.0/5.0
Thesis: On Playback Delay in Streaming Communication — **Best Thesis Award**
Advisors: Gregory Wornell, Yuval Kochman

Indian Institute of Technology Bombay — Jul 2005–Jun 2010

B. Tech & M. Tech in Electrical Engineering; GPA 9.77/10.0 — **Institute Gold Medal**
Thesis: On Relay-Assisted Cellular Networks
Advisor: Abhay Karandikar

Awards and Honors

- Named as one of MIT Technology Review's **35 innovators under 35** — 2022
- **NSF CAREER Award** — 2021
- **Best Paper Award** at ACM SIGMETRICS — 2020
- Facebook Faculty Research Award — 2020
- NSF CRII Research Initiation Award — 2019
- **Distinguished Student Paper Award**, NeurIPS FedML Workshop — 2019
- **Qualcomm Innovation Fellowship**, awarded to my students Jianyu and Ankur — 2019
- Intl. Conf. on Machine Learning (ICML) – **Outstanding Reviewer Award (top 5%)** — 2019
- IBM Faculty Research Award — 2018
- Rising Stars in EECS workshop invited participant — 2015
- William Martin Memorial Award for **Best Masters Thesis** in Computer Science, MIT — 2012
- Morris Joseph Levin Award for **Outstanding Oral Thesis Presentation**, MIT — 2012

- Claude E. Shannon Research Assistantship, MIT 2015–2016
- Schlumberger **Faculty for the Future Fellowship** 2011–2015
- Irwin and Joan Jacobs Presidential Fellowship, MIT 2010–2011
- **Institute Gold Medal** for highest GPA in the undergraduate class, IIT Bombay 2010
- Institute Silver Medal for highest GPA in the masters class, IIT Bombay 2010
- Best student in Communications & Signal processing Award, IIT Bombay 2009
- Selected among top 50 students in India for the International Chemistry Olympiad camp 2005

PhD Student and Postdoc Advising

- Ankur Mallick Fall 2017-Summer 2022
- Jianyu Wang Fall 2017-Summer 2022
- Samarth Gupta (co-advised with Prof. Osman Yağan) Fall 2017-Spring 2022
- Ting-Wu (Rudy) Chin (co-advised with Prof. Diana Marculescu) Spring 2020-Summer 2021
- Ahmet Inci (co-advised with Prof. Diana Marculescu) Spring 2020-Summer 2022
- Yae Jee Cho Fall 2019-present
- Tuhinangshu Choudhury (co-advised with Prof. Weina Wang) Fall 2019-present
- Divyansh Jhunjhunwala Fall 2020-present
- Jiin Woo (co-advised with Prof. Yuejie Chi) Fall 2021-present
- Pranay Sharma, post-doc Fall 2021-present
- Shuli Jiang Spring 2022-present

Teaching and Textbook Development

Courses Taught at Carnegie Mellon:

- *18-461/661: Introductory Machine Learning for Engineers* Spring 2019-present
  Taken by $\sim$ 100 students every semester and offered across the Pittsburgh, SV and Rwanda campuses of Carnegie Mellon
- *18-667: Algorithms for Distributed Machine Learning* Fall 2017-present
  Advanced graduate course covering the state-of-the-art algorithms for distributed machine learning and optimization. First ever course featuring federated learning, to my knowledge.

Textbook Development: Currently writing a book based on 18-667 lecture notes, which will be published in the series of Synthesis lectures on Learning, Networks and Algorithms

Outreach and Mentoring

- Invited Panelist, Networking Women Mentoring Panel, ACM MobiHoc conference 2020
- Mentor, Rising Stars in EECS Workshop for women pursuing academic careers 2019-present
- Organized a de-stress with faculty workshop for CMU ECE women 2018,2019
- Mentor in the EECS REFS group for conflict management and mediation Jan 2014 - Jun 2016

- Completed a 40-hour training course in conflict management and mediation, MIT Jan 2014
- Institute Student Mentor for freshman at IIT Bombay 2009,2010

Selected Publications

Summary: Published more than 70 research papers, including 13 journal papers and 40 peer-reviewed conference papers. According to Google Scholar, my work has been cited $> 4,000$ times with an h-index of 25.

Papers with more than 50 **citations** (full list here):

[14] P. Kairouz, B. McMahan, J. Wang, G. Joshi, et al "Advanced and Open Problems in Federated Learning", *Foundations and Trends in Machine Learning*, June 2021, arXiv:1912.04977

[13] J. Wang and G. Joshi, "Cooperative SGD: A unified Framework for the Design and Analysis of Communication-Efficient SGD Algorithms", *Journal on Machine Learning Research (JMLR)*, Sept 2021, arXiv:1808.07576

[12] J. Wang, Q. Liu, G. Joshi and V. Poor, "Tackling the Objective Inconsistency Problem in Heterogeneous Federated Optimization", *Proceedings of Neural Information Processing Systems (NeurIPS)*, Dec 2020, arXiv:2007.07481

[11] Y. Cho, J. Wang and G. Joshi "Client Selection in Federated Learning: Convergence Analysis and Power-of-Choice Selection Strategies", *International Conference on Artificial Intelligence and Statistics (AISTATS)*, Apr 2022, arXiv:2010.01243

