

Quantum neural estimation of entropies

Ziv Goldfeld,¹ Dhruvil Patel², Sreejith Sreekumar^{3,1} and Mark M. Wilde¹

¹*School of Electrical and Computer Engineering, Cornell University, Ithaca, New York 14850, USA*

²*Department of Computer Science, Cornell University, Ithaca, New York 14850, USA*

³*Institute for Quantum Information, RWTH Aachen University, Aachen, Germany*



(Received 20 July 2023; accepted 27 February 2024; published 25 March 2024)

Entropy measures quantify the amount of information and correlation present in a quantum system. In practice, when the quantum state is unknown and only copies thereof are available, one must resort to the estimation of such entropy measures. Here we propose a variational quantum algorithm for estimating the von Neumann and Rényi entropies, as well as the measured relative entropy and measured Rényi relative entropy. Our approach first parametrizes a variational formula for the measure of interest by a quantum circuit and a classical neural network and then optimizes the resulting objective over parameter space. Numerical simulations of our quantum algorithm are provided, using a noiseless quantum simulator. The algorithm provides accurate estimates of the various entropy measures for the examples tested, which renders it a promising approach for usage in downstream tasks.

DOI: [10.1103/PhysRevA.109.032431](https://doi.org/10.1103/PhysRevA.109.032431)

I. INTRODUCTION

Entropy plays a fundamental role in quantifying uncertainty or the level of disorder in physical systems [1], with its roots tracing back to thermodynamics and statistical mechanics [2]. Since its inception, it has been central to studying various fields of physics and beyond, including cosmology, meteorology, physical chemistry, thermodynamics, and information science. One of its critical applications is in explaining the inefficiency of heat-work dynamics and the irreversibility of physical systems via the second law of thermodynamics [3]. Moreover, the concept of entropy increasing over time provides a basis for causality and the arrow of time [4], which are fundamental to our understanding of the universe.

From an information-theoretic perspective, Shannon developed the concept of entropy to measure the information content of a source [5], and Rényi later extended it to a single-parameter family of information measures [6]. Various discrepancy measures that are intimately connected to entropy were introduced throughout the years in the information theory and mathematical statistics communities, encompassing relative entropy [7] and f -divergences [8–10]. Quantum counterparts of these measures include the von Neumann entropy [2], quantum Rényi entropy, quantum relative entropy [11], and various quantum f -divergences [12–16]. In particular, the entropy of a quantum state quantifies how mixed it is, which is an important concept in quantum mechanics; cf. [17] for a detailed survey of the role of entropy in physics.

While the aforementioned measures are essential for determining the amount of information and correlation present in classical and quantum systems, the state of the system is often unknown in practice and is only accessible through sampling. In such situations, one must resort to estimation of the entropy or divergence terms, which is the task at the core of this work.

A. Contributions

Here we propose a variational quantum algorithm (VQA) for estimating the von Neumann and Rényi entropies of an unknown quantum state (see (4) and (9) for definitions). Our approach also extends to the measured Rényi relative entropy between two unknown quantum states, which in particular captures the measured relative entropy and fidelity as special cases [see (7), (1), and (13), respectively]. The proposed VQA is developed based on variational formulas for these entropy measures [18–20], which represent them as optimizations of certain objectives over the space of Hermitian operators. The key idea behind our approach is to parametrize this optimization domain using two models: (i) a classical neural network to parametrize the spectrum (eigenvalues) of a Hermitian operator and (ii) a quantum circuit to parametrize the eigenvectors. Having that, we present a sampling procedure for execution on a quantum computer, using which the parametrized objective is approximated by employing independent copies of the unknown quantum states. This results in an optimization problem over the parameter space, which is solvable using classical tools and whose solution yields an estimate of the entropy measure of interest.

