Tyler Chen
tyler.chen@nyu.edu

Research Statement

https://chen.pw

My research focuses on the design and analysis of *fast and theoretically justified* algorithms for fundamental linear algebra tasks, with the overarching goal of developing tools to *support the advancement of knowledge in current scientific applications*. Towards this end, I work on problems from numerical linear algebra and nearby related disciplines. To date, this includes (i) work on Krylov subspace methods [CGMM22; CGMM23; CGW23; ACGMM23; XC23], where I have successfully integrated ideas from theoretical computer science, (ii) work on randomized numerical linear algebra [CTU22; CH23; PCM23; CC22; Che+23], including (iii) problems relating to optimization/machine learning [CTU21; BCW22] and (iv) algorithms designed specifically for tasks in computational quantum physics/chemistry [Che23a; CC22; Che+23], and (v) work on problems in the intersection of numerical linear algebra and random matrix theory [CGT23; CT23]. I am proud of the fact that much of my research provides conceptually simple insights into the behavior of widely used methods, which will help scientists more effectively address the problems they are interested in.

My current research program is centered on the following topics, each of which serves to further my overall goal of supporting basic science:

- *Krylov subspace methods methods*: classical problems relating to the stability of the Lanczos algorithm, improving how we think about the convergence of Krylov subspace methods in exact and finite precision arithmetic
- *Randomized numerical linear algebra*: algorithms for (partial) trace estimation, algorithms for low-rank approximation, the interplay between randomized algorithms and matrix functions
- *Computational quantum physics/chemistry*: thermodynamics of spin systems, partition functions, reduced density matrices,
- *Facilitation of cross-disciplinary knowledge transfer*: applications of Krylov subspace methods and randomized numerical linear algebra can be applied to problems in other disciplines, understanding how numerical linear algebra can benefit algorithm already used in these areas

The mentorship of students is an important aspect of my research program. Over the past several years, I have dedicated significant resources to mentoring students in research, with the goal of helping them develop into more independent researchers. My efforts including conceptualizing projects suitable for students without a PhD level background in computational math, as well managing as weekly meetings with student, and searching for opportunities for students to present about their research. While I have been consistently successful in leading projects which lead to publishable results [XC23; Che+23], I am most proud of the fact that of the majority of the students I have worked with are from groups historically underrepresented in math and science. As I progress into a tenure-track role, I intend to continue my current efforts working with undergraduate students while also providing new opportunities for graduate students.

In the remainder of this research statement I provide further details into the topics listed above, specifically on directions I am actively/interested in pursuing.

DESIGN AND ANALYSIS OF KRYLOV SUBSPACE METHODS

Krylov subspace methods (KSMs) are a powerful class of iterative methods, which are the method of choice for solving many large-sparse linear systems and eigenvalue problems [TB97; Gre97b; Saa03]. They are therefore among the most important algorithms in numerical linear algebra and computational science at large.

Methods for matrix function approximation. Approximating $f(\mathbf{A})\mathbf{b}$, the product of a matrix function[1] with a vector is an important linear algebra task. The most common example is the inverse $f(\mathbf{A})\mathbf{b} = \mathbf{A}^{-1}\mathbf{b}$ which corresponds to solving a linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$, but other functions such as the exponential, logarithm, square and inverse square root, and step function each have applications throughout the computational science [Hig08]. In most such applications $\mathbf{A}$ is too large to compute an eigendecomposition, so iterative methods such as the Lanczos method for matrix function approximation (Lanczos-FA) [DK89; Saa92] are the methods of choice.

For linear systems, algorithms such as Conjugate Gradient and MINRES enjoy strong instance optimality grantees, outputting the best possible approximations to $\mathbf{A}^{-1}\mathbf{b}$ from Krylov subspace [Gre97b]. Amazingly, for many other function, algorithms like Lanczos-FA seem to behave nearly as good as the optimal KSM. A more precise understanding of this algorithms behavior is important for guiding future algorithm development; if the algorithms truly are near optimal, then major algorithmic advances necessarily must come from moving beyond KSMs. On the other hand, finding cases for which the algorithm may lead to insights in how to develop better KSMs.

