Time-Efficient Quantum Entropy Estimator via Samplizer

Qisheng Wang
Graduate School of Mathematics, Nagoya University, Japan

Zhicheng Zhang
Centre for Quantum Software and Information, University of Technology Sydney, Ultimo, Australia

Abstract

Entropy is a measure of the randomness of a system. Estimating the entropy of a quantum state is a basic problem in quantum information. In this paper, we introduce a time-efficient quantum approach to estimating the von Neumann entropy $S(\rho)$ and Rényi entropy $S_\alpha(\rho)$ of an N-dimensional quantum state ρ, given access to independent samples of ρ. Specifically, we provide the following quantum estimators.

- A quantum estimator for $S(\rho)$ with time complexity $\widetilde{O}(N^2)$,[1] improving the prior best time complexity $\widetilde{O}(N^6)$ by Acharya, Issa, Shende, and Wagner (2020) and Bavarian, Mehraba, and Wright (2016).
- A quantum estimator for $S_\alpha(\rho)$ with time complexity $\widetilde{O}(N^{4/\alpha-2})$ for $0 < \alpha < 1$ and $\widetilde{O}(N^{4-2/\alpha})$ for $\alpha > 1$, improving the prior best time complexity $\widetilde{O}(N^{6/\alpha})$ for $0 < \alpha < 1$ and $\widetilde{O}(N^6)$ for $\alpha > 1$ by Acharya, Issa, Shende, and Wagner (2020), though at a cost of a slightly larger sample complexity.

Moreover, these estimators are naturally extensible to the low-rank case. We also provide a sample lower bound $\Omega(\max\{N/\varepsilon, N^{1/\alpha-1}/\varepsilon^{1/\alpha}\})$ for estimating $S_\alpha(\rho)$.

Technically, our method is quite different from the previous ones that are based on weak Schur sampling and Young diagrams. At the heart of our construction, is a novel tool called *samplizer*, which can "samplize" a quantum query algorithm to a quantum algorithm with similar behavior using only samples of quantum states; this suggests a unified framework for estimating quantum entropies. Specifically, when a quantum oracle U block-encodes a mixed quantum state ρ, any quantum query algorithm using Q queries to U can be samplized to a δ-close (in the diamond norm) quantum algorithm using $\widetilde{\Theta}(Q^2/\delta)$ samples of ρ. Moreover, this samplization is proven to be *optimal*, up to a polylogarithmic factor.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory; Mathematics of computing → Information theory; Theory of computation → Algorithm design techniques; Theory of computation → Lower bounds and information complexity

Keywords and phrases Quantum computing, entropy estimation, von Neumann entropy, Rényi entropy, sample complexity

Digital Object Identifier 10.4230/LIPIcs.ESA.2024.101

Related Version *Full Version*: `https://arxiv.org/abs/2401.09947` [50]

Funding *Qisheng Wang*: Supported by the MEXT Quantum Leap Flagship Program (MEXT Q-LEAP) under Grant JPMXS0120319794.
Zhicheng Zhang: Supported by the Sydney Quantum Academy, NSW, Australia, and the Australian Research Council Discovery Project under Grant DP220102059.

Acknowledgements The authors thank John Wright for valuable comments and sharing their results [7] on von Neumann entropy estimation. QW also thanks François Le Gall for helpful discussions.

[1] $\widetilde{O}(\cdot)$ suppresses polylogarithmic factors.

32nd Annual European Symposium on Algorithms (ESA 2024).
Editors: Timothy Chan, Johannes Fischer, John Iacono, and Grzegorz Herman; Article No. 101; pp. 101:1–101:15
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

In quantum information theory, entropy is a basic measure of the randomness of a quantum system (cf. [39, 57, 56]), which can be understood as the quantum generalization of the entropy of a probability distribution. Quantum entropy can be used to quantify important quantum properties, e.g., the compressibility of quantum data [43, 28, 36] and the entanglement of quantum states [26, 33]. As an analog to the classical learning task of probability distributions, a natural question is: how can we learn the entropy of a quantum state from its independent samples?

Indeed, this is a real question raised in physics for measuring quantum entanglement, e.g., [15, 22, 27]. Recently, Acharya, Issa, Shende, and Wagner [1] and Bavarian, Mehraba, and Wright [7] proposed sample-efficient quantum algorithms for estimating quantum entropy based on the Empirical Young Diagram (EYD) algorithm [4, 30]. Their algorithms, however, have a large time complexity that is cubic in the sample complexity (i.e., the number of independent samples used in the algorithm), due to the use of weak Schur sampling.[2] By stark contrast, classically estimating the entropy of a probability distribution only takes time linear in the sample complexity [45]. Regarding these, one may ask:

Can we estimate quantum entropy with time complexity linear in the sample complexity?

This is not solely a theoretical question: a time-efficient approach to estimating quantum entropy will benefit many practical applications, e.g., preparing quantum Gibbs states [59, 12, 53] and learning Hamiltonians [5].

2 Main results

We introduce a new quantum approach to estimating the entropy of a quantum state, which takes time *linear* in the sample complexity. For a quantum algorithm[3] that only takes independent samples of a quantum state as input (this input model is called *sample access*), the sample complexity is the number of samples used in the algorithm, and the time complexity is the sum of the number of one- and two-qubit quantum gates and the number of one-qubit measurements in its circuit description.

We will state the sample and time complexity of our von Neumann entropy estimator and Rényi entropy estimator in Section 2.1 and Section 2.2, respectively. In comparison with the additive error ε, we are more interested in the dependence on the size of the input quantum state. For simplicity, we assume constant additive error $\varepsilon = \Theta(1)$ in this section, even though our quantum algorithms are polynomially scalable as $1/\varepsilon$ increases.

In Table 1, we summarize our entropy estimators and compare them with prior best approaches. There are also other approaches for estimating the entropy of a quantum state in the literature, which assume access to the quantum circuit that prepares the purification of ρ (this input model is called *purified quantum query access*), thus very different from our setting that only allows access to independent samples of ρ. This line of work will be reviewed in Section 5.

[2] The quantum algorithms proposed in both [1] and [7] are based on the weak Schur sampling [11] (cf. [38]), so (as noted by [58]) they have quantum time complexity $\widetilde{O}(n^3)$ on input n independent samples of a quantum state, using the current best quantum Fourier transform over symmetric groups [29].

[3] In this paper, we only consider *uniform* quantum algorithms. That is, there is a polynomial-time classical Turing machine that, on input 1^n, outputs the circuit description of the quantum algorithm for the problem of size n.

Table 1 Sample and time complexities for entropy estimation of quantum states.

	Reference	$0 < \alpha < 1$	$\alpha = 1$ (von Neumann)	$\alpha > 1$
Upper Bounds	[1]	$O(N^{2/\alpha})$ samples $\widetilde{O}(N^{6/\alpha})$ time	$O(N^2)$ samples $\widetilde{O}(N^6)$ time	$O(N^2)$ samples $\widetilde{O}(N^6)$ time
	This work	$\widetilde{O}(N^{4/\alpha-2})$ samples $\widetilde{O}(N^{4/\alpha-2})$ time Theorem 2	$\widetilde{O}(N^2)$ samples $\widetilde{O}(N^2)$ time Theorem 1	$\widetilde{O}(N^{4-2/\alpha})$ samples $\widetilde{O}(N^{4-2/\alpha})$ time Theorem 2
Lower Bounds	[1]	$\Omega(N^{1+1/\alpha})$ samples (EYD)	$\Omega(N^2)$ samples (EYD)	$\Omega(N^2)$ samples (EYD)
	This work	$\Omega(N + N^{1/\alpha-1})$ samples Theorem 9	$\Omega(N)$ samples Theorem 9	$\Omega(N)$ samples Theorem 9

(EYD) – These lower bounds are for Empirical Young Diagram algorithms.