Our VQA-based estimation technique, termed quantum neural estimation, builds upon the possibility of VQAs realizing quantum speedup [21,22] and is fully trainable using classical optimizers. Analogous variational methods based on neural network parametrization have seen immense success in recent years in estimating classical entropies and divergences [23–25]. The appeal of such neural estimators stems from their scalability to high-dimensional problems and large data sets, as well as their compatibility with modern gradient-based optimization techniques. We expect the proposed quantum neural estimation framework to enjoy similar virtues. Numerical simulations that demonstrate the performance of our

method on small-scale examples are provided, revealing accurate estimates and convergence of the training algorithm. We also discuss extensions of our method to quantum relative entropy [11] and sandwiched Rényi relative entropy [26,27], surfacing key challenges and outlining potential avenues to overcome them, which are left for future work.

B. Literature review

Estimation of the von Neumann and Rényi entropies has attracted significant interest throughout the years. A naive approach is to use quantum tomography to estimate the entire density matrix and then evaluate the entropic quantities based on the estimate. However, the time complexity of this approach is linear in the dimension of the state (or, equivalently, exponential in the number of qubits) [28–31] and is thus untenable. To address this problem, many quantum algorithms for estimating these quantities have been proposed, and their computational complexities have been investigated under different input models [32–43]. The authors of [34] studied the problem of estimating the von Neumann and Rényi entropies, given access to independent copies of the input quantum state. They demonstrated that the sample complexity (i.e., the number of independent copies of the quantum state) of this task grows exponentially with the number of qubits. A number of other papers have proposed methods for estimating integer Rényi entropies [44–49], with limits on the sample complexities also discussed in [34]. Entropy estimation under the quantum query model, i.e., where one has access to an oracle that prepares the input quantum state, was explored in [35,37]. These works established that the query complexity (i.e., the number of times the oracle is queried) for estimating the von Neumann and Rényi entropies is also exponential in the number of qubits.

Variational methods are often used in physics to find approximate solutions to problems that are hard to solve exactly [50], e.g., computing ground-state energies for quantum systems [51]. As such, they underlie various computational methods, including the Hartree-Fock method [52] and the configuration-interaction method [53]. VQAs are an extension of classical variational methods and have become a prominent research area in quantum computing [21,54]. VQAs are hybrid quantum-classical algorithms that are driven by classical optimizers and only call a quantum subroutine for tasks that are (presumably) hard for a classical machine. These algorithms apply the variational principle to find approximate solutions, by first preparing a quantum trial state and then optimizing its parameters using a classical computer. VQAs have now been applied to a variety of problems, including quantum simulation [55], optimization [56–58], and machine learning [59]. A VQA is also at the heart of the quantum neural estimation method developed herein.

Neural estimation techniques have been at the forefront of research on classical entropy and divergence estimation. The idea is to parametrize a variational form of the measure of interest by a neural network, approximate expectations by sample means, and optimize the resulting empirical objective over parameter space. The appeal of neural estimators stems from excellent scalability and computational efficiency observed in practice, which is in line with the success of neural

nets for language models [60,61], computer vision [62], and generative modeling [63–66]. Various neural estimators of Shannon’s mutual information [24,25], neural net distances [23,67], and classical f -divergences [68,69] have been developed and analyzed for their performance. Neural estimators of quantum entropy measures have not appeared in the literature yet, and we make strides towards harnessing this promising methodology in this work (see also the related work discussed below).

Note on related work. Recently, an independent and related work [70] introduced a method for estimating von Neumann entropy reminiscent of ours but with a few key differences. First, we treat several prominent quantum entropy measures—von Neumann entropy, Rényi entropy, measured relative (Rényi) entropy, and fidelity—while [70] only accounts for the von Neumann entropy. Second, we parametrize the space of all Hermitian operators using parametrized quantum circuits and classical neural networks, whereas [70] rewrites the variational formula for von Neumann entropy in terms of an optimization over parametrized quantum states and uses only the quantum circuit component. We believe that the incorporation of classical neural networks is crucial for scalable estimation of a broad class of quantum entropy and divergence measures.

We also note another independent and related work [71], employing quantum neural estimation as a combination of classical neural networks and parametrized circuits to estimate von Neumann and Rényi entropies. The applications considered in [71] are to classify phases of a Hamiltonian of physical interest, the XXZ chain, and to compare the performance of quantum neural estimation with a variational quantum state eigensolver.