[10] A. Mallick, U. Sheth, G. Palanikumar, M. Chaudhari, and G. Joshi, "Rateless Codes for Near-Perfect Load Balancing in Distributed Matrix-Vector Multiplication", *Proceedings of ACM SIGMETRICS*, June 2020, arXiv:1804.10331 **Best Paper Award**

[9] J. Wang, A. Sahu, G. Joshi, and S. Kar, "MATCHA: Speeding up Decentralized SGD via Matching Decomposition Sampling", *Proceedings of the Neural Information Processing Systems (NeurIPS) Federated Learning workshop*, Dec 2019, **Distinguished Student Paper Award**, arXiv:1905.09435

[8] J. Wang, G. Joshi, "Adaptive Communication Strategies to Achieve the Best Error-Runtime Trade-off in Local-Update SGD", *SysML Conference*, Mar 2019, arXiv:1810.08313

[7] D. Wang, G. Joshi, G. Wornell, "Efficient Straggler Replication in Parallel Computing ", *ACM Transactions on Modeling and Performance Evaluation of Computing Systems*, 2019 arXiv:1503.03128

[6] S. Dutta, G. Joshi, S. Ghosh, P. Dube, P. Nagpurkar, "Slow and Stale Gradients Can Win the Race: Error-Runtime Trade-offs in Distributed SGD", *International Conference on Artificial Intelligence and Statistics (AISTATS)*, Apr 2018.

[5] G. Joshi, E. Soljanin, G. Wornell, "Efficient Redundancy Techniques to Reduce Latency in Cloud Systems", *ACM Transactions on Modeling and Performance Evaluation of Computing Systems*, volume 2, issue 2, May 2017, arXiv:1508.03599

[4] G. Joshi, E. Soljanin, G. Wornell, "Efficient Task Replication to Reduce Latency in Cloud Systems", *Allerton Conference on Communication, Control and Computing*, Sep 2015.

[3] D. Wang, G. Joshi, G. Wornell, "Using Straggler Replication to Reduce Latency in Large-scale Parallel Computing", *SIGMETRICS Workshop on Distributed Cloud Computing*, Jun 2015.

[2] G. Joshi, Y. Liu, E. Soljanin, "On the Delay-Storage Trade-off in Coded Distributed Storage Systems", *IEEE Journal on Selected Areas of Communications*, volume 32, number 5, May 2014, arXiv:1305.3945

[1] G. Joshi, Y. Liu, E. Soljanin, "Coding for Fast Content Download", *Allerton Conference on Communication, Control and Computing*, Oct 2012.

Selected Grants and Contracts awarded to date

Overview: The total share of funding (excluding start-up funding) awarded to my group to-date exceeds $4 million, including 7 NSF Awards and industry awards from Google, IBM, Facebook, and Bosch. Selected awards and contracts are listed below.

- *National Science Foundation (NSF) CPS Frontiers*, "Software-Defined Nanosatellite Constellations: The Foundation of Future Space-Based Cyber-physical Systems": $7,000,000, Award period: June 2022-Jun 2027, PI: Brandon Lucia, co-PIs: Vyas Sekar, Swarun Kumar, Gauri Joshi, Zachary Manchester
- *National Science Foundation (NSF)*, "AI Institute for Future Edge Networks and Distributed Intelligence (AI-EDGE)": $1,100,000, Award period: Oct 2021- Sept 2026, led by PI Ness Shroff from Ohio State, CMU PI: Gauri Joshi, and CMU Senior Personnel: Ameet Talwalkar
- *National Science Foundation (NSF)*, "CAREER: Frontiers of Distributed Machine Learning with Communication, Computation and Data Constraints": $650,000, Award period: Mar 2021-Feb 2026. PI: Gauri Joshi
- *National Science Foundation (NSF)*, "SHF: Medium: Collaborative Research: HERMES: On-Device Distributed Machine Learning via Model-Hardware Co-Design": $636,000, Award period: Oct 2021- Oct 2023. CMU PI: Gauri Joshi, co-PI: Virginia Smith, UT Austin PI: Diana Marculescu, co-PI: Radu Marculescu
- *Carnegie Bosch Institute Research Award*, "Scheduling and Queueing Algorithms for Resource-sharing in Federated Learning": $125,000, Award period: July 2021-July 2022. PI: Gauri Joshi, co-PI: Weina Wang
- *National Science Foundation (NSF)*, "CIF: Small: Efficient Sequential Decision-Making and Inference in the Small Data Regime": $500,000, Award period: Oct 2020-Sept 2023. PI: Gauri Joshi, co-PI: Osman Yagan
- *Lawrence Livermore National Lab DoE sub-contract*, "Explainable and Small Data Machine Learning for Accelerating Feedstock Optimization": $359,462, Award period: Oct 2018-Sept 2021. PI: Gauri Joshi
- *National Science Foundation (NSF)*, "CRII: CIF: Unifying Scheduling and Algorithmic Techniques to Speed Up Distributed Stochastic Gradient Descent": $175,000, Award period: Mar 2019-Oct 2020. PI: Gauri Joshi
- *Carnegie Bosch Institute Grant*, "Privacy-Preserving Inference and Decision-Making with IoT Data": $250,000, Award period: Jan 2019-Jun 2021, PI: Osman Yağan, co-PI: Gauri Joshi