2.1 Von Neumann entropy estimator

Our first result is a time-efficient estimator for von Neumann entropy, defined by (cf. [46])

$$S(\rho) = -\operatorname{tr}(\rho \ln(\rho)).$$

▶ **Theorem 1.** *There is a quantum estimator for the von Neumann entropy $S(\rho)$ of an N-dimensional quantum state ρ with sample and time complexity $\widetilde{O}(N^2)$.*

The prior best quantum estimators for the von Neumann entropy [1, 7] have sample complexity $O(N^2)$ and time complexity $\widetilde{O}(N^6)$.[4] Our estimator is cubicly faster than theirs in the time complexity, while with the same sample complexity (up to a logarithmic factor). Technically, our method is quite different from the previous ones based on weak Schur sampling and Young diagrams. By comparison, our algorithm builds on our new tool – *samplizer* (which will be introduced in Section 3.1), together with the block-encoding techniques (cf. [19]).

Our von Neumann entropy estimator has an advantage in that it can exploit prior knowledge of a relatively low rank r of the quantum state ρ. In this case, our von Neumann entropy estimator has time complexity $\widetilde{O}(r^2)$, which is polynomial in r while only polylogarithmic in N. Note that the work of [1] does not consider the low-rank case. Recently, a von Neumann entropy estimator was proposed in [54] with sample complexity $\widetilde{O}(\kappa^2)$, where κ is the reciprocal of the minimum non-zero eigenvalue of ρ. The rank-dependent version of our algorithm immediately reproduces their result by noting that κ is always an upper bound on the rank r of ρ.

2.2 Rényi entropy estimator

We also provide time-efficient estimators for α-Rényi entropy, defined by (cf. [42])

$$S_\alpha(\rho) = \frac{1}{1-\alpha} \ln(\operatorname{tr}(\rho^\alpha)),$$

with von Neumann entropy a limiting case: $S(\rho) = S_1(\rho)$.

[4] See Footnote 2.

▶ **Theorem 2.** *There is a quantum estimator for the α-Rényi entropy $S_\alpha(\rho)$ of an N-dimensional quantum state ρ with sample and time complexity $\widetilde{O}(N^{4/\alpha-2})$ for $0 < \alpha < 1$ and $\widetilde{O}(N^{4-2/\alpha})$ for $\alpha > 1$.*

The prior best quantum estimators for the α-Rényi entropy [1] have sample complexity $O(N^{2/\alpha})$ and time complexity $\widetilde{O}(N^{6/\alpha})$ for $\alpha < 1$, and sample complexity $O(N^2)$ and time complexity $\widetilde{O}(N^6)$ for $\alpha > 1$.[5] By comparison, our estimators for the α-Rényi entropy and von Neumann entropy are faster (in time) than the approaches of [1] for any constant $\alpha > 0$.[6] It can be seen that there is a trade-off between the sample and time complexities: our algorithms are more time-efficient, while the approaches of [1] are more sample-efficient. Like our von Neumann entropy estimator, our Rényi entropy estimator is also extensible to the low-rank case, resulting in a time complexity polynomial in the rank r of quantum state ρ.

3 Techniques

The design of our quantum algorithms is based on a novel tool – *samplizer*. Roughly speaking, the samplizer allows us to design a quantum algorithm with access to samples of quantum states by instead designing a quantum query algorithm (namely, a quantum algorithm with access to a quantum unitary oracle). We first introduce the samplizer in Section 3.1 and then show how to design our quantum entropy estimators using the samplizer in Section 3.2 (for von Neumann entropy) and in Section 3.3 (for Rényi entropy).

3.1 Samplizer

Throughout this paper, we use the following concepts and notations for quantum query algorithms and quantum sample algorithms. A *quantum query algorithm* with query access to oracle U is described by a quantum circuit family $C = \{C[U]\}$, where $C[U]$ can use queries to (controlled-)U and (controlled-)$U^\dagger$. A *quantum sample algorithm* with sample access to state ρ is described by a quantum channel family $\mathcal{E} = \{\mathcal{E}[\rho]\}$, where $\mathcal{E}[\rho]$ is implemented by a quantum circuit with ancilla input state of the form $\rho^{\otimes k} \otimes |0\rangle\langle 0|^{\otimes \ell}$. We will use $\mathcal{C}[U]$ to denote the quantum channel $\mathcal{C}[U](\varrho) = C[U]\varrho C[U]^\dagger$ induced by $C[U]$. To justify the concepts defined here, we note that any quantum entropy estimator using independent samples of quantum states is indeed a quantum sample algorithm.

Now we are able to introduce the notion of samplizer.

▶ **Definition 3** (Samplizer). *A samplizer $\mathsf{Samplize}_*\langle * \rangle$ is a converter from a quantum circuit family to a quantum channel family with the following property: for any $\delta > 0$, quantum circuit family $C = \{C[U]\}$, and quantum state ρ, there exists a unitary operator U_ρ that is a block-encoding of $\rho/2$ such that*[7]

$$\|\mathsf{Samplize}_\delta\langle C \rangle[\rho] - \mathcal{C}[U_\rho]\|_\diamond \leq \delta,$$

where $\|\cdot\|_\diamond$ denotes the diamond norm between quantum channels. Here, U is a block-encoding of A if the matrix A is in the upper left corner in the matrix representation of U.

---

[5] See Footnote 2.

[6] For integer $\alpha > 1$, the approach of [1] has sample complexity $O(N^{2-2/\alpha})$ and time complexity $\widetilde{O}(N^{6-6/\alpha})$. Our algorithm is faster with the only exception that $\alpha = 2$. To address this special case, we provide a simple algorithm for estimating the 2-Rényi entropy $S_2(\rho)$ with sample and time complexity $\widetilde{O}(N^2)$ via the SWAP test [6, 10] in the full version of this paper [50].

[7] The scaling factor $1/2$ is due to technical reasons.

The definition of samplizer is inspired by existing quantum query algorithms wherein the output depends only on the matrix block-encoded in the oracle, e.g., the quantum algorithms for Hamiltonian simulation and quantum Gibbs sampling in [19] and solving systems of linear equations in [13]. For any such quantum query algorithm C, we can use the samplizer to construct a quantum sample algorithm $\mathsf{Samplize}_\delta\langle C\rangle$ that simulates the behavior of C when the density matrix of a (mixed) quantum state is block-encoded in the oracle. The existence of the samplizer will allow us to design quantum sample algorithms by just designing quantum query algorithms instead. In the following, we provide an efficient samplizer, demonstrating its existence.

▶ **Theorem 4** (Optimal samplizer). *There is an optimal samplizer* $\mathsf{Samplize}_*\langle * \rangle$ *such that for any* $\delta > 0$ *and quantum query algorithm* C *with query complexity* Q*, the quantum sample algorithm* $\mathsf{Samplize}_\delta\langle C\rangle$[8] *has sample complexity* $\mathsf{S} = \widetilde{\Theta}(Q^2/\delta)$ *and incurs an extra time complexity of* $O(n\mathsf{S})$ *over* C *if the quantum oracle of* C *acts on* n *qubits.*

We design our samplizer based on quantum principal component analysis [35, 31]. The idea is inspired by the recent quantum algorithms for estimating fidelity [18] and trace distance [49]. These algorithms can be employed to construct a quantum circuit that (approximately) block-encodes a quantum state given its independent samples, using quantum singular value transformation [19]. Based on this idea, a lifting theorem was discovered in [48] that relates quantum sample complexity to quantum query complexity. In this paper, we further extend this technique to general quantum query algorithms. This is done by replacing each oracle query with a quantum channel that simulates the oracle that is implemented by (samples of) the quantum state block-encoded in the oracle (if applicable). After the replacement, we obtain a quantum sample algorithm that simulates the original quantum query algorithm.