II. QUANTUM ENTROPIES AND DIVERGENCES

In this section, we define the quantum entropy and discrepancy measures of interest and present their variational forms, which are subsequently used for the proposed quantum neural estimation method. Throughout, we fix $d \in \mathbb{N}$ and let ρ and σ be d -dimensional quantum states.

A. Measured relative entropy and von Neumann entropy

The measured relative entropy between ρ and σ is defined as [72,73]

$$D_M(\rho||\sigma) := \sup_{\mathcal{X}, (\Lambda_x)_{x \in \mathcal{X}}} \sum_{x \in \mathcal{X}} \text{Tr}[\Lambda_x \rho] \ln \left(\frac{\text{Tr}[\Lambda_x \rho]}{\text{Tr}[\Lambda_x \sigma]} \right), \quad (1)$$

where the supremum is over all finite sets \mathcal{X} and positive-operator-valued measures (POVMs) $(\Lambda_x)_{x \in \mathcal{X}}$ are indexed by \mathcal{X} .¹ More specifically, it is equal to the largest classical relative entropy between the probability distributions $(\text{Tr}[\Lambda_x \rho])_{x \in \mathcal{X}}$ and $(\text{Tr}[\Lambda_x \sigma])_{x \in \mathcal{X}}$ that result from performing

¹Reference [72] defined the quantity with an optimization over just projective measurements, but Ref. [73] generalized the definition to include an optimization over all possible measurements.

a measurement $(\Lambda_x)_{x \in \mathcal{X}}$ on the respective states ρ and σ .² As a consequence of the data-processing inequality for the classical relative entropy, it suffices to perform the optimization in (1) over rank-one measurements [20] (i.e., where each Λ_x , for $x \in \mathcal{X}$, is a rank-one operator). A variational form for $D_M(\rho\|\sigma)$ was derived in Eq. (11.69) in [18] (see also Lemma 1 and Theorem 2 in [20]), whereby

$$D_M(\rho\|\sigma) = \sup_H \{ \text{Tr}[H\rho] - \ln \text{Tr}[e^H \sigma] \} \quad (2)$$

$$= \sup_H \{ \text{Tr}[H\rho] - \text{Tr}[e^H \sigma] \} + 1, \quad (3)$$

where the supremum in (2) and (3) is over all $d \times d$ Hermitian operators, i.e., H such that $H = H^\dagger$. Equations (2) and (3) can be understood as a quantum generalization of the Donsker-Varadhan formula for classical relative entropy (Lemma 2.1 in [74]; see also Lemma 1 and Eq. (8) in [75]). Throughout our paper and numerical simulations, we use the variational formula in (3).

An important special case of the measured relative entropy is the von Neumann entropy

$$H(\rho) := -\text{Tr}[\rho \ln \rho]. \quad (4)$$

It is obtained from $D_M(\rho\|\sigma)$ by setting σ to be the maximally mixed state $\pi_d := I/d$, where I is the identity operator, namely, we have

$$H(\rho) = \ln d - D_M(\rho\|\pi_d), \quad (5)$$

which follows from Sec. 16 in [76] (see also [77]). Inserting the variational representation from (3) into (5), the von Neumann entropy is expressed in the following variational form:

$$H(\rho) = \ln d - 1 - \sup_H (\text{Tr}[H\rho] - \text{Tr}[e^H \pi_d]). \quad (6)$$

