

Kumar Kshitij Patel

✉ kpatel@ttic.edu • 🌐 kpatel.ttic.edu

Research Goal: I aim to develop privacy-preserving machine learning systems that learn effectively from sensitive, heterogeneous data. Domains with strict data-sharing limits, for example, healthcare, require collaborative methods that protect individual data while maintaining accuracy, robustness to distribution shifts, and communication efficiency. This motivates my work on federated learning, multi-task learning, and scalable privacy mechanisms, and it supports my broader effort to build safe and trustworthy AI. I also work across disciplines to help create sustainable data-sharing frameworks and equitable data markets for future AI systems.

Education

- **Toyota Technological Institute at Chicago** **Chicago, USA**
Ph.D., Advisors: Nathan Srebro, Lingxiao Wang 2019–2025
Thesis title: What Makes Local Updates Effective: The Role of Data Heterogeneity and Smoothness
- **École Polytechnique Fédérale de Lausanne** **Lausanne, Switzerland**
Two Academic Exchange Semesters 2017–2018
- **Indian Institute of Technology Kanpur** **Kanpur, India**
B.Tech., Computer Science and Engineering 2015–2019

Work Experience

- **Program Fellow, University of California, Berkeley** **Berkeley, CA**
Simon's Institute for the Theory of Computing Jan 2026–May 2026
Incoming program fellow for the research semester on Federated and Collaborative Learning, at the Simon's institute where I will be involved in research efforts along with the program participants.
- **Postdoctoral Associate, Yale University** **New Haven, CT**
Advisor: Dr. Dan Spielman, Institute for Foundations of Data Science Jul 2025–Jun 2028
I am working across the campus on various applications of Federated Learning in healthcare and science. I am also working on fundamental theoretical questions in data valuation and multi-distribution learning.
- **Research Intern, Sony AI** **Tokyo, JP**
Advisors: Dr. Nidham Gazagnadou, Dr. Lingjuan Lyu, PPML Team Jun 2023–Sept 2023
Worked on personalization in federated learning (FL) and provided new convergence guarantees for personalized local SGD, highlighting its advantages over other baselines, theoretically and empirically.
- **Applied Scientist Intern, Amazon AWS** **Seattle, USA**
Advisors: Dr. Srinivasan Sengamedu, Dr. Omer Tripp, Codeguru Team Jun 2020–Sept 2020
Worked on using deep language models: BERT and GPT-2, for detecting leakage of sensitive information in Java code, in symbiosis with program analysis tools. My project significantly advanced the adoption of deep learning to taint analysis in CodeGuru and got assimilated into their customer-facing application.

Awards and Achievements

- **Distinguished Paper Award** (equivalent to a best paper award) at IJCAI'24.
- **Best Paper Honorable Mention Award** at the Federated Learning workshop at ICML'23.
- **Top Reviewer Award** at AISTATS'23, ICML'22, NeurIPS'22,'25.
- **Best Poster Presentation Award** at TTIC student workshop 2023.
- **Best Talk Award** at TTIC student workshop 2022.
- Recipient of **Honda Young Engineer and Scientist's (Y-E-S) Fellowship 2017**, awarded to only 14 undergraduates in India for their academic and research work.

- **Academic Excellence Award**, 2015 at IIT Kanpur.
- **Represented India** as a part of the Youth Delegation to Nepal organized by the Govt. of India.
- **All India Rank 200** in JEE-Adv. 2015, and **99.9 %**-tile in JEE-Mains 2015 out of 1.3M students.
- Among **top 1%** in national standard tests; qualified for **Indian National Chemistry Olympiad**.

Invited Talks and Tutorials

- Delivered the UAI'23 tutorial on **Online Optimization for Federated Learning**. [Talk]
- Delivered talks on optimization for federated learning at the **Google Federated Learning Seminar** and the **Federated Learning One World Seminar (FLOW)**. [Talk]
- Delivered talks on online Combinatorial optimization with group fairness at the sessions on **New Directions in Revenue Management** at INFORMS'23, and **Operations Management Problems with Social Interactions** at INFORMS'24.