We prove the optimality of the samplizer by observing that any samplizer can samplize a quantum query algorithm for Hamiltonian simulation (e.g., [19, 37]) to a quantum sample algorithm for sample-based Hamiltonian simulation [35]. Then, we use the quantum sample lower bound for the latter problem [31] to derive a matching lower bound for the samplizer.

▶ Remark 5. Our samplizer studies the sample complexity of simulating quantum query algorithms, which is a generalization of [48, Theorem 1.1] (for Q-dependence only) and [18, Corollary 21] (for δ-dependence only). In [48], they showed a tight Q-dependence but did not consider the dependence on the overall error δ in the sample/time complexity (they only consider the case when δ is a constant). The δ-dependence is extremely important, as the time complexity of the samplizer grows polynomially in $1/\delta$. In [18], they did not consider the Q-dependence and did not show the optimality of the δ-dependence. In Theorem 4, we show matching upper and lower bounds with respect to both Q and δ.

3.2 Von Neumann entropy estimator

Now we present a quantum estimator for the von Neumann entropy in Algorithm 1 through the samplizer provided in Theorem 4. As demonstrated, the samplizer is convenient and useful in designing quantum sample algorithms in a modular fashion.

[8] Our samplizer is uniform. That is, there is a polynomial-time deterministic Turing machine that, on input the description of quantum circuit family $C = \{C[U]\}$ and the unary encodings of Q and δ, outputs the quantum circuit description of the implementation of the quantum channel family $\mathsf{Samplize}_\delta\langle C\rangle$.

Algorithm 1 `estimate_von_Neumann_main`(ε, δ) – *quantum sample algorithm.*

Resources: Access to independent samples of N-dimensional quantum state ρ of rank r.
Input: $\varepsilon \in (0,1)$ and $\delta \in (0,1)$.
Output: $\widetilde{S}$ such that $|\widetilde{S} - S(\rho)| \leq \varepsilon$ with probability $\geq 1-\delta$.

1: **function** `von_Neumann_subroutine`$(\delta_p, \varepsilon_p, \delta_Q)$ – *quantum query algorithm*
   Resources: Unitary oracle U_A that is a block-encoding of A.
2: Let $p(x)$ be a polynomial of degree $d_p = O\left(\frac{1}{\delta_p} \log\left(\frac{1}{\varepsilon_p}\right)\right)$ such that $|p(x)| \leq \frac{1}{2}$ for $x \in [-1,1]$ and $\left|p(x) - \frac{\ln(1/x)}{4\ln(2/\delta_p)}\right| \leq \varepsilon_p$ for $x \in [\delta_p, 1]$ (by the polynomial approximation of logarithms [17, Lemma 11]).
3: Construct unitary operator $U_{p(A)}$ that is a $(1, a, \delta_Q)$-block-encoding of $p(A)$, using $O(d_p)$ queries to U_A (by quantum singular value transformation [19, Theorem 31]).
4: **return** $U_{p(A)}$.
5: **end function**
6: $\delta_p \leftarrow \frac{\varepsilon}{128r\ln(32r/\varepsilon)}$, $\varepsilon_p \leftarrow \frac{\varepsilon}{32\ln(2/\delta_p)}$, $\delta_Q \leftarrow \frac{\varepsilon}{32r\ln(2/\delta_p)}$, $\delta_a \leftarrow \frac{\varepsilon}{64\ln(2/\delta_p)}$, $\varepsilon_H \leftarrow \delta_a$, $k \leftarrow \left\lceil \frac{1}{2\varepsilon_H^2} \ln\left(\frac{2}{\delta}\right) \right\rceil$.
7: **for** $i = 1 \ldots k$ **do**
8: Perform the Hadamard test on $\mathsf{Samplize}_{\delta_a}\langle$`von_Neumann_subroutine`$(\delta_p, \varepsilon_p, \delta_Q)\rangle[\rho]$ and ρ (by the Hadamard test [18, Lemma 9]). Let $X_i \in \{0,1\}$ be the outcome.
9: **end for**
10: $\widetilde{S} \leftarrow 4(2\sum_{i\in[k]} X_i/k - 1)\ln(2/\delta_p) - \ln(2)$.
11: **return** $\widetilde{S}$.

The framework of our quantum estimator for the von Neumann entropy is inspired by the quantum query algorithm in [47] for estimating the von Neumann entropy. In Algorithm 1, we first design a quantum query algorithm

$$\texttt{von\_Neumann\_subroutine}(\delta_p, \varepsilon_p, \delta_Q)[U_A] = U_{p(A)},$$

which implements a block-encoding $U_{p(A)}$ of $p(A)$, using queries to a block-encoding U_A of A, where $p(\cdot)$ is a polynomial defined in Line 2 of Algorithm 1 that approximates the logarithm (up to some constant factor) in certain regime. If `von_Neumann_subroutine`$(\delta_p, \varepsilon_p, \delta_Q)[U_\rho]$ can be implemented as desired for every quantum state ρ, then we can estimate the von Neumann entropy through the Hadamard test [3]. To see this, we provide the following lemma.

▶ **Lemma 6.** *Suppose that U_ρ is a block-encoding of $\rho/2$ where ρ is a quantum state of rank r. Let random variable $X \in \{0,1\}$ be the output of the Hadamard test (as in Line 8 of Algorithm 1) on the unitary operator* `von_Neumann_subroutine`$(\delta_p, \varepsilon_p, \delta_Q)[U_\rho]$ *and the quantum state ρ. Then,*

$$\left|\left(4(2\mathbb{E}[X]-1)\ln\left(\frac{2}{\delta_p}\right) - \ln(2)\right) - S(\rho)\right| \leq 4(2r\delta_p + \varepsilon_p + r\delta_Q)\ln\left(\frac{2}{\delta_p}\right).$$

Using the samplizer provided in Theorem 4, we are able to construct its "samplized" version

$$\mathsf{Samplize}_{\delta_a}\langle\texttt{von\_Neumann\_subroutine}(\delta_p, \varepsilon_p, \delta_Q)\rangle[\rho],$$

which only uses independent samples of the input quantum state ρ. Let $X' \in \{0,1\}$ be the output of the Hadamard test on $\mathsf{Samplize}_{\delta_a}\langle$`von_Neumann_subroutine`$(\delta_p, \varepsilon_p, \delta_Q)\rangle[\rho]$ and the quantum state ρ, as analogous to Lemma 6. It can be shown that $|\mathbb{E}[X'] - \mathbb{E}[X]| \le \delta_a$, which implies that

$$\left|\left(4(2\mathbb{E}[X'] - 1)\ln\left(\frac{2}{\delta_p}\right) - \ln(2)\right) - S(\rho)\right| \le 4(2r\delta_p + \varepsilon_p + r\delta_Q + 2\delta_a)\ln\left(\frac{2}{\delta_p}\right).$$

Therefore, once an estimate p of $\mathbb{E}[X']$ is obtained, we can use $4(2p-1)\ln(2/\delta_p) - \ln(2)$ as an estimate of $S(\rho)$. By choosing appropriate values for the parameters such as $\delta_p, \varepsilon_p, \delta_Q, \delta_a, \varepsilon_H, k$ as in Algorithm 1, we can obtain an ε-estimate of the von Neumann entropy $S(\rho)$ with sample and time complexity $\widetilde{O}(r^2/\varepsilon^5)$.