B. Measured Rényi relative entropy and Rényi entropy

The measured Rényi relative entropy of order $\alpha \in (0, 1) \cup (1, \infty)$ between ρ and σ is defined as [see Eqs. (3.116) and (3.117) in [78]]

$$D_{M,\alpha}(\rho\|\sigma) := \sup_{\mathcal{X}, (\Lambda_x)_{x \in \mathcal{X}}} \frac{1}{\alpha - 1} \ln \sum_{x \in \mathcal{X}} \text{Tr}[\Lambda_x \rho]^\alpha \text{Tr}[\Lambda_x \sigma]^{1-\alpha}, \quad (7)$$

where the supremum is over all finite sets \mathcal{X} and POVMs $(\Lambda_x)_{x \in \mathcal{X}}$. Akin to the measured relative entropy case, $D_{M,\alpha}(\rho\|\sigma)$ is equal to the largest classical Rényi relative entropy between probability distributions of the form $(\text{Tr}[\Lambda_x \rho])_{x \in \mathcal{X}}$ and $(\text{Tr}[\Lambda_x \sigma])_{x \in \mathcal{X}}$ that result from performing a measurement $(\Lambda_x)_{x \in \mathcal{X}}$ on ρ and σ . As before, it suffices to perform the optimization in (1) over rank-one measurements. The definition in (7) recovers the measured relative entropy $D_M(\rho\|\sigma)$ by taking the limit $\alpha \rightarrow 1$. The following variational representation was established in Lemma 3 and Theorem 4 in [20]:

$$D_{M,\alpha}(\rho\|\sigma) = \sup_H \left(\frac{\alpha}{\alpha - 1} \ln \text{Tr}[e^{(\alpha-1)H} \rho] - \ln \text{Tr}[e^{\alpha H} \sigma] \right). \quad (8)$$

Equation (8) is a quantum generalization of the more recently derived variational formula for classical Rényi relative entropy (Theorem 3.1 in [79]), i.e., the classical case arises by plugging in commuting density operators ρ and σ .

Instantiating σ in (7) as the maximally mixed state gives rise to the quantum Rényi entropy

$$H_\alpha(\rho) := \frac{1}{1 - \alpha} \ln \text{Tr}[\rho^\alpha], \quad (9)$$

which admits the form

$$H_\alpha(\rho) = \ln d - D_{M,\alpha}(\rho\|\pi_d). \quad (10)$$

Leveraging the variational formula in (8), we obtain

$$H_\alpha(\rho) = \ln d - \sup_H \left(\frac{\alpha}{\alpha - 1} \ln \text{Tr}[e^{(\alpha-1)H} \rho] - \ln \text{Tr}[e^{\alpha H} \pi_d] \right). \quad (11)$$

Another interesting special case of the measured Rényi relative entropy is obtained for $\alpha = 1/2$,

$$\begin{aligned} D_{M,1/2}(\rho\|\sigma) &= -\ln \inf_{\mathcal{X}, (\Lambda_x)_{x \in \mathcal{X}}} \left(\sum_{x \in \mathcal{X}} \sqrt{\text{Tr}[\Lambda_x \rho] \text{Tr}[\Lambda_x \sigma]} \right)^2 \\ &= -\ln F(\rho, \sigma), \end{aligned} \quad (12)$$

where the fidelity of ρ and σ is defined as [80]

$$F(\rho, \sigma) := \|\sqrt{\rho} \sqrt{\sigma}\|_1^2. \quad (13)$$

The equality in (12) was established in [81], which indicates that the fidelity of quantum states ρ and σ is achieved by a measurement, i.e., minimizing the classical fidelity of the distributions $\{\text{Tr}[\Lambda_x \rho]\}_{x \in \mathcal{X}}$ and $\{\text{Tr}[\Lambda_x \sigma]\}_{x \in \mathcal{X}}$ over all possible measurements. It thus follows from (8) and (12) that the negative logarithm of the fidelity has the following variational form:

$$-\ln F(\rho, \sigma) = -\inf_H \{ \ln \text{Tr}[e^{-H} \rho] + \ln \text{Tr}[e^H \sigma] \}. \quad (14)$$

Alternatively, by making use of the variational form in Eq. (20) in [20], we find that the root fidelity has the form

$$\sqrt{F(\rho, \sigma)} = \frac{1}{2} \inf_H (\text{Tr}[e^{-H} \rho] + \text{Tr}[e^H \sigma]), \quad (15)$$

which coincides with the expression from Eq. (6) in [19].