My past work [CGMM22; ACGMM23; XC23] incorporate ideas from theoretical computer science, such as reductions and instance optimality, to shed new light on the impressive performance of Lanczos-FA. In particular, these works help explain the fact that Lanczos-FA benefits when $\mathbf{A}$ has a favorable spectrum (i.e. large gaps, outlying eigenvalues, etc.). This is important, because many previous bounds for Lanczos-FA do not take into account fine-grained spectral properties, and therefore are sometimes misinterpreted as begin indicative of the typical behavior of the algorithm. My work has also resulted in the development of new optimal (in a certain norm) KSMs [CGMM23; CGW23] for a broad class of rational matrix functions.

Better bounds and better algorithms. Owing to the ubiquity of matrix functions in the sciences, there are still many theoretical and practical topics of study in this area. These range from theoretical tasks like deriving stronger optimality guarantees to more practical problems like developing tools for user friendly error estimates and optimized variants of existing algorithms.

Stability in finite precision arithmetic. Algorithms like Conjugate Gradient and Lanczos-FA are very computationally efficient: they access $\mathbf{A}$ using just matrix-vector products, and require storing only a 3 or 4 vectors of the dimension of $\mathbf{A}$. However, the Lanczos algorithm, on which they are based, behaves very differently in finite precision arithmetic than exact arithmetic; i.e. it is unstable. Because of this, *there is a widespread hesitance towards Lanczos-based approaches for problems involving matrix functions*, at least without the

[1] Given eigendecomposition $\mathbf{A} = \sum_{i=1}^{n} \lambda_i \mathbf{u}_i \mathbf{u}_i^\mathsf{T}$, the matrix function is $f(\mathbf{A}) = \sum_{i=1}^{n} f(\lambda_i)\mathbf{u}_i\mathbf{u}_i^\mathsf{T}$. Note we assume the symmetric case for simplicity of exposition, but many of the ideas are relevant to the non-symmetric case as well.

use of computationally expensive reorthogonalization techniques [JP94; SRVK96; ADEL03; WWAF06; UCS17; GWG19]. This hesitance has resulted in alternative algorithms being used in place of Lanczos-based methods, even though the alternatives have less desirable convergence properties, even in finite precision arithmetic!

Convergence of Lanczos-FA. There is actually a lot of work showing that the Lanczos algorithm is still effective for many tasks, even in finite precision arithmetic [Pai70; Pai72; Pai76; Pai80; Gre89; DK91; Kni96; MMS18]. One of the major themes of my thesis was addressing how we think about the tradeoffs between Lanczos-based methods and their alternatives. As I discuss more in a subsequent section, trying to disseminate this knowledge is a large aspect of my work facilitating cross-disciplinary knowledge transfer [CTU22; Che23a].

Maximum attainable accuracy. A distinct but related question is bounding the final accuracy of CG. That is, the minimum possible error (or residual) attained after the algorithm has stagnated due to the effects of finite precision arithmetic. A reasonable goal is to try and show a linear dependence of the error on the condition number of the system. ller than past analyses show [Gre89; Gre97a; MMS18], but seems to align with numerical experiments. As with all finite precision analyses, subtle details such as whether the analysis is for the standard CG implementation or a Lanczos-based implementation are important. I first became interested in this problem after discussions with Yuji Nakatsukasa about the stability of randomized algorithms for least squares problems [MNTW23].

Backwards stability. A career-long goal of mine is to find a simpler and more intuitive proof of Greenbaum's seminal backwards stability analysis for Lanczos [Gre89]. We made incremental progress towards this objective in [CT23], in which we proved a stronger backwards stability result for a restricted class of matrices by significantly extending an overlooked work of Knizhernman [Kni96].

RANDOMIZED NUMERICAL LINEAR ALGEBRA

Randomized numerical linear algebra involves introducing randomness into linear algebra algorithms in order to get them to run faster (asymptotically and practically) on the vast majority of problem instances. This paradigm has enabled the study of previously intractable scientific problems [HMT11; MT20].

Low-rank approximation. Obtaining a low-rank matrix which approximates to some target matrix is one of the key tasks in randomized numerical linear algebra. The most powerful randomized low-rank approximation methods are based on the Lanczos algorithm [MM15; TW23].