3.3 Rényi entropy estimator

We also provide quantum estimators for the α-Rényi entropy for every $\alpha \in (0,1) \cup (1,+\infty)$ through the samplizer provided in Theorem 4. As an illustrative example, we mainly introduce the estimator for $\alpha > 1$ in Algorithm 2. The idea for $0 < \alpha < 1$ is similar, which is presented in Algorithm 3. The framework of our quantum estimators for the Rényi entropy of quantum states is recursive, which is inspired by the quantum query algorithm in [34] for estimating the Rényi entropy of discrete probability distributions. We denote $P_\alpha(\rho) = \operatorname{tr}(\rho^\alpha)$.

3.3.1 $\alpha > 1$

In Algorithm 2, two main functions are explained as follows.

- `estimate_Rényi_gt1`$(\alpha, \varepsilon, \delta)$: return an estimate $\widetilde{P}$ such that $(1-\varepsilon)\widetilde{P} \le P_\alpha(\rho) \le (1+\varepsilon)\widetilde{P}$ with probability $\ge 1-\delta$.
- `estimate_Rényi_gt1_promise`$(\alpha, P, \varepsilon, \delta)$: return an estimate $\widetilde{P}$ such that $(1-\varepsilon)\widetilde{P} \le P_\alpha(\rho) \le (1+\varepsilon)\widetilde{P}$ with probability $\ge 1-\delta$, given a promise that $P \le P_\alpha(\rho) \le 10P$.

It can be seen that by letting $\widetilde{P} \leftarrow$ `estimate_Rényi_gt1`$(\alpha, (\alpha-1)\varepsilon/2, \delta)$ as in Line 31 of Algorithm 2, $\widetilde{S} \leftarrow \frac{1}{1-\alpha}\ln(\widetilde{P})$ is then an ε-estimate of $S_\alpha(\rho)$.

The main observation is that `estimate_Rényi_gt1`$(\alpha, \varepsilon, \delta)$ can be computed recursively. This is done by two steps:

1. With probability $\ge 1-\delta/2$, obtain an estimate P such that $P \le P_\alpha(\rho) \le 10P$. This is done by reducing to another entropy estimation task with smaller α as in Line 26 of Algorithm 2.
2. With probability $\ge 1-\delta/2$, obtain an estimate $\widetilde{P}$ such that $(1-\varepsilon)\widetilde{P} \le P_\alpha(\rho) \le (1+\varepsilon)\widetilde{P}$ by calling `estimate_Rényi_gt1_promise`$(\alpha, P, \varepsilon, \delta/2)$ as in Line 29 of Algorithm 2.

To implement the function `estimate_Rényi_gt1_promise`$(\alpha, P, \varepsilon, \delta)$, we first design a quantum query algorithm

`Rényi_gt1_subroutine`$(\alpha, P, \delta_p, \varepsilon_p, \delta_Q)[U_A] = U_{p(A)}$,

which implements a block-encoding $U_{p(A)}$ of $p(A)$, using queries to a block-encoding U_A of A, where $p(\cdot)$ is a polynomial defined in Line 3 of Algorithm 2 that approximates the positive power function (up to some constant factor). If `Rényi_gt1_subroutine`$(\alpha, P, \delta_p, \varepsilon_p, \delta_Q)[U_\rho]$ can be implemented as desired for every quantum state ρ, then we can estimate $P_\alpha(\rho)$ by applying it on $\rho \otimes |0\rangle\langle 0|^{\otimes a}$. To see this, we provide the following lemma.

▶ **Lemma 7.** *Suppose that U_ρ is a block-encoding of $\rho/2$ where ρ is a quantum state of rank r. Let $U_{p(\rho/2)} =$ `Rényi_gt1_subroutine`$(\alpha, P, \delta_p, \varepsilon_p, \delta_Q)[U_\rho]$. Let random variable $X = 1$ if the measurement outcome of $U_{p(\rho/2)}(\rho \otimes |0\rangle\langle 0|^{\otimes a})U_{p(\rho/2)}^\dagger$ in the computational basis (on the last a qubits) is $|0\rangle^{\otimes a}$, and 0 otherwise (as in Line 14 of Algorithm 2). Then,*

$$\left|16(4\beta)^{\alpha-1}\mathbb{E}[X] - P_\alpha(\rho)\right| \le 5r(2\delta_p)^\alpha + 32(4\beta)^{\alpha-1}(\varepsilon_p + r\delta_Q).$$

Algorithm 2 `estimate_Rényi_gt1_main`$(\alpha, \varepsilon, \delta)$ – *quantum sample algorithm.*

Resources: Access to independent samples of N-dimensional quantum state ρ of rank r.
Input: $\alpha > 1$, $\varepsilon \in (0,1)$, and $\delta \in (0,1)$.
Output: $\widetilde{S}$ such that $|\widetilde{S} - S_\alpha(\rho)| \le \varepsilon$ with probability $\ge 1-\delta$.

function `Rényi_gt1_subroutine`$(\alpha, P, \delta_p, \varepsilon_p, \delta_Q)$ – *quantum query algorithm*
  Resources: Unitary oracle U_A that is a block-encoding of A.
  $\beta \leftarrow \min\{(10P)^{1/\alpha}, 1/2\}$, $c \leftarrow (\alpha-1)/2$.
  Let $p(x)$ be a polynomial of degree $d_p = O\left(\frac{1}{\delta_p}\log\left(\frac{1}{\delta_p\varepsilon_p}\right)\right)$ such that $|p(x)| \le \frac{1}{2}\left(\frac{\delta_p}{2\beta}\right)^c$ for $x \in [0, \delta_p]$, $\left|p(x) - \frac{1}{4}\left(\frac{x}{2\beta}\right)^c\right| \le \varepsilon_p$ for $x \in [\delta_p, \beta]$, and $|p(x)| \le \frac{1}{2}$ for $x \in [-1,1]$ (by the polynomial approximation of positive power functions, e.g., [52, Lemma 6]).
  Construct unitary operator $U_{p(A)}$ that is a $(1, a, \delta_Q)$-block-encoding of $p(A)$, using $O(d_p)$ queries to U_A (by quantum singular value transformation [19, Theorem 31]).
  return $U_{p(A)}$.
end function
function `estimate_Rényi_gt1_promise`$(\alpha, P, \varepsilon, \delta)$
  $\beta \leftarrow \min\{(10P)^{1/\alpha}, 1/2\}$, $m \leftarrow \lceil 8\ln(1/\delta)\rceil$, $\delta_p \leftarrow \frac{1}{2}\left(\frac{P\varepsilon}{40r}\right)^{1/\alpha}$.
  $\varepsilon_p \leftarrow \frac{(4\beta)^{1-\alpha}P\varepsilon}{256}$, $\delta_Q \leftarrow \frac{(4\beta)^{1-\alpha}P\varepsilon}{128r}$, $\delta_a \leftarrow \frac{(4\beta)^{1-\alpha}P\varepsilon}{128}$, and $k \leftarrow \left\lceil \frac{65536}{(4\beta)^{1-\alpha}P\varepsilon^2}\right\rceil$.
  for $j = 1 \dots m$ **do**
    for $i = 1 \dots k$ **do**
      Let $\sigma = \mathsf{Samplize}_{\delta_a}\langle$`Rényi_gt1_subroutine`$(\alpha, P, \delta_p, \varepsilon_p, \delta_Q)\rangle[\rho](\rho \otimes |0\rangle\langle 0|^{\otimes a})$.
      Measure σ in the computational basis.
      Let X_i be 1 if the outcome is $|0\rangle^{\otimes a}$, and 0 otherwise.
    end for
    $\hat{P}_j \leftarrow 16(4\beta)^{\alpha-1}\sum_{i\in[k]} X_i/k$.
  end for
  $\widetilde{P} \leftarrow$ the median of $\hat{P}_j$ for $j \in [m]$.
  return $\widetilde{P}$.
end function
function `estimate_Rényi_gt1`$(\alpha, \varepsilon, \delta)$
  $\lambda \leftarrow 1 + 1/\ln(r)$.
  if $\alpha \le \lambda$ **then**
    $P \leftarrow e^{-1}$.
  else
    $P' \leftarrow$ `estimate_Rényi_gt1`$(\alpha/\lambda, 1/4, \delta/2)$.
    $P \leftarrow (4P'/5)^\lambda e^{-1}$.
  end if
  return `estimate_Rényi_gt1_promise`$(\alpha, P, \varepsilon, \delta/2)$.
end function
$\widetilde{P} \leftarrow$ `estimate_Rényi_gt1`$(\alpha, (\alpha-1)\varepsilon/2, \delta)$.
$\widetilde{S} \leftarrow \frac{1}{1-\alpha}\ln(\widetilde{P})$.
return $\widetilde{S}$.

