Department of Mathematics, University of Utah
Topics in Randomized Numerical Linear Algebra
MTH7870 – Section 001 – Spring 2026

End-of-semester Presentations

A required component of this course is the in-class presentation of either a paper, or a summary of a project that you are investigating or have completed. The topic of your presentation must intersect with the content of the course in some (possibly speculative) way. Potential types of presentations are:

- Overview of a research paper (see below for suggestions). This includes a high-level summary of the problem the paper is attempting to solve, the main ideas of the approach, a description of the numerical results (if appropriate), and a discussion of the advantages/disadvantages and open problems or future directions from this work. You need not be an expert in the content of the paper, and you need not fully understand all technical details, but you absolutely should be well-versed enough to present the paper and answer some basic questions. You are encouraged to read related references as well to gain a fuller picture of the landscape of the paper's focus.
- Presentation of a project you have pursued or are actively pursuing. "Project" is loosely defined – it can be a numerical implementation of an algorithm to solve a problem, an empirical investigation of several algorithms, a presentation of theory that you have learned, etc. The major goal is to have a well-crafted message and presentation.

In terms of topics, I am quite flexible – anything RandNLA "related" is fine. See below for some paper suggestions. (They are only suggestions, you may choose other options.)
In general I am happy to meet to discuss anything related to these projects, including deciding on a topic, settling on what to present, how much detail to present, etc. Please do reach out if you'd like to discuss at any point.

Timeline

Here is a proposed timeline for these presentations:

- Now until March 20 – Decide on a topic/paper <u>and</u> finalize a presentation date.
- By March 20 – Communicate with me your decision on topic and your preferred presentation date. We can discuss over email, virtually over Zoom, or in-person.
- March 27 (or before) – I will announce the presentation lineup and topics
- April 2 - April 21 – In-class presentations. I plan on three per class meeting.

You can interpret the above to mean that **March 20 is the deadline for deciding on a project topic + presentation date**. I am aiming to have three presentation per class meeting, which means that I will fill people into slots on a first-come first-served basis.

Presentation "Guidelines" and Preparation

I am leery of providing rules on presentations since I don't want to impose too many constraints. But here is a general outline of what I expect:

- Presentation length: 20-25 minutes. 20 minutes is probably the minimum amount of time required to adequately give background on the project. As a general rule of thumb, most folks generally underestimate the amount of time they need to speak.
- Presentation format: you are welcome to use electronic slides, or present on the whiteboard. (We do not have a chalkboard in our classroom.) I would recommend slides if you're going to show numerical results/pictures. The class has an HDMI input, and I will always have a USB-C adapter + HDMI cable. If you need other adapters, please let me know as soon as you can so I can arrange to have them available. If you prefer to present virtually, please let me know well ahead of time so we can iron out logistics for that.
- Group presentations: You are welcome to present in groups, I'd suggest no more than 2 individuals per group; to work in groups of 3 or more you'd have to present me with a compelling argument. Presentation length for group projects will be adjusted to ensure each person has adequate time to speak.
- I would encourage you to share with me either a draft of your slides or an outline of your presentation before your talk to give me notice about what you'll talk about. However, this is not necessary.

Potential papers/topics

The ultimate goal of this is for you to present on a topic that piques your interest. Below I list some topics that either I am familiar with, or would be interesting for me, but you should prioritize your own interest in choosing your topics. In particular, <u>do not</u> consider the list below as a comprehensive collection of materials: There exist appropriate references that I have not listed below, and there exist entire sub-topics that I have omitted.
I can provide you with PDF soft copies of any of the following papers if you have trouble accessing them.

A <u>non-exhaustive</u> list of potential topics

- [3], *Blendenpik: Supercharging LAPACK's Least-Squares Solver*
- [28], *A fast randomized algorithm for overdetermined linear least-squares regression*
- [35], *A fast randomized algorithm for the approximation of matrices*
- [8], *An Improved Approximation Algorithm for the Column Subset Selection Problem*
- [30], *A Randomized Kaczmarz Algorithm with Exponential Convergence*
- [12], *Sublinear Randomized Algorithms for Skeleton Decompositions*
- [22], *Hutch++: Optimal stochastic trace estimation*
- [13], *On Randomized Trace Estimates for Indefinite Matrices with an Application to Determinants*
- [25], *Improved variants of the hutch++ algorithm for trace estimation*
- [17], *XTrace: Making the most of every sample in stochastic trace estimation*

