Quantum algorithms for spectral sums

We present quantum algorithms to estimate *four* important quantities connected to a matrix: the log-determinant, the von Neumann entropy of a graph, the trace of inverse, and the Schatten p-norm. All these quantities are *spectral sums* of a matrix. The spectral sum of a positive semi-definite (PSD) matrix A, for a function $f : \mathbb{R} \mapsto \mathbb{R}$, is defined as $\mathcal{S}_f(A) = \sum_j f(\lambda_j)$, where λ_j are the eigenvalues of A. This notion can be generalized to non-square matrices using singular value decomposition. As a corollary of our work, we show how to use our results and techniques for *three* problems in spectral graph theory: we discuss how to estimate the number of triangles in a graph (with two algorithms that work in two different regimes), estimate the effective resistance between nodes of an electric network, and count the spanning trees in a graph.

As in classical algorithms, we leverage the idea that the sum of the eigenvalues of a matrix is equal to its trace. Thus, we perform trace estimation of a suitably created matrix which we obtain with quantum singular value transformation techniques([1, 2, 3, 4, 5, 6]). Quantum algorithms for computing the trace are well-studied in many models of quantum computation (see [7, 8]) and can be generalized to the gate-based model, and in particular to the block-encoding model [9, 10], which we use. In this model, we have access to a unitary (and its inverse) that encodes a scaled matrix A in the top-left corner. We rewrite and formalize — using the notation of block-encodings — an algorithm that returns an estimate of the trace with absolute and relative error, for which we offer a useful formal statement that can be used in other settings. The algorithms that work with relative error use the algorithm that returns an estimate with absolute error, without knowing a lower bound of the quantity being estimated (see [9, 10]). In our work, we also use subroutines to estimate the trace of $A^T A$ with relative error [11]. To our knowledge, this idea was not previously used for similar problems. Its utility is twofold: on one side it allows working with non-square matrices, and on the other side allows for a quadratic speedup in the size of the matrix. The cost for using this subroutine is preparing a block-encoding for $\sqrt{f(A)}$. We show how to use these ideas for the problem of counting triangles and Schatten p-norms. We add to the list of useful quantum subroutines to manipulate block-encodings as an extension of [4, Lemma 16], which allows to take the product of k preamplified block-encodings. When working with classical data, we can build block-encodings within the *quantum memory device* model (i.e. using a QMD, a quantum memory device [12] — e.g. a QRAM [13]). Lastly, our algorithms can work in a *quantum arithmetic model*, which is a convention that we use to specify decimal numbers using fixed-precision. In this model, we point out that the algorithm for obtaining a relative error estimate from a subroutine that returns an absolute error, can be slightly simplified [14]. In the conclusion of the manuscript we discuss an ample list of possible future research directions.

In our work we discuss the relationship between the algorithms presented in this work and the *one clean qubit model* of computation: a model where one is given only a single clean qubit and some of the qubits in the maximally mixed state. This model — which was inspired by past experiments in MRI — is less powerful than the general quantum circuit model, but is believed to be hard to simulate classically, and started to gain interest again in the community [15, 16]. We also discuss our works in relation to recent research in *limited-depth quantum computation*: a setting where we want to reap the benefits of quantum computing but we don't have access to computers that can run a circuit for the depth that is necessary to execute the full algorithm. In particular, all our algorithm can go from using $O(\frac{1}{\epsilon})$ queries to the input with a circuit of the same depth, to an algorithm that uses a circuit of depth $O(\frac{1}{\epsilon^{(1-\beta)}})$ for $O(\frac{1}{\epsilon^{(1+\beta)}})$ times.

In the following, let κ be the condition number of the matrix A, and α be a quantity upper bounded by the Frobenius norm of the matrix (both are precisely defined in the manuscript). The notation $\widetilde{O}(\cdot)$ hides polylogarithmic factors from the asymptotic complexity. For all the following algorithms we show *tight* bounds on the error of the block-encoding we receive as input for our algorithms. While in the manuscript we report all the polylogarithmic factors in the runtime, our results can be summarised and simplified as follows.