Low-memory algorithms. Unfortunately, unlike Lanczos-FA, if reorthogonlization is not used, current Lanczos-based low-rank approximation fail completely. Therefore, I am currently working on developing variants of these methods that enjoy the same theoretical guarantees as state of the art algorithms, but that can be implemented stably in a computationally efficient way; i.e. with the same computational profile as Lanczos without reorthogonalization.

Krylov-aware algorithms. A somewhat more specific task is obtaining a low-rank approximation of $f(\mathbf{A})$, and it is commonly suggested to simply combine existing low-rank algorithms with Lanczos-FA. However, as we show in [CH23], there are significant efficiencies to be gained by taking a more careful look at how low-rank algorithms interact with Lanczos-

FA. The theoretical convergence of the algorithm with respect to various parameters was analyzed more carefully in [PCM23], but a number of practical questions regarding how these parameters should be set and whether they can be determined adaptively remain. I am interested in exploring this further in order to obtain a more practical and usable implementation, and I believe this would be a particularly suitable project for student involvement.

Trace estimation.

Stochastic trace estimation involves estimating the trace of a matrix only using matrix-vector products [Gir87; Hut89; Ski89]. This is useful in cases where the matrix of interest might not be know explicitly (e.g. because it is a matrix function, corresponds to solving a differential equation, etc.).

Probing and query methods. The most basic trace estimators make use of isotropic test vectors (e.g. vectors with independent and identically distributed Gaussian entires) to construct an unbiased estimator It is possible to use structured test vectors to construct estimators with lower variance than Gaussian test vectors, at least when the matrix of interest has some structure (e.g. is tridiagonal) [TS11; HT22; FRS23].

I am most interested in developing probing techniques for problems arising in quantum physics (described in the next section), where certain types of kronecker structure is common. I designed Algorithm 1 in [MA23] for exact recovery of a matrix with a certain kind of Kronecker structure. It would be interesting to extend the method to matrices which only approximately have the relevant structure.

Easy-to-use software. In [Che23a] I introduced implementation of the well-known kernel polynomial method (KPM) based on the Lanczos algorithm. This implementation avoids the need for certain hyperparameters to be determined a priori, making it feasible to test out many different choices of parameters, such as the orthogonal polynomial family, which can improve the qualitative features of the approximation.

While there are presently a number of one-off packages which make use of KPM approximations, particularly in condensed matter physics [GWAW14; RF20; MMF20], these lack the flexibility of the method in [Che23a]. Thus, I am currently developing the `spectral_density` package [Che23b] which allows users to compute KPM approximations from the output of their favorite Lanczos implementation. I hope to eventually include more tools, including practical error estimates for related algorithms for trace estimation.

Sampling methods. For the task of solving linear least squares problems, randomized algorithms such as stochastic gradient descent and randomized Kaczmarz can theoretically beat the runtime (but not iteration count) of classical KSMs in some cases. In [BCW22], we show that, with a large enough batch size, minibatch stochastic gradient descent can in fact match the iteration count of of classical KSMs.

Methods for matrix functions. I'm also currently in the early stages of a project with a student on applying stochastic optimization methods (like randomized Kaczmarz or stochastic gradient descent) to compute samples of a Gaussian vector whose covariance is determined by the sample covariance matrix of some data. Broadly, I'm interested in understanding the extent to which randomization can be used to accelerate Lanczos-based methods.

QUANTUM PHYSICS/CHEMISTRY

While all of my research broadly aims to support advancement of knowledge in current scientific applications, I also work on a algorithms designed specifically to study scientific questions in quantum physics and chemistry.

Dissemination / implementation. A number of algorithms related to trace estimation have been developed more-or-less independently by the Physics community and the numerical linear algebra and theoretical computer science communities [LSY16; Jin+21]. Over the past several years I have worked on bridging the gap between these communities. For instance, in [CTU22], we provide a unified analysis of stochastic Lanczos quadrature (SLQ) and the kernel polynomial method (KPM). These are two widely used methods for approximating $\mathrm{tr}(f(\mathbf{A}))$, with SLQ more widely used in applied math and KPM more widely used in physics. Subsequently, in [Che23a], we show how the Lanczos algorithm can be used to stably implement the KPM, avoiding the need for choosing certain hyperparameters ahead of time. I also gave talks discussing connections between these communities at the "Perspectives on Matrix Computations: Theoretical Computer Science Meets Numerical Analysis" workshop, as well as a recent seminar talk to a physics department.