- [10], *Faster Linear Algebra Algorithms with Structured Random Matrices*
- [31], *Comparison theorems for the minimum eigenvalue of a random positive-semidefinite matrix*
- [11], *Randomly pivoted Cholesky: Practical approximation of a kernel matrix with few entry evaluations*
- [24], *Fast and accurate randomized algorithms for linear systems and eigenvalue problems*
- [19], *Matrix Concentration for Products*
- [16], *Efficient error and variance estimation for randomized matrix computations*
- [20], *Approximate Gaussian Elimination for Laplacians - Fast, Sparse, and Simple*
- [6], *Pseudospectral Shattering, the Sign Function, and Diagonalization in Nearly Matrix Multiplication Time*
- [1], *Fast Direct Methods for Gaussian Processes*
- [2], *Living on the edge: Phase transitions in convex programs with random data*
- [4], *Sharp analysis of low-rank kernel matrix approximations*
- [29], *Improved Approximation Algorithms for Large Matrices via Random Projections*
- [14], *Asymptotics for Sketching in Least Squares Regression*
- [15], *Randomized QR with Column Pivoting*
- [18], *Accelerated stochastic matrix inversion: general theory and speeding up BFGS rules for faster second-order optimization*
- [21], *Fast Randomized Iteration: Diffusion Monte Carlo through the Lens of Numerical Linear Algebra*
- [23], *Recursive Sampling for the Nystrom Method*
- [27], *Stochastic Reformulations of Linear Systems: Algorithms and Convergence Theory*
- [26], *Iterative Hessian Sketch: Fast and Accurate Solution Approximation for Constrained Least-Squares*
- [32], *Concentration inequalities for random tensors*
- [33], *Efficient algorithms for cur and interpolative matrix decompositions*
- [34], *Scalable Kernel K-Means Clustering with Nystrom Approximation: Relative-Error Bounds*
- [5], *Randomized Cholesky QR factorizations*
- [7], *A Practical Randomized CP Tensor Decomposition*
- [9], *Norm and Trace Estimation with Random Rank-one Vectors*

References

[1] S. Ambikasaran, D. Foreman-Mackey, L. Greengard, D. W. Hogg, and M. O'Neil. "Fast Direct Methods for Gaussian Processes". In: *IEEE Trans. Pattern Anal. Mach. Intell.* 38.2 (Feb. 2016), pp. 252–265.

[2] D. Amelunxen, M. Lotz, M. B. McCoy, and J. A. Tropp. *Living on the edge: Phase transitions in convex programs with random data.* arXiv:1303.6672 [cs]. Apr. 2014.

[3] H. Avron, P. Maymounkov, and S. Toledo. "Blendenpik: Supercharging LAPACK's Least-Squares Solver". In: *SIAM Journal on Scientific Computing* 32.3 (Jan. 2010), pp. 1217–1236.

[4] F. Bach. “Sharp analysis of low-rank kernel matrix approximations”. In: *Proceedings of the 26th Annual Conference on Learning Theory*. PMLR, June 2013, pp. 185–209.

[5] O. Balabanov. *Randomized Cholesky QR factorizations*. arXiv:2210.09953 [math]. Oct. 2022.

[6] J. Banks, J. Garza-Vargas, A. Kulkarni, and N. Srivastava. “Pseudospectral Shattering, the Sign Function, and Diagonalization in Nearly Matrix Multiplication Time”. In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. ISSN: 2575-8454. Nov. 2020, pp. 529–540.

[7] C. Battaglino, G. Ballard, and T. G. Kolda. “A Practical Randomized CP Tensor Decomposition”. In: *SIAM Journal on Matrix Analysis and Applications* 39.2 (2018), pp. 876–901.

[8] C. Boutsidis, M. W. Mahoney, and P. Drineas. “An Improved Approximation Algorithm for the Column Subset Selection Problem”. In: (Dec. 2008).

[9] Z. Bujanovic and D. Kressner. “Norm and Trace Estimation with Random Rank-one Vectors”. In: *SIAM Journal on Matrix Analysis and Applications* 42.1 (2021), pp. 202–223.

[10] C. Camaño, E. N. Epperly, R. A. Meyer, and J. A. Tropp. *Faster Linear Algebra Algorithms with Structured Random Matrices*. arXiv:2508.21189 [cs]. Aug. 2025.

[11] Y. Chen, E. N. Epperly, J. A. Tropp, and R. J. Webber. “Randomly pivoted Cholesky: Practical approximation of a kernel matrix with few entry evaluations”. In: *Communications on Pure and Applied Mathematics* 78.5 (2025). arxiv: 2207.06503, pp. 995–1041.

[12] J. Chiu and L. Demanet. “Sublinear Randomized Algorithms for Skeleton Decompositions”. In: *SIAM Journal on Matrix Analysis and Applications* 34.3 (Jan. 2013), pp. 1361–1383.

[13] A. Cortinovis and D. Kressner. “On Randomized Trace Estimates for Indefinite Matrices with an Application to Determinants”. In: *Foundations of Computational Mathematics* 22.3 (June 2022), pp. 875–903.

[14] E. Dobriban and S. Liu. “Asymptotics for Sketching in Least Squares Regression”. In: *Advances in Neural Information Processing Systems*. Vol. 32. Curran Associates, Inc., 2019.