III. QUANTUM NEURAL ESTIMATION OF ENTROPIES

We develop variational estimators for the entropy and measured relative entropy terms defined in the preceding section. Our approach assumes access to a black-box procedure for repeatedly preparing the quantum states ρ and σ . The key idea is to parametrize the set of Hermitian operators in (11) using a classical neural network and a quantum circuit. The parametrization procedure, sampling step, and quantum algorithm to optimize the resulting objective are described next.

²Here we identify a probability distribution p on \mathcal{X} with the $|\mathcal{X}|$ -dimensional simplex vector $(p(x))_{x \in \mathcal{X}}$.

A. Measured relative entropy

Parametrization. The spectral theorem implies that any Hermitian operator H can be decomposed as

$$H = \sum_{i=1}^d \lambda_i |\lambda_i\rangle \langle \lambda_i|, \quad (16)$$

where $\lambda_1, \dots, \lambda_d$ are the eigenvalues and $|\lambda_1\rangle, \dots, |\lambda_d\rangle$ are the corresponding (orthonormal) eigenvectors.

Our first step is to approximate the set of eigenvalues using a classical neural network $f_w : \{1, \dots, d\} \rightarrow \mathbb{R}$ with a parameter vector $w \in \mathbb{R}^p$, where $p \in \mathbb{N}$. The neural network output $f_w(i)$, for $i \in \{1, \dots, d\}$, serves as a proxy for the eigenvalue λ_i . We keep the architecture of the neural net (viz., nonlinearity, width, depth, etc.) implicit, in order to maintain flexibility of the approach. Next we approximate the set of eigenvectors using a parametrized quantum circuit $U(\theta)$ with a parameter vector $\theta \in [0, 2\pi]^q$, where $q \in \mathbb{N}$. In practice, the total number of neural network and quantum circuit parameters, i.e., $p + q$, should scale like $O(\text{poly}(\log d))$ so that the optimization of the ensuing VQA is efficient with respect to the number of qubits specifying ρ and σ . See [82–84] for a similar approach for parametrizing the set of mixed quantum states and Eq. (3) in [82] for a similar approach to parametrizing Hermitian operators.

The above procedure defines a set of parametrized Hermitian operators, specified as

$$H(w, \theta) = \sum_{i=1}^d f_w(i) U(\theta) |i\rangle \langle i| U^\dagger(\theta), \quad (17)$$

where $(w, \theta) \in \mathbb{R}^p \times [0, 2\pi]^q$ and $\{|i\rangle\}_{i=1}^d$ denotes the computational basis. Using this parametrization, we approximate the measured relative entropy from below as

$$D_M(\rho \| \sigma) \geq \sup_{w, \theta} \{\text{Tr}[H(w, \theta)\rho] - \text{Tr}[e^{H(w, \theta)}\sigma]\} + 1, \quad (18)$$

which follows from (3) and because $\{H(w, \theta)\}_{w, \theta}$ is a subset of all $d \times d$ Hermitian operators.

To further simplify the parametrized objective and arrive at a form that lands well for sampling on a quantum computer, we define the following probability distributions on $\{1, \dots, d\}$:

$$p_\theta^\rho(i) := \text{Tr}[|i\rangle \langle i| U^\dagger(\theta) \rho U(\theta)], \quad (19)$$

$$q_\theta^\sigma(i) := \text{Tr}[|i\rangle \langle i| U^\dagger(\theta) \sigma U(\theta)]. \quad (20)$$

Using (17), (19), and (20), we can write the trace terms from the right-hand side of (18) as

$$\text{Tr}[H(w, \theta)\rho] = \sum_{i=1}^d p_\theta^\rho(i) f_w(i), \quad (21)$$

$$\text{Tr}[e^{H(w, \theta)}\sigma] = \sum_{i=1}^d q_\theta^\sigma(i) e^{f_w(i)}, \quad (22)$$

where (22) follows because

$$e^{\beta H(w, \theta)} = \sum_{i=1}^d e^{\beta f_w(i)} U(\theta) |i\rangle \langle i| U^\dagger(\theta) \quad (23)$$

for all $\beta \in \mathbb{R}$. Inserting (21) and (22) into the right-hand side of (18) yields an objective function that is readily estimated using a quantum computer, as described next.