Algorithms for partial traces. In [CC22], we introduced a simple matrix-free stochastic estimator for the *partial trace*[2], an operation ubiquitous in quantum physics. To the best of my knowledge, this is the first stochastic estimator for partial traces, perhaps due to the simple fact that most linear algebraists have never heard of the partial trace! In subsequent work [Che+23], we improved the estimator of [CC22] by introducing a variance reduction technique based on ideas from standard trace estimation [Gir87]. This project was done with undergraduate students at NYU and one high-school student from Stuyvesant HS.

Both [CC22; Che+23] focus on applications to quantum thermodynamics. In particular, the state of an open quantum systems with Hamiltonian $\mathbf{A}$ held in thermal equilibrium due to weak coupling with a "super bath" is described by a total system density matrix. To understand how one part of the system of interest behaves (e.g. its entanglement with other parts of the total system, its energy, heat capacity, etc.), one must compute the partial trace of the total system density matrix with respect to the part of the system not of interest [CZT10; IHT09; TH20].

In many cases, $\mathbf{A}$ is highly structured. However, the algorithms of [CC22; Che+23] do not use this structure in a fine-grained way. Taking advantage of the structure inherent to $\mathbf{A}$ provides a potential approach to developing more efficient algorithms.

Residual estimation methods. The most basic example of structure in $\mathbf{A}$ is if $\mathbf{A} = \mathbf{A}_1 + \theta\mathbf{A}_2$, where θ is some parameter of interest (e.g. the magnetic field strength). If we are interested in understanding the partial trace of $f(\mathbf{A})$ for a range of θ, then we may hope to do something better than applying existing algorithms to each value of θ independently. In particular, the partial trace and matrix exponential are continuous functions of θ, and so we can try to use some kind of residual estimation methods [Gir87; MMMW21; DM21]. This project would be suitable for involvement of student researchers.

Methods for dynamics. Besides considering systems in thermal equilibrium, physicists are

[2]The partial trace can intuitively be though of as the quantum analog to integrating out irrelevant degrees of freedom; e.g. computing a marginal density from a joint density

interested in the time evolution of quantum systems. While there are a number of tensor network methods for this task [Cyg+22; CSK22, etc.], such methods are not particularly suitable for systems with long-range interactions if the initial state is is a thermal state at high temperature. In this setting classical randomized Krylov subspace methods and offer speedups over dense liner algebra methods.

CROSS-DISCIPLINARY KNOWLEDGE TRANSFER

I enjoy acting as an intermediary between disciplines in order to facilitate the transfer of what would otherwise remain domain specific knowledge, and my efforts have resulted in progress on topics which would have been intractable without communication between disciplines.

Theoretical computer science. During my PhD I worked to incorporate perspectives from theoretical computer science into my research, and initiated collaboration with theoretical computer scientists Christoper Musco and Cameron Musco, who have since become some of my closest collaborators. Our collaborations [CGMM22; CGMM23; ACGMM23] incorporate ideas from theoretical computer science such as instance optimality and reductions in order to analyze classical algorithms in numerical linear algebra.

Optimization and theoretically justified machine learning. Following discussions with Rachel Ward about how CG and related KSMs are viewed in the optimization and machine learning communities, as well as some open questions I was interested in resolving, we began a collaboration analyzing stochastic minibatch gradient descent, a widely used algorithm for training machine learning models [BCW22]. I am broadly interested in understanding how stochastic sampling techniques from optimization into algorithms for computing $f(\mathbf{A})\mathbf{b}$.

In addition my work [CTU21], published in the proceedings for the International Conference on Machine Learning, provides the first theoretical grantees for a spectrum estimation algorithm used to study the Hessians of neural networks. The quality of the work is evidence by the fact it was one of roughly 3% of submitted papers selected for a long-presentation.