Quantum algorithm for log-determinants It is hard to underestimate the importance of algorithms for estimating the (logarithm of the) determinant of an PSD matrix A, which is defined as $\text{logdet}(A) = \sum_{j=1}^{n} \log \lambda_j$. Notable examples can be found in machine learning and other computational sciences [17, 18, 19, 20, 21]. Our result reads as follows. While many classical and quantum algorithms exist, we improve (polynomially) the previous quantum algorithm for the log-determinant [22, 23].

Theorem (Informal - Algorithm for estimating log-determinants)**.** *Let U_A be a α-block-encoding of an PSD matrix $A \in \mathbb{R}^{n\times n}$. There is a quantum algorithm that returns an estimate $\overline{\text{logdet}(A)}$ of the log-determinant with relative error and high probability using $\widetilde{O}\left(\frac{\alpha\kappa}{\epsilon}\right)$ queries to U_A.*

Two Quantum algorithms for Schatten p-norm Let $A \in \mathbb{R}^{m\times n}$ be a matrix with singular values $\{\sigma_i\}_{i=1}^{\min(m,n)}$. The Schatten p-norm of A is $\|A\|_p = \left(\sum_{i=1}^{\min(m,n)} \sigma_i^p\right)^{1/p}$. There are many applications for the algorithms that compute the Schatten p-norms, for example in completion algorithms, theoretical chemistry, image processing, and many others [24, 25, 26, 27, 28, 29]. We propose two different algorithms for estimating the Schatten p-norm of a matrix. Both use the subroutine for estimating the trace of B^TB for a suitably contracted matrix B. For the second, for big p, we use an efficient polynomial approximation of monomials in $[-1, 1]$ that allows to have a quadratic improvement in $\sqrt{2}^p$. This is better formulated in the main manuscript. We also report the results of experiments showing the asymptotic scaling of certain parameters on real-world datasets. Our results improve polynomially over previous quantum algorithms [30].

Theorem (Informal - Algorithm for Schatten p-norms)**.** *Let U_A be a α-block-encoding of a matrix $A \in \mathbb{R}^{m\times n}$ with $n \leq m$. There is a quantum algorithm that returns an estimate $\overline{\|A\|_p}$ with relative error and high probability using U_A $\widetilde{O}\left(\frac{p\sqrt{n}(\sqrt{2}\|A\|)^{p/2}}{\epsilon\|A\|_p^{p/2}}\left(\frac{\alpha}{\|A\|}\right)\right)$ times if p is even, or $\widetilde{O}\left(\frac{\sqrt{n}\alpha^{1.5}(\sqrt{2})^{p/2}(\|A\|)^{p/2-1}}{\epsilon\|A\|_p^{p/2}}(p+\kappa)\right)$ times if p is odd.*

Quantum algorithm for von Neumann entropy of graphs We show the runtime of an algorithm to estimate the von Neumann entropy of a matrix without having access to a purification that allows the creation of a block-encoding of the density matrix. This is especially relevant for matrices that are obtained graph Laplacian. For a graph $G = (V, E)$ a graph Laplacian $\mathcal{L}$ is defined as $\mathcal{L} = \Delta(G) - A(G)$ We can associate to G a density matrix $\rho_G = \mathcal{L}/\text{Tr}[\mathcal{L}]$. Then, the von Neumann entropy of the graph G is defined as $H(G) := H(\rho_G) = -\text{Tr}[\rho_G \log \rho_G]$. In the following $s(G)$ is the number of edges in the graph G, i.e. $|E|$. Von Neumann entropy is used in financial data analysis, genomics, complex network analysis, and pattern recognition [31, 32, 33, 34, 35, 36]. There is a vast literature of quantum algorithms for von Neumann entropies of density matrices ρ, but they work in a model where you have access to a purification of the density matrix, which results in a 1-block encoding of ρ [14, 37, 38, 39, 40].

