# Zachary Frangella

zfrangella@alumni.stanford.edu

zfrangella.github.io | zfrangella


## Education

**Stanford University** 2022–25
**Ph.D.**, Management Science & Engineering

**Cornell University**, 2019-2022
**M.Sc.**, Applied Mathematics

**Rensselaer Polytechnic Institute**, 2015-2018
**B.Sc.**, Mathematics, *Magna Cum Laude*

## Awards & Honors

**Dean's Honors List**, Rensselaer Polytechnic Institute 2015-2018

**Max Hirsch Prize**, Rensselaer Polytechnic Institute 2019
*This Prize is awarded to a Senior in the Department of Mathematical Sciences who has demonstrated outstanding ability in his or her academic work and also gives promise of outstanding success in a career in mathematical sciences.*

## Research Experience

**Stanford University**, Stanford, CA 2022–2025
*Research Assistant*

- Designed preconditioned stochastic gradient algorithms leveraging randomized numerical linear algebra to accelerate large-scale machine learning, with applications to GLMs and neural network training.
- Introduced RPCholesky and KRILL preconditioners for kernel ridge regression (KRR), both adopted in RandLAPACK for large-scale randomized numerical linear algebra.
- Developed block-access algorithms for solving linear systems involving massive kernel matrices ($n > 10^6$), enabling scalable KRR and Gaussian Process inference on datasets with $n \sim 10^8$ using a single GPU.
- Analyzed the optimization landscape of physics-informed neural networks (PINNs), explaining training challenges; proposed NysNewton-CG to help address these challenges, which has now been integrated into the DeepXDE library for scientific machine learning.
- Implemented and benchmarked algorithms in PyTorch, JAX, Python, and Matlab.

**Cornell University**, Ithaca, NY 2019–2022
*Research Assistant*

- Investigated algorithms for large-scale linear systems and convex composite optimization, with applications to GLMs, portfolio optimization, (kernel) ridge regression, and SVMs.
- Analyzed the Leave-one-out Cross-Validation (LOOCV) objective for ridge regression; proved quasiconvexity under mild assumptions, enabling global optimization for model selection despite general non-convexity.
- Developed Nyström Preconditioned Conjugate Gradient (Nyström PCG), adopted in RandLAPACK (a LAPACK extension for randomized numerical linear algebra), for solving large-scale linear systems with low-rank structure.
- Proposed NysADMM, combining Nyström PCG with linearized ADMM, achieving state-of-the-art performance over solvers such as GLMNet, LIBSVM, and SAGA.
- Implemented and benchmarked algorithms in Python and Matlab for empirical validation.

**Rensselaer Polytechnic Institute**, Troy, NY Summer 2018
*Undergraduate Researcher*

- Developed finite element schemes for parabolic PDEs with non-smooth initial data.
- Investigated Rannacher time-stepping for classic FEM schemes and extended methods to Discontinuous Galerkin formulations.
- Implemented prototypes in MATLAB and conducted numerical experiments.

## Publications (Google Scholar)

† → Equal contribution

### Journal Articles

J1. **Frangella, Zachary**, Tropp, J. A. & Udell, M. Randomized Nyström Preconditioning. *SIAM Journal on Matrix Analysis and Applications* **44,** 718–752 (2023).

J2. **Frangella, Zachary**, Rathore, P., Zhao, S. & Udell, M. Promise: Preconditioned Stochastic Optimization Methods by Incorporating Scalable Curvature Estimates. *Journal of Machine Learning Research* **25,** 1–57 (2024).

J3. **Frangella, Zachary**, Rathore, P., Zhao, S. & Udell, M. SketchySGD: Reliable Stochastic Optimization via Randomized Curvature Estimates. *SIAM Journal on Mathematics of Data Science* **6,** 1173–1204 (2024).

### Peer-reviewed Conference Proceedings

C1. Stephenson, W., **Frangella, Zachary**, Udell, M. & Broderick, T. *Can We Globally Optimize Cross-Validation Loss? Quasiconvexity in Ridge Regression* in *Advances in Neural Information Processing Systems* (2021), 24352–24364.

C2. Zhao†, *S.*, **Frangella, Zachary**† & Udell, M. *NysADMM: faster composite convex optimization via low-rank approximation* in *International Conference on Machine Learning* (2022), 26824–26840.

C3. Rathore, P., Lei, W., **Frangella, Zachary**, Lu, L. & Udell, M. *Challenges in Training PINNs: A Loss Landscape Perspective* in *International Conference on Machine Learning* (2024), 42159–42191. **Oral, Top 1.5% of submissions**.

C4. Feng, M., **Frangella, Zachary** & Pilanci, M. *CRONOS: Enhancing Deep Learning with Scalable GPU Accelerated Convex Neural Networks* in *Advances in Neural Information Processing Systems* (2024).