Using the samplizer provided in Theorem 4, we are able to construct its "samplized" version

$$\mathsf{Samplize}_{\delta_a}\langle\texttt{Rényi\_gt1\_subroutine}(\alpha, P, \delta_p, \varepsilon_p, \delta_Q)\rangle[\rho],$$

which only uses independent samples of the input quantum state ρ. Let random variable $X' = 1$ if the measurement outcome of $\mathsf{Samplize}_{\delta_a}\langle\texttt{Rényi\_gt1\_subroutine}(\alpha, P, \delta_p, \varepsilon_p, \delta_Q)\rangle[\rho](\rho \otimes |0\rangle\langle 0|^{\otimes a})$ in the computational basis (on the last a qubits) is $|0\rangle^{\otimes a}$, and $X' = 0$ otherwise, as analogous to Lemma 6. It can be shown that $|\mathbb{E}[X'] - \mathbb{E}[X]| \leq \delta_a$, which implies that

$$\left|16(4\beta)^{\alpha-1}\mathbb{E}[X'] - P_\alpha(\rho)\right| \leq 5r(2\delta_p)^\alpha + 16(4\beta)^{\alpha-1}(2\varepsilon_p + 2r\delta_Q + \delta_a).$$

Therefore, once an estimate p of $\mathbb{E}[X']$ is obtained, we can use $16(4\beta)^{\alpha-1}p$ as an estimate of $P_\alpha(\rho)$. By choosing appropriate values for the parameters such as $\delta_p, \varepsilon_p, \delta_Q, \delta_a, k$ as in Algorithm 2, we can obtain an ε-estimate of the Rényi entropy $S_\alpha(\rho)$ with sample and time complexity $\widetilde{O}(r^{4-2/\alpha}/\varepsilon^{3+2/\alpha})$.

3.3.2 $0 < \alpha < 1$

Although the structure and the analysis of Algorithm 3 are similar to those of Algorithm 2, we introduce them here for completeness and for noting the differences in detail. The two main functions are explained as follows.

- `estimate_Rényi_lt1`$(\alpha, \varepsilon, \delta)$: return an estimate $\widetilde{P}$ such that $(1-\varepsilon)P_\alpha(\rho) \leq \widetilde{P} \leq (1+\varepsilon)P_\alpha(\rho)$ with probability $\geq 1-\delta$.
- `estimate_Rényi_lt1_promise`$(\alpha, P, \varepsilon, \delta)$: return an estimate $\widetilde{P}$ such that $(1-\varepsilon)P_\alpha(\rho) \leq \widetilde{P} \leq (1+\varepsilon)P_\alpha(\rho)$ with probability $\geq 1-\delta$, given a promise that $P \leq P_\alpha(\rho) \leq 10P$.

The key part is the implementation of the function `estimate_Rényi_lt1_promise`$(\alpha, P, \varepsilon, \delta)$. To this end, we first design a quantum query algorithm

$$\texttt{Rényi\_lt1\_subroutine}(\alpha, P, \delta_p, \varepsilon_p, \delta_Q)[U_A] = U_{p(A)},$$

which implements a block-encoding $U_{p(A)}$ of $p(A)$, using queries to a block-encoding U_A of A, where $p(\cdot)$ is a polynomial defined in Line 2 of Algorithm 3 that approximates the negative power function (up to some constant factor). Similar to the analysis for $\alpha > 1$, if one can implement `Rényi_lt1_subroutine`$(\alpha, P, \delta_p, \varepsilon_p, \delta_Q)[U_\rho]$ for every quantum state ρ, then we can estimate $P_\alpha(\rho)$ by applying it on $\rho \otimes |0\rangle\langle 0|^{\otimes a}$. To see this, we provide the following lemma.

▶ **Lemma 8.** *Suppose that U_ρ is a block-encoding of $\rho/2$ where ρ is a quantum state of rank r. Let $U_{p(\rho/2)} =$ `Rényi_lt1_subroutine`$(\alpha, P, \delta_p, \varepsilon_p, \delta_Q)[U_\rho]$. Let random variable $X = 1$ if the measurement outcome of $U_{p(\rho/2)}(\rho \otimes |0\rangle\langle 0|^{\otimes a})U^\dagger_{p(\rho/2)}$ in the computational basis (on the last a qubits) is $|0\rangle^{\otimes a}$, and 0 otherwise (as in Line 13 of Algorithm 3). Then,*

$$\left|16(2\delta_p)^{\alpha-1}\mathbb{E}[X] - P_\alpha(\rho)\right| \leq 5r(2\delta_p)^\alpha + 32(2\delta_p)^{\alpha-1}(\varepsilon_p + r\delta_Q).$$

Using the samplizer provided in Theorem 4, we are able to construct its "samplized" version

$$\mathsf{Samplize}_{\delta_a}\langle\texttt{Rényi\_lt1\_subroutine}(\alpha, P, \delta_p, \varepsilon_p, \delta_Q)\rangle[\rho],$$

which only uses independent samples of the input quantum state ρ. Let random variable $X' = 1$ if the measurement outcome of $\mathsf{Samplize}_{\delta_a}\langle\texttt{Rényi\_lt1\_subroutine}(\alpha, P, \delta_p, \varepsilon_p, \delta_Q)\rangle[\rho](\rho \otimes$

Algorithm 3 `estimate_Rényi_lt1_main`$(\alpha, \varepsilon, \delta)$ – *quantum sample algorithm.*

Resources: Access to independent samples of N-dimensional quantum state ρ of rank r.
Input: $0 < \alpha < 1$, $\varepsilon \in (0, 1)$, and $\delta \in (0, 1)$.
Output: $\widetilde{S}$ such that $|\widetilde{S} - S_\alpha(\rho)| \leq \varepsilon$ with probability $\geq 1 - \delta$.