Theorem (Informal - Algorithm for estimating von Neumann entropy)**.** *Let U_A be a α-block-encoding of an PSD matrix $A \in \mathbb{R}^{n\times n}$. There is a quantum algorithm that returns an estimate $\overline{H(A)}$ with absolute error and high probability using $\widetilde{O}\left(\frac{n\alpha\kappa}{\epsilon s(G)}\right)$ queries to U_A.*

Quantum algorithm for trace of the inverse For a square matrix A, the trace of the inverse is $I(A) = \text{Tr}[A^{-1}] = \sum_{i=1}^{n} \lambda_i^{-1}$. The trace of the inverse finds application in lattice quantum chromodynamics, generalized cross validation, and uncertainty quantification [41, 42, 43]. To our knowledge, there are only previous works for the trace of the inverse for classical algorithms, while there are no previous quantum algorithms for the same problem.

Theorem (Informal - Algorithm for trace of inverse)**.** *Let $\epsilon > 0$, and A be a matrix $A \in \mathbb{R}^{n\times n}$ with. There is a quantum algorithm that returns an estimate $\overline{I(A)}$ with relative error ϵ using $\widetilde{O}\left(\frac{\mu^2\kappa^2}{\epsilon}\right)$ queries to U_A.*

Applications in spectral graph theory We discuss three applications of our algorithms to problems in spectral graph theory. All the problems we discuss have important applications in many computational sciences. First, we show how to use the techniques from the algorithm for computing the Schatten p-norms for counting the number of triangles in a graph. We discuss *two* algorithms. The runtime of the first algorithm is $O(\frac{\alpha n}{\epsilon\Delta(G)})$, and the runtime of the second algorithm is $O(\frac{\alpha\sqrt{n}\kappa}{\epsilon\sqrt{\Delta(G)}})$. Algorithms for counting triangles find many applications in network analysis [44, 45], and importantly, generalize to the problem of finding k-cycles, another fundamental problem in computer science. Second, we show how to give a relative error estimate of the number of spanning trees in a graph G, assuming block-encoding access to the graph Laplacian. Our quantum algorithm makes $\widetilde{O}(n\alpha\kappa/\epsilon)$ calls to the block-encoding of the Laplacian matrix of A. There are many applications in machine learning [46], genomics, and network theory [47, 48]. Third, we show how to use the algorithm for estimating the log-determinant for computing the effective resistance between two nodes in a graph using the block-encoding of (a modified) graph Laplacian $\mathcal{L}$ for $\widetilde{O}(n\alpha\kappa/\epsilon)$ times. Applications include graph clustering, graph sparsification and analysis, graph neural networks, and many others [49, 50, 51, 52, 53].

References

[1] S. Chakraborty, A. Gilyén, and S. Jeffery. The power of block-encoded matrix powers: improved regression techniques via faster hamiltonian simulation. In *Proc. 46th International Colloquium on Automata, Languages, and Programming*, pages 33:1–33:14, 2019. 1

[2] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 193–204, 2019. 1

[3] Guang Hao Low and Isaac L Chuang. Hamiltonian simulation by qubitization. *Quantum*, 3:163, 2019. 1

[4] Shantanav Chakraborty, Aditya Morolia, and Anurudh Peduri. Quantum regularized least squares. *arXiv preprint arXiv:2206.13143*, 2022. 1

[5] Ewin Tang and Kevin Tian. A cs guide to the quantum singular value transformation. In *2024 Symposium on Simplicity in Algorithms (SOSA)*, pages 121–143. SIAM, 2024. 1

[6] Danial Motlagh and Nathan Wiebe. Generalized quantum signal processing. *arXiv preprint arXiv:2308.01501*, 2023. 1

[7] Peter W Shor and Stephen P Jordan. Estimating jones polynomials is a complete problem for one clean qubit. *Quantum Information & Computation*, 8(8):681–714, 2008. 1

[8] Chris Cade and Ashley Montanaro. The quantum complexity of computing schatten p-norms. In *13th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018. 1

[9] Anirban N Chowdhury, Rolando D Somma, and Yiğit Subaşı. Computing partition functions in the one-clean-qubit model. *Physical Review A*, 103(3):032422, 2021. 1

[10] Sathyawageeswar Subramanian and Min-Hsiu Hsieh. Quantum algorithm for estimating α-renyi entropies of quantum states. *Physical review A*, 104(2):022428, 2021. 1

[11] Joran Van Apeldoorn, András Gilyén, Sander Gribling, and Ronald de Wolf. Quantum sdp-solvers: Better upper and lower bounds. *Quantum*, 4:230, 2020. 1