Random matrix theory. Testing numerical algorithms on random matrices is a widespread practice. However, as noted by Edelman and Rao [ER05], "It is a mistake to link psychologically a random matrix with the intuitive notion of a 'typical' matrix or the vague concept of 'any old matrix"'. In [CGT23], we prove a version of Crouzeix's conjecture [CP17] for a class of random matrices, and in [CT23], we prove that the Lanczos algorithm is forward stable when run on many random matrix models. These provide further evidence that care must be taken when using random matrices as test matrices, as the behavior can be very different from worst case behavior.

Computational quantum physics. As noted in the previous section, I am especially interested in applications in computational quantum physics, and have been successful in developing algorithms designed specifically for studying problems in this domain [CC22; Che+23; Che23a].

Tyler Chen
tyler.chen@nyu.edu

Bibliography

https://chen.pw

[ACGMM23] N. Amsel, **T. Chen**, A. Greenbaum, C. Musco, and C. Musco. "Near-Optimality Guarantees for Approximating Rational Matrix Functions by the Lanczos Method". 2023.

[ADEL03] M. Aichhorn, M. Daghofer, H. G. Evertz, and W. von der Linden. "Low-temperature Lanczos method for strongly correlated systems". In: *Physical Review B* 67.16 (2003).

[BCW22] R. Bollapragada, **T. Chen**, and R. Ward. "On the fast convergence of minibatch heavy ball momentum". 2022.

[CC22] **T. Chen** and Y.-C. Cheng. "Numerical computation of the equilibrium-reduced density matrix for strongly coupled open quantum systems". In: *The Journal of Chemical Physics* 157.6 (2022).

[CGMM22] **T. Chen**, A. Greenbaum, C. Musco, and C. Musco. "Error Bounds for Lanczos-Based Matrix Function Approximation". In: *SIAM Journal on Matrix Analysis and Applications* 43.2 (2022).

[CGMM23] **T. Chen**, A. Greenbaum, C. Musco, and C. Musco. "Low-memory Krylov subspace methods for optimal rational matrix function approximation". In: *SIAM Journal on Matrix Analysis and Applications* (2023). to appear.

[CGT23] **T. Chen**, A. Greenbaum, and T. Trogdon. "GMRES, pseudospectra, and Crouzeix's conjecture for shifted and scaled Ginibre matrices". 2023.

[CGW23] **T. Chen**, A. Greenbaum, and N. Wellen. "Optimal Polynomial Approximation to Rational Matrix Functions Using the Arnoldi Algorithm". 2023.

[CH23] **T. Chen** and E. Hallman. "Krylov-Aware Stochastic Trace Estimation". In: *SIAM Journal on Matrix Analysis and Applications* 44.3 (2023).

[Che+23] **T. Chen**, R. Chen, K. Li, S. Nzeuton, Y. Pan, and Y. Wang. "Faster randomized partial trace estimation". 2023.

[Che23a] **T. Chen**. "A spectrum adaptive kernel polynomial method". In: *The Journal of Chemical Physics* 159.11 (2023).

[Che23b] T. Chen. *`spectral_density`: a python package for computing spectral density approximations*. Version 0.1.0. 2023.

[CP17] M. Crouzeix and C. Palencia. "The Numerical Range is a $(1 + \sqrt{2})$-Spectral Set". In: *SIAM Journal on Matrix Analysis and Applications* 38.2 (2017).

[CSK22] Y.-F. Chiu, A. Strathearn, and J. Keeling. "Numerical evaluation and robustness of the quantum mean-force Gibbs state". In: *Physical Review A* 106.1 (2022).

[CT23] **T. Chen** and T. Trogdon. "Stability of the Lanczos algorithm on matrices with regular spectral distributions". 2023.

[CTU21] **T. Chen**, T. Trogdon, and S. Ubaru. "Analysis of stochastic Lanczos quadrature for spectrum approximation". In: *Proceedings of the 38th International Conference on Machine Learning*. Vol. 139. Proceedings of Machine Learning Research. PMLR, 2021.

[CTU22] **T. Chen**, T. Trogdon, and S. Ubaru. "Randomized matrix-free quadrature for spectrum and spectral sum approximation". 2022.

[Cyg+22] M. Cygorek, M. Cosacchi, A. Vagov, V. M. Axt, B. W. Lovett, J. Keeling, and E. M. Gauger. "Simulation of open quantum systems by automated compression of arbitrary environments". In: *Nature Physics* 18.6 (2022).