Publications

- Golrezaei N.*, Niazadeh R.*, **Patel K. K.*** & Susan F.*, Online combinatorial optimization with group fairness constraints. **Distinguished Paper**, IJCAI'24. *Alphabetical ordering [PDF]
- **Patel K. K.***, Glasgow M.*, Wang L., Joshi N. & Srebro N., On the still unreasonable effectiveness of federated averaging for heterogeneous distributed learning. FL Workshop, ICML'23. **Best paper, honorable mention**. *Equal Contribution [PDF]
- **Patel K. K.**, Zindari A., Wang L. & Stich S.U., Revisiting Consensus Error: A Fine-grained Analysis of Local SGD under Second-order Data Heterogeneity, NeurIPS'25. [PDF]
- Kundu R., Shi X., **Patel K. K.**, Ohno-Machado L., Salvatore M., Song P.X.K. & Mukherjee B., Privacy-Enhancing Sequential Learning under Heterogeneous Selection Bias in Multi-Site EHR Data. Under review. [PDF]
- **Patel K. K.**, Zhang W. & Wang L., Personalized Federated Training of Diffusion Models with Privacy Guarantees. Under review. [PDF]
- **Patel K. K.**, Woodworth B., Gazagnadou N., Wang L., & Lyu L., Personalization Mitigates the Perils of Local SGD for Heterogeneous Distributed Learning. Under review. [Talk]
- Manoj N. & **Patel K. K.**, A Second-Order Algorithm for Empirical Group Distributionally Robust Regression. OPT'24 Workshop, NeurIPS'24. Extended version, under review.
- **Patel K. K.**, Glasgow M., Zindari A., Wang L., Stich S.U., Cheng Z., Joshi N. & Srebro N., The Limits and Potentials of Local SGD for Distributed Heterogeneous Learning with Intermittent Communication., COLT'24. [Talk][PDF]
- Kabra A.* & **Patel K. K.***, The Limitations of Model Retraining in the Face of Performativity. Humans, Algorithmic Decision-Making and Society: Modeling Interactions and Impact Workshop, ICML, 2024. [PDF]
- Yunis D., **Patel K. K.**, Wheeler S., Savarese P., Vardi G., Frankle J., Livescu K., Maire M., Walter M., Approaching Deep Learning through the Spectral Dynamics of Weights. Under review, ICLR'25. [PDF]
- Wang L., Zou D., **Patel K. K.**, Wu J. & Srebro N., Private Overparameterized Linear Regression Without Suffering in High Dimensions. Under review. [PDF]
- Han M.*, Shao H.*, **Patel K. K.*** & Wang L.*, On the effect of defection in federated learning and how to prevent it. Under review, JMLR. *Alphabetical ordering [PDF]
- Wang L., Zhou X., **Patel K. K.**, Tang L. & Saha A., Efficient Private Federated Non-Convex Optimization With Shuffled Model. Privacy Regulation and Protection in Machine Learning Workshop. ICLR'24. [PDF]
- **Patel K. K.**, Wang L., Saha A. & Srebro N., Federated online and bandit convex optimization. ICML'23 and OPT'22 Workshop, NeurIPS'22. [PDF]