[12] Jonathan Allcock, Jinge Bao, João F Doriguello, Alessandro Luongo, and Miklos Santha. Constant-depth circuits for uniformly controlled gates and boolean functions with application to quantum memory circuits. *arXiv preprint arXiv:2308.08539*, 2023. 1

[13] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Architectures for a quantum random access memory. *Physical Review A*, 78(5):052310, 2008. 1

[14] Sathyawageeswar Subramanian and Min-Hsiu Hsieh. Quantum algorithm for estimating α-renyi entropies of quantum states. *Physical Review A*, 104(2):022428, 2021. 1, 2

[15] Jonathan A Jones. Controlling nmr spin systems for quantum computation. *Progress in Nuclear Magnetic Resonance Spectroscopy*, 2024. 1

[16] Tao Xin, Bi-Xue Wang, Ke-Ren Li, Xiang-Yu Kong, Shi-Jie Wei, Tao Wang, Dong Ruan, and Gui-Lu Long. Nuclear magnetic resonance for quantum computing: Techniques and recent achievements. *Chinese Physics B*, 27(2):020308, 2018. 1

[17] Havard Rue and Leonhard Held. *Gaussian Markov random fields: theory and applications*. CRC press, 2005. 1

[18] Carl Edward Rasmussen and Hannes Nickisch. Gaussian processes for machine learning (GPML) toolbox. *Journal of machine learning research*, 11(Nov):3011–3015, 2010. 1

[19] Andrzej Cichocki, Sergio Cruces, and Shun-ichi Amari. Log-determinant divergences revisited: Alpha-beta and gamma log-det divergences. *Entropy*, 17(5):2988–3034, 2015. 1

[20] Iordanis Kerenidis, Alessandro Luongo, and Anupam Prakash. Quantum expectation-maximization for gaussian mixture models. In *International Conference on Machine Learning*, pages 5187–5197. PMLR, 2020. 1

[21] Douglas A Reynolds. Gaussian mixture models. *Encyclopedia of biometrics*, 741(659-663), 2009. 1

[22] Zhikuan Zhao, Jack K Fitzsimons, Michael A Osborne, Stephen J Roberts, and Joseph F Fitzsimons. Quantum algorithms for training gaussian processes. *Physical Review A*, 100(1):012304, 2019. 1

[23] Liming Zhao, Zhikuan Zhao, Patrick Rebentrost, and Joseph Fitzsimons. Compiling basic linear algebra subroutines for quantum computers. *Quantum Machine Intelligence*, 3(2):1–10, 2021. 1

[24] Praneeth Netrapalli, UN Niranjan, Sujay Sanghavi, Animashree Anandkumar, and Prateek Jain. Non-convex robust PCA. In *Advances in Neural Information Processing Systems*, pages 1107–1115, 2014. 2

[25] Ivan Gutman. The energy of a graph: old and new results. In *Algebraic combinatorics and applications*, pages 196–211. Springer, 2001. 2

[26] Yu Lu and Sahand N Negahban. Individualized rank aggregation using nuclear norm regularization. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1473–1479. IEEE, 2015. 2

[27] Feiping Nie, Heng Huang, and Chris Ding. Low-rank matrix recovery via efficient schatten p-norm minimization. In *Twenty-sixth AAAI conference on artificial intelligence*, 2012. 2

[28] Angshul Majumdar and Rabab K Ward. An algorithm for sparse mri reconstruction by schatten p-norm minimization. *Magnetic resonance imaging*, 29(3):408–417, 2011. 2

[29] Yuan Xie, Shuhang Gu, Yan Liu, Wangmeng Zuo, Wensheng Zhang, and Lei Zhang. Weighted schatten p-norm minimization for image denoising and background subtraction. *IEEE transactions on image processing*, 25(10):4842–4857, 2016. 2

[30] Ashley Montanaro. Quantum speedup of Monte Carlo methods. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 471(2181):20150301, 2015. 2

[31] Petre Caraiani. The predictive power of singular value decomposition entropy for stock market dynamics. *Physica A: Statistical Mechanics and its Applications*, 393:571–578, 2014. 2