function `Rényi_lt1_subroutine`$(\alpha, P, \delta_p, \varepsilon_p, \delta_Q)$ – *quantum query algorithm*
Resources: Unitary oracle U_A that is a block-encoding of A.
Let $p(x)$ be a polynomial of degree $d_p = O\left(\frac{1}{\delta_p} \log\left(\frac{1}{\varepsilon_p}\right)\right)$ such that $\left| p(x) - \frac{1}{4}\left(\frac{x}{\delta_p}\right)^{\frac{\alpha-1}{2}} \right| \leq \varepsilon_p$ for $x \in [\delta_p, 1]$, and $|p(x)| \leq \frac{1}{2}$ for $x \in [-1, 1]$ (by the polynomial approximation of negative power functions [19, Corollary 67 in the full version]).
Construct unitary operator $U_{p(A)}$ that is a $(1, a, \delta_Q)$-block-encoding of $p(\rho)$, using $O(d_p)$ queries to U_A (by quantum singular value transformation [19, Theorem 31]).
return $U_{p(A)}$.
end function
function `estimate_Rényi_lt1_promise`$(\alpha, P, \varepsilon, \delta)$
$m \leftarrow \lceil 8 \ln(1/\delta) \rceil$, $\delta_p \leftarrow \frac{1}{2}\left(\frac{P\varepsilon}{40r}\right)^{1/\alpha}$.
$\varepsilon_p \leftarrow \frac{(2\delta_p)^{1-\alpha} P\varepsilon}{256}$, $\delta_Q \leftarrow \frac{(2\delta_p)^{1-\alpha} P\varepsilon}{128r}$, $\delta_a \leftarrow \frac{(2\delta_p)^{1-\alpha} P\varepsilon}{128}$, and $k \leftarrow \left\lceil \frac{65536}{(2\delta_p)^{1-\alpha} P\varepsilon^2} \right\rceil$.
for $j = 1 \ldots m$ **do**
for $i = 1 \ldots k$ **do**
Let $\sigma = \mathsf{Samplize}_{\delta_a}\langle$`Rényi_lt1_subroutine`$(\alpha, P, \delta_p, \varepsilon_p, \delta_Q)\rangle[\rho](\rho \otimes |0\rangle\langle 0|^{\otimes a})$.
Measure σ in the computational basis.
Let X_i be 1 if the outcome is $|0\rangle^{\otimes a}$, and 0 otherwise.
end for
$\hat{P}_j \leftarrow 16(2\delta_p)^{\alpha-1} \sum_{i \in [k]} X_i / k$.
end for
$\widetilde{P} \leftarrow$ the median of $\hat{P}_j$ for $j \in [m]$.
return $\widetilde{P}$.
end function
function `estimate_Rényi_lt1`$(\alpha, \varepsilon, \delta)$
$\lambda \leftarrow 1 - 1/\ln(r)$.
if $\alpha \geq \lambda$ **then**
$P \leftarrow 1$.
else
$P' \leftarrow$ `estimate_Rényi_lt1`$(\alpha/\lambda, 1/4, \delta/2)$.
$P \leftarrow (4P'/5)^\lambda$.
end if
return `estimate_Rényi_lt1_promise`$(\alpha, P, \varepsilon, \delta/2)$.
end function
$\widetilde{P} \leftarrow$ `estimate_Rényi_lt1`$(\alpha, (1-\alpha)\varepsilon/2, \delta)$.
$\widetilde{S} \leftarrow \frac{1}{1-\alpha} \ln(\widetilde{P})$.
return $\widetilde{S}$.

$|0\rangle\langle 0|^{\otimes a}$) in the computational basis (on the last a qubits) is $|0\rangle^{\otimes a}$, and $X' = 0$ otherwise, as analogous to Lemma 8. It can be shown that $|\mathbb{E}[X'] - \mathbb{E}[X]| \leq \delta_a$, which implies that

$$\left|16(2\delta_p)^{\alpha-1}\mathbb{E}[X'] - P_\alpha(\rho)\right| \leq 5r(2\delta_p)^\alpha + 16\delta_p^{\alpha-1}(2\varepsilon_p + 2r\delta_Q + \delta_a).$$

Therefore, once an estimate p of $\mathbb{E}[X']$ is obtained, we can use $16(2\delta_p)^{\alpha-1}p$ as an estimate of $P_\alpha(\rho)$. By choosing appropriate values for the parameters such as $\delta_p, \varepsilon_p, \delta_Q, \delta_a, k$ as in Algorithm 3, we can obtain an ε-estimate of the Rényi entropy $S_\alpha(\rho)$ with sample and time complexity $\widetilde{O}(r^{4/\alpha-2/\alpha}/\varepsilon^{1+4/\alpha})$.

4 Lower bounds

For completeness, we prove lower bounds on the sample complexity for estimating the von Neumann entropy and Rényi entropy.

▶ **Theorem 9.** *For every constant $\alpha > 0$, any quantum estimator for the α-Rényi entropy of an N-dimensional quantum state within additive error ε requires sample complexity $\Omega(\max\{N/\varepsilon, N^{1/\alpha-1}/\varepsilon^{1/\alpha}\})$. In particular, estimating the von Neumann entropy ($\alpha = 1$) requires sample complexity $\Omega(N/\varepsilon)$.*

To the best of our knowledge, we are not aware of any general sample lower bounds for estimating the von Neumann entropy or Rényi entropy that are explicitly stated in the literature. Nevertheless, we note that the sample lower bound for the mixedness testing problem of quantum states given in [41, Theorem 1.10] actually implies an $\Omega(N)$ sample lower bound for entropy estimation. In Theorem 9, our contribution is that we give a better sample lower bound for $0 < \alpha < 1/2$, and that we further consider the ε-dependence in the lower bounds. This is achieved by reducing the task of estimating the α-Rényi entropy of quantum states to the mixed testing problem of quantum states in [41] and to the distinguishing problem of a special probability distribution used in [2, 1].

We note that in [1], they provided sample lower bounds $\Omega(\max\{N^2/\varepsilon, N^{1+1/\alpha}/\varepsilon^{1/\alpha}\})$ for any empirical Young diagram algorithms that estimate the α-Rényi entropy for $\alpha > 0$ (including $\alpha = 1$ for the von Neumann entropy). Compared to the lower bounds given in Theorem 9, their lower bounds do not apply to general algorithms that are not based on empirical Young diagrams (which is noted by [58]).

We discuss the limiting cases $\alpha = 0$ and $\alpha = \infty$ of Theorem 9 as follows.

- For the case of $\alpha = 0$, $S_0(\rho) = \ln(\mathrm{rank}(\rho))$ is the Max (Hartley) entropy. We further show that there is no estimator for the Max entropy (within constant additive error). To see this, consider the problem of distinguishing the two quantum states $\rho_0 = |0\rangle\langle 0|$ and $\rho_\delta = (1-\delta)|0\rangle\langle 0| + \delta \cdot \frac{I}{N}$, where $\delta > 0$ can be arbitrarily close to 0. Note that $\mathrm{rank}(\rho_0) = 1$ and $\mathrm{rank}(\rho_\delta) = N$, and thus $S_0(\rho_0) = 0$ and $S_0(\rho_\delta) = \ln(N)$. On the other hand, according to the upper bound on the success probability of quantum state discrimination [23, 25], distinguishing between ρ_0 and ρ_δ requires $\Omega(1/\delta)$ samples, which can be arbitrarily large and is independent of N.
- For the case of $\alpha = \infty$, $S_\infty(\rho) = -\ln(\|\rho\|)$ is the Min entropy. An estimator with sample complexity $O(N^2/\varepsilon^2)$ is implied by [40, Theorem 1.18].[9] On the other hand, the proof for $\alpha > 1$ also applies to $\alpha = \infty$, thus an $\Omega(N/\varepsilon)$ sample lower bound also holds for estimating $S_\infty(\rho)$.

[9] In [40], they proposed a quantum algorithm that finds the top-k eigenvalues of an N-dimensional quantum state ρ to δ-accuracy in ℓ_2^2 distance with sample complexity $O(k/\delta)$. Note that $\|\rho\|$ is the largest (i.e., top-1) eigenvalue of ρ and $1/N \leq \|\rho\| \leq 1$. To obtain an estimate of $S_\infty(\rho)$ within additive error ε, an estimate of $\|\rho\|$ with multiplicative error ε suffices. This can be done by taking $k = 1$ and $\delta = \varepsilon^2/N^2$, resulting in a sample complexity of $O(N^2/\varepsilon^2)$.