- Yunis D., **Patel K. K.**, Savarese P., Vardi G., Livescu K. & Walter M., On convexity and linear mode connectivity. OPT'22 Workshop, NeurIPS'22. [PDF]
- **Patel K. K.**, Wang L., Woodworth B., Bullins B. & Srebro N., Towards optimal communication complexity in distributed non-convex optimization. NeurIPS'22 [Talk][PDF]
- Bullins B.*, **Patel K. K.***, Shamir O.*, Srebro N.* & Woodworth B.*, A stochastic newton algorithm for distributed convex optimization. NeurIPS'21. *Alphabetical ordering [PDF]
- Woodworth B., **Patel K. K.** & Srebro N., Minibatch vs local SGD for heterogeneous distributed learning. NeurIPS'20. [Talk][PDF]
- Woodworth B., **Patel K. K.**, Stich S.U., Dai Z., Bullins B., McMahan B., Shamir O. & Srebro N., Is local SGD better than minibatch SGD? ICML'20. [Talk] [PDF]
- Lin T., Stich S.U., **Patel K. K.** & Jaggi M., Don't use large mini-batches, use local SGD. ICLR'20. [Talk][PDF]
- **Patel K. K.*** & Dieuleveut A.*, Communication trade-offs for synchronized distributed SGD with large step-size. NeurIPS'19. *Equal Contribution [PDF]
- Kapoor S., **Patel K. K.** & Kar, P., Corruption-tolerant bandit learning. Machine Learning Journal, Springer. [Journal Link][PDF]

Service and Professional Activities

Organization

- Co-organizer of the Simon's Institute workshop on **Learning from Heterogeneous Sources**, as a part of the research semester on Federated and Collaborative Learning. [Link]
- Co-organizer of TTIC Summer 2025 workshop on **Incentives for Data Sharing and Collaborative Learning**, focusing on various issues in data markets in AI. [Link]
- Co-organizer of the **New Frontiers in Federated Learning** Workshop at TTIC (Summer'23). Our workshop resulted in several open problems and new collaborations. [Link]
- **AI hub ambassador** at IJCAI'24; interviewed two keynote speakers at the conference.
- Co-organizer of **Machine Learning and Optimization reading group** at TTIC, Winter'21-'22.
- Co-organizer of the TTIC Student Workshop 2021. [Link]
- Session chair for the session on **AI for drug-discovery** at ICML'22.
- Co-started the TTIC/UChicago Student Theory Seminar.
- Volunteer for IJCAI'24, NeurIPS'22, ICML'20,'22, and ICLR'20.

Teaching

- Taught a mini-course on algorithmic game theory at TTIJ, Nagoya, Japan (Fall'24).
- Teaching assistant for the Mathematical Toolkit course at TTIC (Spring'23).
- Teaching assistant for the Convex Optimization course at TTIC (Winter'22, and Winter'24).
- Teaching assistant and co-organizer for Research at TTIC Colloquium (Fall'20 - Winter'21).

Reviewing

- Journals: JMLR, TMLR, Springer Machine Learning Journal.
- Conferences: STOC'21, ICML'21,22,24, NeurIPS'21,22,23,24, ICLR'22,23,24,25, AISTATS'22,23,24.
- Workshops: OPT'22,'23 @ NeurIPS'22,'23, FL workshop at NeurIPS'23 and IJCAI'23.

Committees

- Student member of the Sexual Misconduct Policy Committee at TTIC (2021).

Mentorship

- Peer Mentor to a first-year PhD student at TTIC (2020).
- Mentor to first-year computer science undergraduates for a machine learning project (2016).
- Student guide to six students at the counseling service at IIT Kanpur (2016).

Participation

- Participant in the Mathematics of Deep Learning collaboration led by Simons Foundation.
- Attended the Machine Learning Summer School 2020 organized by Max Plank Institute for Intelligent Systems, Tübingen, Germany.

Skills

Relevant Coursework

- **ML Theory:** convex optimization, statistical and computational learning theory, online learning and optimization, bayesian machine learning, information theory and coding.
- **ML applications:** machine learning and systems, deep learning, computer vision, unsupervised learning and speech, natural language processing.
- **Mathematics and statistics:** real analysis, measure theory, numerical linear algebra, time series, stochastic processes.
- **Computer science:** algorithms, theory of computation, operating systems, database design, compiler design.

Programming

- **Languages:** Python, C, R, \LaTeX , HTML-CSS, C++, Matlab, SQLite, Assembly.
- **Packages:** PyTorch, Keras, Scikit, Gensim, NLTK, XGBoost, CVXPY.