[32] Monami Banerjee and Nikhil R Pal. Feature selection with svd entropy: Some modification and extension. *Information Sciences*, 264:118–134, 2014. 2

[33] Orly Alter, Patrick O Brown, and David Botstein. Singular value decomposition for genome-wide expression data processing and modeling. *Proceedings of the National Academy of Sciences*, 97(18):10101–10106, 2000. 2

[34] Giorgia Minello, Luca Rossi, and Andrea Torsello. On the von Neumann entropy of graphs. *Journal of Complex Networks*, 7(4):491–514, 2018. 2

[35] Lin Han, Francisco Escolano, Edwin R Hancock, and Richard C Wilson. Graph characterizations from von neumann entropy. *Pattern Recognition Letters*, 33(15):1958–1967, 2012. 2

[36] Filippo Passerini and Simone Severini. The von neumann entropy of networks. *Available at SSRN 1382662*, 2008. 2

[37] Tongyang Li and Xiaodi Wu. Quantum query complexity of entropy estimation. *IEEE Transactions on Information Theory*, 65(5):2899–2921, 2018. 2

[38] András Gilyén and Tongyang Li. Distributional property testing in a quantum world. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020. 2

[39] Qisheng Wang, Ji Guan, Junyi Liu, Zhicheng Zhang, and Mingsheng Ying. New quantum algorithms for computing quantum entropies and distances. *arXiv preprint arXiv:2203.13522*, 2022. 2

[40] Tom Gur, Min-Hsiu Hsieh, and Sathyawageeswar Subramanian. Sublinear quantum algorithms for estimating von neumann entropy. *arXiv preprint arXiv:2111.11139*, 2021. 2

[41] Masoumeh Dashti and Andrew M Stuart. Uncertainty quantification and weak approximation of an elliptic inverse problem. *SIAM Journal on Numerical Analysis*, 49(6):2524–2542, 2011. 2

[42] Gene H Golub and Urs Von Matt. Generalized cross-validation for large-scale problems. *Journal of Computational and Graphical Statistics*, 6(1):1–34, 1997. 2

[43] Andreas Stathopoulos, Jesse Laeuchli, and Kostas Orginos. Hierarchical probing for estimating the trace of the matrix inverse on toroidal lattices. *SIAM Journal on Scientific Computing*, 35(5):S299–S322, 2013. 2

[44] Siddharth Suri and Sergei Vassilvitskii. Counting triangles and the curse of the last reducer. In *Proceedings of the 20th international conference on World wide web*, pages 607–614, 2011. 2

[45] David Easley, Jon Kleinberg, et al. Networks, crowds, and markets: Reasoning about a highly connected world. *Significance*, 9:43–44, 2012. 2

[46] Marina Meila and Michael I Jordan. Learning with mixtures of trees. *Journal of Machine Learning Research*, 1(Oct):1–48, 2000. 2

[47] Guan-Wei Wang, Chun-Xia Zhang, and Jian Zhuang. Clustering with prim's sequential representation of minimum spanning tree. *Applied Mathematics and Computation*, 247:521–534, 2014. 2

[48] Edward C Kirby, Roger B Mallion, Paul Pollak, and Paweł J Skrzyński. What kirchhoff actually did concerning spanning trees in electrical networks and its relationship to modern graph-theoretical work. *Croatica Chemica Acta*, 89(4):403–417, 2016. 2

[49] Vedat Levi Alev, Nima Anari, Lap Chi Lau, and Shayan Oveis Gharan. Graph clustering using effective resistance. *arXiv preprint arXiv:1711.06530*, 2017. 2

[50] Santo Fortunato. Community detection in graphs. *Physics reports*, 486(3-5):75–174, 2010. 2

[51] Tasweer Ahmad, Lianwen Jin, Luojun Lin, and GuoZhi Tang. Skeleton-based action recognition using sparse spatio-temporal gcn with edge effective resistance. *Neurocomputing*, 423:389–398, 2021. 2

[52] Leo Katz. A new status index derived from sociometric analysis. *Psychometrika*, 18(1):39–43, 1953. 2

[53] Daniel A Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 563–568, 2008. 2