[15] J. A. Duersch and M. Gu. “Randomized QR with Column Pivoting”. In: *SIAM Journal on Scientific Computing* 39.4 (2017), pp. C263–C291.

[16] E. N. Epperly and J. A. Tropp. “Efficient error and variance estimation for randomized matrix computations”. In: *SIAM Journal on Scientific Computing* 46.1 (2024). arxiv: 2207.06342, A508–A528.

[17] E. N. Epperly, J. A. Tropp, and R. J. Webber. “XTrace: Making the most of every sample in stochastic trace estimation”. In: *SIAM Journal on Matrix Analysis and Applications* 45.1 (2024). arxiv: 2301.07825, pp. 1–23.

[18] R. Gower, F. Hanzely, P. Richtárik, and S. U. Stich. “Accelerated stochastic matrix inversion: general theory and speeding up BFGS rules for faster second-order optimization”. In: *Advances in Neural Information Processing Systems* 31 (2018).

[19] D. Huang, J. Niles-Weed, J. A. Tropp, and R. Ward. “Matrix Concentration for Products”. In: *Foundations of Computational Mathematics* 22.6 (Dec. 2022), pp. 1767–1799.

[20] R. Kyng and S. Sachdeva. “Approximate Gaussian Elimination for Laplacians - Fast, Sparse, and Simple”. In: IEEE Computer Society, Oct. 2016, pp. 573–582.

[21] L.-H. Lim and J. Weare. “Fast Randomized Iteration: Diffusion Monte Carlo through the Lens of Numerical Linear Algebra”. In: *SIAM Review* 59.3 (2017), pp. 547–587.

[22] R. A. Meyer, C. Musco, C. Musco, and D. P. Woodruff. “Hutch++: Optimal stochastic trace estimation”. In: *2021 symposium on simplicity in algorithms (SOSA)*. arxiv: 2010.09649, pp. 142–155.

[23] C. Musco and C. Musco. “Recursive Sampling for the Nystrom Method”. In: *Advances in Neural Information Processing Systems*. Vol. 30. Curran Associates, Inc., 2017.

[24] Y. Nakatsukasa and J. A. Tropp. “Fast and accurate randomized algorithms for linear systems and eigenvalue problems”. In: *SIAM Journal on Matrix Analysis and Applications* 45.2 (2024). arxiv: 2111.00113, pp. 1183–1214.

[25] D. Persson, A. Cortinovis, and D. Kressner. "Improved variants of the hutch++ algorithm for trace estimation". In: *SIAM Journal on Matrix Analysis and Applications* 43.3 (2022), pp. 1162–1185.

[26] M. Pilanci and M. J. Wainwright. "Iterative Hessian Sketch: Fast and Accurate Solution Approximation for Constrained Least-Squares". In: *Journal of Machine Learning Research* 17.53 (2016), pp. 1–38.

[27] P. Richtárik and M. Takáč. "Stochastic Reformulations of Linear Systems: Algorithms and Convergence Theory". In: *SIAM Journal on Matrix Analysis and Applications* 41.2 (2020), pp. 487–524.

[28] V. Rokhlin and M. Tygert. "A fast randomized algorithm for overdetermined linear least-squares regression". In: *Proceedings of the National Academy of Sciences* 105.36 (Sept. 2008), pp. 13212–13217.

[29] T. Sarlos. "Improved Approximation Algorithms for Large Matrices via Random Projections". In: *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*. ISSN: 0272-5428. Oct. 2006, pp. 143–152.

[30] T. Strohmer and R. Vershynin. "A Randomized Kaczmarz Algorithm with Exponential Convergence". In: *Journal of Fourier Analysis and Applications* 15.2 (Apr. 2008), p. 262.

[31] J. A. Tropp. *Comparison theorems for the minimum eigenvalue of a random positive-semidefinite matrix*. arXiv:2501.16578 [math]. Jan. 2025.

[32] R. Vershynin. "Concentration inequalities for random tensors". In: *Bernoulli* 26.4 (Nov. 2020), pp. 3139–3162.

[33] S. Voronin and P.-G. Martinsson. "Efficient algorithms for cur and interpolative matrix decompositions". In: *Advances in Computational Mathematics* 43.3 (June 2017), pp. 495–516.

[34] S. Wang, A. Gittens, and M. W. Mahoney. "Scalable Kernel K-Means Clustering with Nystrom Approximation: Relative-Error Bounds". In: *Journal of Machine Learning Research* 20.12 (2019), pp. 1–49.

[35] F. Woolfe, E. Liberty, V. Rokhlin, and M. Tygert. "A fast randomized algorithm for the approximation of matrices". In: *Applied and Computational Harmonic Analysis* 25.3 (Nov. 2008), pp. 335–366.