[CZT10] M. Campisi, D. Zueco, and P. Talkner. "Thermodynamic anomalies in open quantum systems: Strong coupling effects in the isotropic XY model". In: *Chemical Physics* 375.2-3 (2010).

[DK89] V. Druskin and L. Knizhnerman. "Two polynomial methods of calculating functions of symmetric matrices". In: *USSR Computational Mathematics and Mathematical Physics* 29.6 (1989).

[DK91] V. L. Druskin and L. A. Knizhnerman. "Error Bounds in the Simple Lanczos Procedure for Computing Functions of Symmetric Matrices and Eigenvalues". In: *Comput. Math. Math. Phys.* 31.7 (1991). ISSN: 0965-5425.

[DM21] P. Dharangutte and C. Musco. "Dynamic Trace Estimation". In: *Advances in Neural Information Processing Systems*. Ed. by M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan. Vol. 34. Curran Associates, Inc., 2021.

[ER05] A. Edelman and N. R. Rao. "Random matrix theory". In: *Acta Numerica* 14 (2005).

[FRS23] A. Frommer, M. Rinelli, and M. Schweitzer. "Analysis of stochastic probing methods for estimating the trace of functions of sparse symmetric matrices". 2023.

[Gir87] D. Girard. "Un algorithme simple et rapide pour la validation croisée généralisée sur des problèmes de grande taille". 1987.

[Gre89] A. Greenbaum. "Behavior of slightly perturbed Lanczos and conjugate-gradient recurrences". In: *Linear Algebra and its Applications* 113 (1989). ISSN: 0024-3795.

[Gre97a] A. Greenbaum. "Estimating the Attainable Accuracy of Recursively Computed Residual Methods". In: *SIAM Journal on Matrix Analysis and Applications* 18.3 (1997).

[Gre97b] A. Greenbaum. *Iterative Methods for Solving Linear Systems*. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 1997. ISBN: 0-89871-396-X.

[GWAW14] C. W. Groth, M. Wimmer, A. R. Akhmerov, and X. Waintal. "Kwant: a software package for quantum transport". In: *New Journal of Physics* 16.6 (2014).

[GWG19] D. Granziol, X. Wan, and T. Garipov. "Deep Curvature Suite". 2019.

[Hig08] N. J. Higham. *Functions of Matrices*. Society for Industrial and Applied Mathematics, 2008.

[HMT11] N. Halko, P. G. Martinsson, and J. A. Tropp. "Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions". In: *SIAM Review* 53.2 (2011).

[HT22] D. Halikias and A. Townsend. “Structured matrix recovery from matrix-vector products”. 2022.

[Hut89] M. Hutchinson. “A Stochastic Estimator of the Trace of the Influence Matrix for Laplacian Smoothing Splines”. In: *Communications in Statistics - Simulation and Computation* 18.3 (1989).

[IHT09] G.-L. Ingold, P. Hänggi, and P. Talkner. “Specific heat anomalies of open quantum systems”. In: *Physical Review E* 79.6 (2009).

[Jin+21] F. Jin, D. Willsch, M. Willsch, H. Lagemann, K. Michielsen, and H. De Raedt. “Random State Technology”. In: *Journal of the Physical Society of Japan* 90.1 (2021). ISSN: 1347-4073.

[JP94] J. Jaklič and P. Prelovšek. “Lanczos method for the calculation of finite-temperature quantities in correlated systems”. In: *Physical Review B* 49.7 (1994).

[Kni96] L. A. Knizhnerman. “The Simple Lanczos Procedure: Estimates of the Error of the Gauss Quadrature Formula and Their Applications”. In: *Comput. Math. Math. Phys.* 36.11 (1996). ISSN: 0965-5425.

[LSY16] L. Lin, Y. Saad, and C. Yang. “Approximating Spectral Densities of Large Matrices”. In: *SIAM Review* 58.1 (2016).

[MA23] R. A. Meyer and H. Avron. “Hutchinson’s Estimator is Bad at Kronecker-Trace-Estimation”. 2023.

[MM15] C. Musco and C. Musco. “Randomized Block Krylov Methods for Stronger and Faster Approximate Singular Value Decomposition”. Montreal, Canada, 2015.