### In the pipeline

P1. Diamandis, T., **Frangella, Zachary**, Zhao, S., Stellato, B. & Udell, M. GeNIOS: an (almost) second-order operator-splitting solver for large-scale convex optimization. *Submitted* (2023).

P2. **Frangella, Zachary**, Diamandis, T., Stellato, B. & Udell, M. On the (linear) convergence of Generalized Newton Inexact ADMM. *Submitted* (2023).

P3. Díaz[†], *M.*, Epperly[†], *E.N.*, **Frangella, Zachary**[†], Tropp, J. A. & Webber[†], *R.J.*. Robust, randomized preconditioning for kernel ridge regression. *Submitted* (2024).

P4. Fazliani, S., Frangella, Z. & Udell, M. Enhancing Physics-Informed Neural Networks Through Feature Engineering. *Submitted* (2025).

P5. Rathore, P., **Frangella, Zachary**, Yang, J., Dereziński, M. & Udell, M. Have ASkotch: A Neat Solution for Large-scale Kernel Ridge Regression. *Submitted* (2025).

P6. Rathore, P. *et al.* Turbocharging Gaussian Process Inference with Approximate Sketch-and-Project. *Submitted* (2025).

P7. Sun, J., **Frangella, Zachary** & Udell, M. SAPPHIRE: Preconditioned Stochastic Variance Reduction for Faster Large-Scale Statistical Learning. *Submitted* (2025).

## Talks and Posters

**ICME Seminar on Linear Algebra and Optimization**, Stanford 2024
*Faster Convex Optimization via Randomized Numerical Linear Algebra*

**MILA Tensor Network Reading Group**, Montreal *2024*
*SketchySGD: Reliable Stochastic Optimization via Randomized Curvature Estimates*

**International Congress of Industrial and Applied Mathematics**, Tokyo 2023
*Faster Convex Optimization via Randomized Numerical Linear Algebra*

**ICME Xpo**, Stanford *2023*
*SketchySGD: Reliable Stochastic Optimization via Randomized Curvature Estimates*

**SIAM Conference on Optimization**, Seattle 2023
*Low-rank Approximation for Faster Optimization*

**International Conference on Machine Learning**, Baltimore 2022
*NysADMM: Faster Composite Convex Optimization via Low-Rank Approximation*

**SciMLCon**, Virtual 2022
*Speeding up A\b with Randomized Preconditioners*

## Technical Skills

**Programming Languages & Frameworks :**

- *Proficient:* Python, Jax, PyTorch, NumPy, Cython, MATLAB, LaTeX
- *Familiar:* Pandas, C/C++, Julia

## Academic Mentoring

### Graduate

Jingruo Sun, M.Sc. Stanford MS&E — *Summer 2024 – Spring 2025*
*Stochastic Proximal Preconditioning*

Weimu Lei, M.Sc. Stanford ICME — *Fall 2023 – Fall 2024*
*Optimization for PINNs, GPU Acclerated Optimization Software*

### Undergraduate

Yingxi Li, B.Sc. Cornell ORIE — *Spring 2022*
*Faster Sparse Composite Optimization via SAFE coordinate selection*

## Teaching

### Rensselaer Polytechnic Institute

**MATH 2400: Differential Equations** — *Spring 2019*
*Teaching Assistant*

### Cornell University

**MATH 1910: Calculus for Engineers** — *Fall 2019, 2020*
*Teaching Assistant*

**MATH 1920: Multivariable Calculus for Engineers** — *Spring 2021*
*Teaching Assistant*

**MATH 2930: Differential Equations for Engineers** — *Spring 2020, 2022*
*Teaching Assistant*

### Stanford University

**CME 307/MS&E 311: Optimization** — *Spring 2023, Winter 2024, Fall 2024*
*Teaching Assistant, Guest Lecturer*

## Academic Service

### Organized Sessions

**ICME Seminar on Linear Algebra and Optimization**, Stanford *Winter 2024*
*with Michael Saunders*

**INFORMS: Advances in Optimization for Machine Learning**, Seattle *2024*
*with Pratik Rathore and Madeleine Udell*

### Conference Reviewer

Advances in Neural Information Processing Systems (NeurIPS)

International Conference on Machine Learning (ICML)

International Conference on Learning Representations (ICLR)

International Conference on Artificial Intelligence and Statistics (AISTATS)

### Journal Reviewer

Journal of Machine Learning Research

SIAM Journal on Mathematics of Data Science

SIAM Journal on Matrix Analysis

Automatica

Optimization and Engineering

Statistics & Computing

## Relevant Courses

Numerical Computing, Numerical Ordinary Differential Equations, Numerical Partial Differential Equations, Matrix Computations, Data-sparse Matrix Computations, Numerical Methods for Data Science, Reinforcement Learning, Probability Theory, Quantum Physics I-II

---

Last updated: September 17, 2025