5 Related Work

As not aforementioned, there are quantum query algorithms for estimating the entropy of a quantum state ρ, given purified access to ρ. It was shown in [17] that the von Neumann entropy $S(\rho)$ can be estimated with quantum query complexity $\widetilde{O}(N)$. The estimation of $S(\rho)$ was shown to be useful as a subroutine in variational quantum algorithms [12], where they showed that $S(\rho)$ can be estimated with query complexity $\widetilde{O}(\kappa^2)$, and κ is the reciprocal of the minimum non-zero eigenvalue of ρ. A quantum query algorithm for estimating $S(\rho)$ with multiplicative error was proposed in [20]. It was shown in [44] that the α-Rényi entropy $S_\alpha(\rho)$ can be estimated with quantum query complexity $\widetilde{O}(\kappa N^{\max\{\alpha,1\}})$, which was later shown in [52] to be $\widetilde{O}(N^{1/2+1/2\alpha})$ for $0 < \alpha < 1$ and $\widetilde{O}(N^{3/2-1/2\alpha})$ for $\alpha > 1$. When ρ is low-rank, it was shown in [47] that the quantum query complexity of estimating $S(\rho)$ and $S_\alpha(\rho)$ is $\operatorname{poly}(r)$. Other than upper bounds, it was shown in [16] that estimating the entropy of shallow circuit outputs is hard. In addition to quantum approaches, a classical approach for estimating the von Neumann entropy was proposed in [32]. For probability distributions, quantum algorithms for estimating their entropy were investigated in [34].

6 Discussion

In this paper, we provide time-efficient quantum estimators for the von Neumann entropy and Rényi entropy of quantum states using their independent samples. They are designed under the unified framework of our novel tool – samplizer. Very different from the prior approaches [1, 7] that are based on weak Schur sampling and Young diagrams, our quantum entropy estimators build on the samplizer and quantum singular value transformation, demonstrating that block-encoding techniques [19] are also useful to obtain efficient quantum estimators that take only independent samples of quantum states as input.

We conclude by mentioning several open questions related to our work.

- Can we improve the logarithmic factors in the sample complexity of the samplizer given in Theorem 4? The current upper and lower bounds on the sample complexity of the samplizer are only tight up to polylogarithmic factors.
- All of the existing estimators for the von Neumann entropy, including the estimators based on the EYD (empirical Young diagram) by [1, 7] and ours (Theorem 1), have sample complexity $\widetilde{O}(N^2)$. It was also shown in [1] that any quantum EYD estimator for the von Neumann entropy has sample complexity $\Omega(N^2)$. We conjecture that the same sample lower bound also holds for any von Neumann entropy estimator (that is not necessarily based on the EYD), though we can only prove a lower bound $\Omega(N)$ in Theorem 9.
- Although our Rényi entropy estimator (Theorem 2) is more time-efficient than the estimator proposed in [1], its sample complexity is worse. Can we improve the sample-time tradeoff or prove any sample-time lower bound for Rényi entropy estimators?
- We believe that the samplizer can be useful to design quantum algorithms for quantum property testing, especially for those concerning quantum states. For example, we think that it could be used to simplify the fidelity estimator in [18] and the trace distance estimator in [49]. Except for these direct applications, can we find new quantum sample algorithms for other computational tasks of interest through the samplizer?

References

1 Jayadev Acharya, Ibrahim Issa, Nirmal V. Shende, and Aaron B. Wagner. Estimating quantum entropy. *IEEE Journal on Selected Areas in Information Theory*, 1(2):454–468, 2020. doi:10.1109/JSAIT.2020.3015235.

2 Jayadev Acharya, Alon Orlitsky, Ananda Theertha Suresh, and Himanshu Tyagi. Estimating Renyi entropy of discrete distributions. *IEEE Transactions on Information Theory*, 63(1):38–56, 2017. doi:10.1109/TIT.2016.2620435.

3 Dorit Aharonov, Vaughan Jones, and Zeph Landau. A polynomial quantum algorithm for approximating the Jones polynomial. *Algorithmica*, 55(3):395–421, 2009. doi:10.1007/s00453-008-9168-0.

4 Robert Alicki, Sławomir Rudnicki, and Sławomir Sadowski. Symmetry properties of product states for the system of n n-level atoms. *Journal of Mathematical Physics*, 29(5):1158–1162, 1988. doi:10.1063/1.527958.

5 Anurag Anshu, Srinivasan Arunachalam, Tomotaka Kuwahara, and Mehdi Soleimanifar. Sample-efficient learning of interacting quantum systems. *Nature Physics*, 17(8):931–935, 2021. doi:10.1038/s41567-021-01232-0.

6 Adriano Barenco, André Berthiaume, David Deutsch, Artur Ekert, Richard Jozsa, and Chiara Macchiavello. Stabilization of quantum computations by symmetrization. *SIAM Journal on Computing*, 26(5):1541–1557, 1997. doi:10.1137/S0097539796302452.

7 Mohammad Bavarian, Saeed Mehraban, and John Wright. Learning entropy. A manuscript on von Neumann entropy estimation, private communication, 2016.

8 Christian Beck and Friedrich Schögl. *Thermodynamics of Chaotic Systems: An Introduction*. Cambridge University Press, 1993.

9 Fernando G. S. L. Brandão, Amir Kalev, Tongyang Li, Cedric Yen-Yu Lin, Krysta M. Svore, and Xiaodi Wu. Quantum SDP solvers: large speed-ups, optimality, and applications to quantum learning. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming*, pages 27:1–27:14, 2019. doi:10.4230/LIPIcs.ICALP.2019.27.

10 Harry Buhrman, Richard Cleve, John Watrous, and Ronald de Wolf. Quantum fingerprinting. *Physical Review Letters*, 87(16):167902, 2001. doi:10.1103/PhysRevLett.87.167902.

11 Andrew M. Childs, Aram W. Harrow, and Paweł Wocjan. Weak Fourier-Schur sampling, the hidden subgroup problem, and the quantum collision problem. In *Proceedings of the 24th Annual Symposium on Theoretical Aspects of Computer Science*, pages 598–609, 2007. doi:10.1007/978-3-540-70918-3_51.

12 Anirban N. Chowdhury, Guang Hao Low, and Nathan Wiebe. A variational quantum algorithm for preparing quantum Gibbs states. ArXiv e-prints, 2020. arXiv:2002.00055.

13 Pedro C. S. Costa, Dong An, Yuval R. Sanders, Yuan Su, Ryan Babbush, and Dominic W. Berry. Optimal scaling quantum linear-systems solver via discrete adiabatic theorem. *PRX Quantum*, 3(4):040303, 2022. doi:10.1103/PRXQuantum.3.040303.

14 William Feller. *An Introduction to Probability Theory and Its Applications, Volume 1*. John Wiley & Sons, 1968.

15 F. Franchini, A. R. Its, and V. E. Korepin. Renyi entropy of the XY spin chain. *Journal of Physics A: Mathematical and Theoretical*, 41(2):025302, 2008. doi:10.1088/1751-8113/41/2/025302.

16 Alexandru Gheorghiu and Matty J. Hoban. Estimating the entropy of shallow circuit outputs is hard. ArXiv e-prints, 2020. arXiv:2002.12814.

17 András Gilyén and Tongyang Li. Distributional property testing in a quantum world. In *Proceedings of the 11th Innovations in Theoretical Computer Science Conference*, pages 25:1–25:19, 2020. doi:10.4230/LIPIcs.ITCS.2020.25.

18 András Gilyén and Alexander Poremba. Improved quantum algorithms for fidelity estimation. ArXiv e-prints, 2022. arXiv:2203.15993.

19 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. doi:10.1145/3313276.3316366.

20 Tom Gur, Min-Hsiu Hsieh, and Sathyawageeswar Subramanian. Sublinear quantum algorithms for estimating von Neumann entropy. ArXiv e-prints, 2021. arXiv:2111.11139.

21 Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Physical Review Letters*, 103(15):150502, 2009. doi:10.1103/PhysRevLett.103.150502.