[MMF20] D. Moldovan, Miša Anđelković, and Francois Peeters. “pybinding v0.9.5: a Python package for tight-binding calculations”. 2020.

[MMMW21] R. A. Meyer, C. Musco, C. Musco, and D. P. Woodruff. “Hutch++: Optimal Stochastic Trace Estimation”. In: *Symposium on Simplicity in Algorithms (SOSA)*. Society for Industrial and Applied Mathematics, 2021.

[MMS18] C. Musco, C. Musco, and A. Sidford. “Stability of the Lanczos Method for Matrix Function Approximation”. In: Society for Industrial and Applied Mathematics, 2018.

[MNTW23] M. Meier, Y. Nakatsukasa, A. Townsend, and M. Webb. “Are sketch-and-precondition least squares solvers numerically stable?” 2023.

[MT20] P.-G. Martinsson and J. A. Tropp. “Randomized numerical linear algebra: Foundations and algorithms”. In: *Acta Numerica* 29 (2020).

[Pai70] C. C. Paige. “Practical use of the symmetric Lanczos process with re-orthogonalization”. In: *BIT* 10.2 (1970).

[Pai72] C. C. Paige. “Computational Variants of the Lanczos Method for the Eigenproblem”. In: *IMA Journal of Applied Mathematics* 10.3 (1972).

[Pai76] C. C. Paige. “Error Analysis of the Lanczos Algorithm for Tridiagonalizing a Symmetric Matrix”. In: *IMA Journal of Applied Mathematics* 18.3 (1976). ISSN: 0272-4960.

[Pai80] C. C. Paige. “Accuracy and effectiveness of the Lanczos algorithm for the symmetric eigenproblem”. In: *Linear Algebra and its Applications* 34 (1980). ISSN: 0024-3795.

[PCM23] D. Persson, **T. Chen**, and C. Musco. “Randomized block Krylov subspace methods for low rank approximation of matrix functions”. 2023.

[RF20] B. M. F. de Resende and L. da F. Costa. “Characterization and comparison of large directed networks through the spectra of the magnetic Laplacian”. In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 30.7 (2020).

[Saa03] Y. Saad. *Iterative methods for sparse linear systems*. 2nd ed. Philadelphia: SIAM, 2003. ISBN: 9780898715347.

[Saa92] Y. Saad. “Analysis of Some Krylov Subspace Approximations to the Matrix Exponential Operator”. In: *SIAM Journal on Numerical Analysis* 29.1 (1992).

[Ski89] J. Skilling. “The Eigenvalues of Mega-dimensional Matrices”. In: *Maximum Entropy and Bayesian Methods*. Springer Netherlands, 1989.

[SRVK96] R. Silver, H. Roeder, A. Voter, and J. Kress. “Kernel Polynomial Approximations for Densities of States and Spectral Functions”. In: *Journal of Computational Physics* 124.1 (1996).

[TB97] L. N. Trefethen and D. Bau Iii. *Numerical Linear Algebra*. en. Philadelphia, PA: Society for Industrial and Applied Mathematics, 1997. ISBN: 9780898713619.

[TH20] P. Talkner and P. Hänggi. “Colloquium : Statistical mechanics and thermodynamics at strong coupling: Quantum and classical”. In: *Reviews of Modern Physics* 92.4 (2020).

[TS11] J. M. Tang and Y. Saad. “A probing method for computing the diagonal of a matrix inverse”. In: *Numerical Linear Algebra with Applications* 19.3 (2011).

[TW23] J. A. Tropp and R. J. Webber. “Randomized algorithms for low-rank matrix approximation: Design, analysis, and applications”. 2023.

[UCS17] S. Ubaru, J. Chen, and Y. Saad. “Fast Estimation of tr($f(A)$) via Stochastic Lanczos Quadrature”. In: *SIAM Journal on Matrix Analysis and Applications* 38.4 (2017).

[WWAF06] A. Weiße, G. Wellein, A. Alvermann, and H. Fehske. “The kernel polynomial method”. In: *Reviews of Modern Physics* 78.1 (2006).

[XC23] Q. Xu and **T. Chen**. “A posteriori error bounds for the block-Lanczos method for matrix function approximation”. 2023.