Jiaming Yang

jiamyang@umich.edu · (+1) 734-834-4148 · Ann Arbor, MI · Google Scholar · Homepage · LinkedIn

Summary

Final-year Computer Science Ph.D. candidate at the University of Michigan (graduating May 2027), working in **randomized numerical linear algebra**: developing provably efficient algorithms that exploit spectral structure to accelerate linear solvers, kernel methods, second-order optimization and beyond. Publications at **STOC, SODA, ICML, ICLR, JMLR** and **SIMAX**. Seeking research roles at the intersection of randomized algorithms, optimization and efficient machine learning.

Education

University of Michigan, Ann Arbor Aug. 2022 – May 2027 (expected)
Ph.D. Candidate, Computer Science and Engineering. GPA 4.0/4.0. Advisor: Prof. Michał Dereziński.

Peking University Sep. 2018 – Jul. 2022
B.S. in Statistics, School of Mathematical Sciences.

Research Experience

University of Michigan, Ann Arbor Aug. 2022 – Present
Graduate Student Research Assistant, advised by Prof. Michał Dereziński

- **Fast linear-system solvers** (STOC 2024, SODA 2025). Designed a randomized block coordinate descent solver for dense $n \times n$ systems running in $\tilde{O}(n^2 + nk^{\omega-1})$ time (k = number of large outlying singular values; ω = fast matrix multiplication exponent; ignoring log factors), improving on both direct solvers and preconditioned CG for flat-tailed spectra; proved sharp convergence via DPP sampling, majorization theory and Markov chain coupling. Developed *multi-level sketched Nyström preconditioning*, giving the first $\tilde{O}(n^2 + d_\lambda^\omega)$-time solver for regularized PSD systems (d_λ = effective dimension) and faster Schatten-norm approximation.
- **Randomized Kaczmarz with beyond-Krylov convergence** (SIMAX 2025). Introduced *Kaczmarz++*, a block Kaczmarz solver with adaptive momentum, Tikhonov-regularized projections and block memoization; it captures large singular-value outliers provably faster than popular Krylov methods. Its PSD variant *CD++* is competitive with PCG/GMRES in arithmetic cost on benchmarks.
- **Accelerated second-order optimization** (Acc-RSRN, under review; with RIKEN AIP / U. Tokyo). Developed an accelerated randomized subspace regularized Newton method with a safeguarded global-to-local switch, proved conditional local acceleration guarantees, and validated it on kernel logistic regression and CNN training.
- **Scalable kernel ridge regression** (JMLR 2026). Co-developed *ASkotch*, an accelerated Nyström-preconditioned block coordinate descent solver for full KRR with linear convergence guarantees; outperforms state-of-the-art full and inducing-points KRR solvers on 23 large-scale regression and classification tasks.
- **Efficient GNN training** (NeurIPS 2024 M3L Workshop). Co-designed *HERTA*, which converges to the original unfolded GNN model's optimum and has nearly-linear-time training guarantees under spectral and dimensional assumptions, together with a new spectral sparsifier for normalized, regularized graph Laplacians.

Remote Research Collaboration — *supervised by Dr. Zhao Song* Jun. 2021 – Jun. 2022

- **Pixelated Butterfly** (ICLR 2022, **spotlight**). Sparse training via butterfly + low-rank weight structure: sparse ViT, MLP-Mixer and GPT-2 models train up to 2.5× faster than dense baselines on ImageNet and WikiText-103 with no accuracy loss.
- **Federated adversarial learning** (ICML 2023). Established convergence guarantees for federated adversarial training of over-parameterized ReLU networks, handling the nested min–max structure and multiple local updates.

Columbia University (Remote) — *supervised by Prof. Josh Alman* Jun. 2021 – Dec. 2021

- Designed and analyzed a recursive algorithm for Kronecker product regression, exploiting the problem's structured matrix representation to compute exact solutions; proved its optimality and applied the resulting method to Kronecker low-rank approximation.

Publications

(α-β) = alphabetical author order

[1] **Randomized Kaczmarz Methods with Beyond-Krylov Convergence.** M. Dereziński, D. Needell, E. Rebrova, **J. Yang** *(α-β)*. *SIMAX 2025*.
[2] **Have ASkotch: A Neat Solution for Large-scale Kernel Ridge Regression.** P. Rathore, Z. Frangella, **J. Yang**, M. Dereziński, M. Udell. *JMLR 2026*.
[3] **Faster Linear Systems and Matrix Norm Approximation via Multi-level Sketched Preconditioning.** M. Dereziński, C. Musco, **J. Yang** *(α-β)*. *SODA 2025*.
[4] **Solving Dense Linear Systems Faster than via Preconditioning.** M. Dereziński, **J. Yang** *(α-β)*. *STOC 2024*.
[5] **HERTA: A High-Efficiency and Rigorous Training Algorithm for Unfolded Graph Neural Networks.** Y. Yang, **J. Yang**, W. Hu, M. Dereziński. *NeurIPS 2024 M3L Workshop*.
[6] **Federated Adversarial Learning: A Framework with Convergence Analysis.** X. Li, Z. Song, **J. Yang** *(α-β)*. *ICML 2023*.
[7] **Pixelated Butterfly: Simple and Efficient Sparse Training for Neural Network Models.** B. Chen, T. Dao, K. Liang, **J. Yang**, Z. Song, A. Rudra, C. Ré. *ICLR 2022*, **spotlight**.
[P1] **Acc-RSRN: Accelerated Randomized Subspace Regularized Newton. J. Yang**, P.-L. Poirion, A. Takeda, M. Dereziński. Under review.

Skills, Teaching, Service & Honors

Programming: Python (NumPy/SciPy, PyTorch), MATLAB, R, LaTeX.
Expertise: randomized numerical linear algebra, stochastic optimization, kernel methods, machine learning theory, random matrix theory.
Languages: Chinese (native), English (fluent). GRE: V 159, Q 170, AW 4.0; TOEFL 107.
Teaching: Graduate Student Instructor, EECS 376: Foundations of Computer Science, University of Michigan (Winter 2026).
Reviewer: NeurIPS 2024–2025, ICML 2025–2026, ICLR 2027, AISTATS 2025–2026; UMich CSE Ph.D. admissions.
Honors: Department finalist, Two Sigma Ph.D. Fellowship (2025); student travel grants from STOC 2024, SODA 2025, RMT+NLA II 2025 and Rackham Graduate School (2023–2026). Talks at INFORMS 2026, ICCOPT 2025 and SODA 2025.