22 Matthew B. Hastings, Iván González, Ann B. Kallin, and Roger G. Melko. Measuring Renyi entanglement entropy in quantum Monte Carlo simulations. *Physical Review Letters*, 104(15):157201, 2010. doi:10.1103/PhysRevLett.104.157201.

23 Carl W. Helstrom. Detection theory and quantum mechanics. *Information and Control*, 10(3):254–291, 1967. doi:10.1016/S0019-9958(67)90302-6.

24 Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. doi:10.1080/01621459.1963.10500830.

25 Alexander S. Holevo. Statistical decision theory for quantum systems. *Journal of Multivariate Analysis*, 3(4):337–394, 1973. doi:10.1016/0047-259X(73)90028-6.

26 Ryszard Horodecki, Paweł Horodecki, Michał Horodecki, and Karol Horodecki. Quantum entanglement. *Reviews of Modern Physics*, 81(2):865, 2009. doi:10.1103/RevModPhys.81.865.

27 Rajibul Islam, Ruichao Ma, Philipp M. Preiss, M. Eric Tai, Alexander Lukin, Matthew Rispoli, and Markus Greiner. Measuring entanglement entropy in a quantum many-body system. *Nature*, 528(7580):77–83, 2015. doi:10.1038/nature15750.

28 Richard Jozsa and Benjamin Schumacher. A new proof of the quantum noiseless coding theorem. *Journal of Modern Optics*, 41(12):2343–2349, 1994. doi:10.1080/09500349414552191.

29 Yasuhito Kawano and Hiroshi Sekigawa. Quantum Fourier transform over symmetric groups — improved result. *Journal of Symbolic Computation*, 75:219–243, 2016. doi:10.1016/j.jsc.2015.11.016.

30 M. Keyl and R. F. Werner. Estimating the spectrum of a density operator. *Physical Review A*, 64(5):052311, 2001. doi:10.1103/PhysRevA.64.052311.

31 Shelby Kimmel, Cedric Yen-Yu Lin, Guang Hao Low, Maris Ozols, and Theodore J. Yoder. Hamiltonian simulation with optimal sample complexity. *npj Quantum Information*, 3(1):1–7, 2017. doi:10.1038/s41534-017-0013-7.

32 Eugenia-Maria Kontopoulou, Gregory-Paul Dexter, Wojciech Szpankowski, Ananth Grama, and Petros Drineas. Randomized linear algebra approaches to estimate the von Neumann entropy of density matrices. *IEEE Transactions on Information Theory*, 66(8):5003–5021, 2020. doi:10.1109/TIT.2020.2971991.

33 Nicolas Laflorencie. Quantum entanglement in condensed matter systems. *Physics Reports*, 646:1–59, 2016. doi:10.1016/j.physrep.2016.06.008.

34 Tongyang Li and Xiaodi Wu. Quantum query complexity of entropy estimation. *IEEE Transactions on Information Theory*, 65(5):2899–2921, 2019. doi:10.1109/TIT.2018.2883306.

35 Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. Quantum principal component analysis. *Nature Physics*, 10(9):631–633, 2014. doi:10.1038/nphys3029.

36 Hoi-Kwong Lo. Quantum coding theorem for mixed states. *Optics Communications*, 119(5-6):552–556, 1995. doi:10.1016/0030-4018(95)00406-X.

37 Guang Hao Low and Isaac L. Chuang. Hamiltonian simulation by qubitization. *Quantum*, 3:163, 2019. doi:10.22331/q-2019-07-12-163.

38 Ashley Montanaro and Ronald de Wolf. A survey of quantum property testing. In *Theory of Computing Library*, number 7 in Graduate Surveys, pages 1–81. University of Chicago, 2016. doi:10.4086/toc.gs.2016.007.

39 Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2010.

40 Ryan O'Donnell and John Wright. Efficient quantum tomography II. In *Proceedings of the 49th Annual ACM Symposium on Theory of Computing*, pages 962–974, 2017. `doi:10.1145/3055399.3055454`.

41 Ryan O'Donnell and John Wright. Quantum spectrum testing. *Communications in Mathematical Physics*, 387(1):1–75, 2021. `doi:10.1007/s00220-021-04180-1`.

42 Alfréd Rényi. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, pages 547–561, 1961. URL: `https://projecteuclid.org/ebook/Download?urlid=bsmsp/1200512181&isFullBook=false`.

43 Benjamin Schumacher. Quantum coding. *Physical Review A*, 51(4):2738, 1995. `doi:10.1103/PhysRevA.51.2738`.

44 Sathyawageeswar Subramanian and Min-Hsiu Hsieh. Quantum algorithm for estimating α-Renyi entropies of quantum states. *Physical Review A*, 104(2):022428, 2021. `doi:10.1103/PhysRevA.104.022428`.

45 Gregory Valiant and Paul Valiant. Estimating the unseen: an $n/\log(n)$-sample estimator for entropy and support size, shown optimal via new CLTs. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, pages 685–694, 2011. `doi:10.1145/1993636.1993727`.

46 John von Neumann. *Mathematische Grundlagen der Quantenmechanik (Mathematical Foundations of Quantum Mechanics)*. Springer, 1932.

47 Qisheng Wang, Ji Guan, Junyi Liu, Zhicheng Zhang, and Mingsheng Ying. New quantum algorithms for computing quantum entropies and distances. *IEEE Transactions on Information Theory*, 70(8):5653–5680, 2024. `doi:10.1109/TIT.2024.3399014`.

48 Qisheng Wang and Zhicheng Zhang. Quantum lower bounds by sample-to-query lifting. ArXiv e-prints, 2023. `arXiv:2308.01794`.

49 Qisheng Wang and Zhicheng Zhang. Fast quantum algorithms for trace distance estimation. *IEEE Transactions on Information Theory*, 70(4):2720–2733, 2024. `doi:10.1109/TIT.2023.3321121`.

50 Qisheng Wang and Zhicheng Zhang. Time-efficient quantum entropy estimator via samplizer. ArXiv e-prints, 2024. The full version of this paper also includes references [8, 9, 14, 21, 24, 51, 55, 56, 57]. `arXiv:2401.09947`.

51 Qisheng Wang, Zhicheng Zhang, Kean Chen, Ji Guan, Wang Fang, Junyi Liu, and Mingsheng Ying. Quantum algorithm for fidelity estimation. *IEEE Transactions on Information Theory*, 69(1):273–282, 2023. `doi:10.1109/TIT.2022.3203985`.

52 Xinzhao Wang, Shengyu Zhang, and Tongyang Li. A quantum algorithm framework for discrete probability distributions with applications to Rényi entropy estimation. *IEEE Transactions on Information Theory*, 70(5):3399–3426, 2024. `doi:10.1109/TIT.2024.3382037`.

53 Youle Wang, Guangxi Li, and Xin Wang. Variational quantum Gibbs state preparation with a truncated Taylor series. *Physical Review Applied*, 16(5):054035, 2021. `doi:10.1103/PhysRevApplied.16.054035`.

54 Youle Wang, Benchi Zhao, and Xin Wang. Quantum algorithms for estimating quantum entropies. *Physical Review Applied*, 19(4):044041, 2023. `doi:10.1103/PhysRevApplied.19.044041`.

55 John Watrous. Limits on the power of quantum statistical zero-knowledge. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, pages 459–468, 2002. `doi:10.1109/SFCS.2002.1181970`.

56 John Watrous. *The Theory of Quantum Information*. Cambridge University Press, 2018.

57 Mark M. Wilde. *Quantum Information Theory*. Cambridge University Press, 2013.

58 John Wright. Private communication, 2022.

59 Jingxiang Wu and Timothy H. Hsieh. Variational thermal quantum simulation via thermofield double states. *Physical Review Letters*, 123(22):220502, 2019. `doi:10.1103/PhysRevLett.123.220502`.