Professional Service and Leadership

Editorial and Organizing Roles

Student Activities Chair, ACM SIGMETRICS	2022
Associate Editor, IEEE/ACM Transactions on Networking	2021–
Co-organizer WiOpT workshop on queueing theory meets reinforcement learning	2021
Co-organizer NSF TRIPODS workshop on communication-efficient distributed optimization	2021
Publicity Chair, ACM SIGMETRICS	2021

Co-organizer of the Federated Learning competition at NeurIPS	2020
Organizer, ICML Workshop on Federated Learning, Austria	2020
Associate Editor of the IEEE Open Journal on Signal Processing	2020–
Area Chair, International Conf. on Machine Learning (ICML)	2020
Publications Chair, ACM MobiHoc	2020
Publicity Chair, MLSys Conference	2020
General Co-chair, ICML workshop on redundancy techniques in machine learning	2019
General Co-Chair, Scalable Deep Learning at IPDPS	2019
Machine Learning in Science and Engineering, ECE Track (24 invited talks)	2019

Technical Program Committee

Area Chair, International Conf. on Learning Representations (ICLR)	2020-present
ACM SIGMETRICS	2020-present
International Conference on Machine Learning (ICML)	2019
Information Theory Workshop (ITW)	2020
ACM MobiHoc Conference	2019, 2020
International Symposium on Information Theory (ISIT)	2018, 2019
Machine Learning and Systems (MLSys) Conference	2019, 2020

Invited Keynote and Seminar Talks

Tackling Computational Heterogeneity in Federated Learning

AI-Edge Institute Seminar	Feb 2022
Texas A&M Computer Engineering and Systems Seminar	Feb 2022
FLOW Seminar	Oct 2021

Communication-Efficient Optimization Methods for Federated Learning

Stochastics Networks, Applied Probability and Performance (SNAPP) Seminar	June 2021
Keynote Talk, SIGMETRICS Distributed Cloud Computing workshop	June 2021
Academia Seminar, Microsoft Research Redmond	May 2021
Communication and Signal Processing Seminar, University of Michigan	March 2021

Adaptive Quantization in Federated Learning

London Symposium on Information Theory, United Kingdom	May 2021

Biased Client Selection in Federated Learning

Allerton Conference on Communication, Control and Computing	Oct 2021
ICLR Workshop on Federated Learning	May 2021

Speeding up Distributed SGD via Communication-Efficient Model Aggregation

University of Virginia, VA	Oct 2020
University of Texas at Austin, TX	Mar 2020

ITA Workshop, UCSD, San Diego, CA — Feb 2020
ML Seminar, Univ. of Washington, Seattle WA — Oct 2019
EE Seminar, Harvard SEAS, Cambridge MA — Oct 2019
RLE Seminar, Massachusetts Inst. Tech, Cambridge MA — Oct 2019

Correlated Multi-armed Bandit Algorithms

INFORMS Annual Meeting, Seattle WA — Oct 2019

Fast Distributed Machine Learning with Slow and Stale Updates

SILO Seminar, University of Wisconsin-Masion — Nov 2018
INFORMS Annual Meeting, Phoenix AR — Nov 2018
Google Inc., Pittsburgh PA — May 2018
Georgia Institute of Technology — Mar 2018
Information Theory and Application Workshop, UCSD — Feb 2018

Efficient Redundancy Techniques to Reduce Latency in Cloud Systems

Rutgers University, New Brunswick NJ — Mar 2017
Indian Institute of Technology Bombay, Mumbai — Dec 2016
Tata Institute of Fundamental Research, Mumbai — Dec 2016
Texas A&M University TX — Oct 2016
Carnegie Mellon University, Pittsburgh PA — Mar 2016
University of Southern California, Los Angeles, CA — Mar 2016
IBM T. J. Watson Research Center, Yorktown Heights NY — Feb 2016
Princeton University, EE Departmental Seminar, NJ — Feb 2016
Graduation Day, Information Theory and Applications Workshop UCSD — Feb 2016
IBM Research Student Workshop on Cloud and Data Services, NY — Dec 2015
Rising Stars in EECS, academic workshop for women, MIT — Nov 2015

Using Coding to Reduce Delay in Content Access

Three-lecture series on Coding Theory at Clemson University — Mar 2018
Joint Mathematics Meeting, Atlanta — Jan 2017
Mobile Networked Systems Group, MIT — Sept 2014
Google Headquarters, Mountain View, CA — Aug 2014
University of California Berkeley CA — Aug 2014
DIMACS Workshop on Green Data Storage, Rutgers NJ — Dec 2013

Throughput-Smoothness Trade-offs in Streaming Communication

BIRS Workshop on Mathematical Coding Theory for Streaming — Oct 2015
University of Toronto, Canada — May 2014
Indian Institute of Technology (IIT) Bombay — Aug 2012