Sampling. We use a quantum computer to sample from the distributions p_θ^ρ and q_θ^σ . As the procedures are similar, we only describe the steps for the former. We prepare the state ρ , act on it with the parametrized unitary $U^\dagger(\theta)$, and then measure in the computational basis to obtain a sample from p_θ^ρ . Repeating this process n times for each distribution, we obtain the samples $i_1(\theta), \dots, i_n(\theta)$ and $j_1(\theta), \dots, j_n(\theta)$ from p_θ^ρ and q_θ^σ , respectively. With that, we approximate the trace values in (21) and (22) by sample means, to arrive at the objective

$$\mathcal{L}_n(w, \theta) := 1 + \frac{1}{n} \sum_{\ell=1}^n f_w(i_\ell(\theta)) - \frac{1}{n} \sum_{\ell=1}^n e^{f_w(j_\ell(\theta))}, \quad (24)$$

and the resulting estimator is thus

$$\hat{D}_M^n : \sup_{w, \theta} \mathcal{L}_n(w, \theta). \quad (25)$$

Algorithm. We present a VQA for performing the optimization in (25). The algorithm uses the parameter-shift rule [85–87] to update the quantum circuit parameters and standard backpropagation [88] to update the weights of the neural network. The pseudocode of our algorithm is as follows.

ALGORITHM 1. VQA for estimating measured relative entropy.

-
- 1: **Input:** number of iterations K , learning rate η , number of samples n , and quantum circuits that prepare ρ and σ .
 - 2: $w^1 \leftarrow$ Random initialization in \mathbb{R}^p .
 - 3: $\theta^1 \leftarrow$ Random initialization in $[0, 2\pi]^q$.
 - 4: **for** $k \in \{1, 2, \dots, K\}$ **do**
 - 5: Evaluate $\nabla_\theta \mathcal{L}_n(w^k, \theta^k)$ using the parameter-shift rule.
 - 6: $\theta^{k+1} \leftarrow \theta^k + \eta \nabla_\theta \mathcal{L}_n(w^k, \theta^k)$.
 - 7: Evaluate $\nabla_w \mathcal{L}_n(w^k, \theta^k)$ using backpropagation.
 - 8: $w^{k+1} \leftarrow w^k + \eta \nabla_w \mathcal{L}_n(w^k, \theta^k)$.
 - 9: **Output:** $\mathcal{L}_n(w^{K+1}, \theta^{K+1})$ as an estimate of $D_M(\rho \| \sigma)$.
-

B. von Neumann entropy

The von Neumann entropy $H(\rho)$ can be estimated via a similar approach by appealing to (6) (in this context, see also Eq. (21) in [82]). This leads to the upper bound

$$H(\rho) \leq \ln d - 1 - \sup_{w, \theta} \{\text{Tr}[H(w, \theta)\rho] - \text{Tr}(e^{H(w, \theta)}\pi_d)\}, \quad (26)$$

and we thus take the estimator to be

$$\hat{H}_n : \ln d - 1 - \sup_{w, \theta} \left(\frac{1}{n} \sum_{\ell=1}^n f_w(i_\ell(\theta)) - \frac{1}{n} \sum_{\ell=1}^n e^{f_w(\tilde{j}_\ell(\theta))} \right), \quad (27)$$

where $\tilde{j}_\ell(\theta), \dots, \tilde{j}_\ell(\theta)$ are samples from the distribution \tilde{q}_θ , with

$$\tilde{q}_\theta(i) : q \text{ Tr}[|i\rangle \langle i| U^\dagger(\theta) \pi_d U(\theta)] = \frac{1}{d}. \quad (28)$$

Given that the distribution \tilde{q}_θ is simply the uniform distribution, a quantum computer is not required to sample from it, and so the term $\frac{1}{n} \sum_{\ell=1}^n e^{f_w(\tilde{j}_\ell(\theta))}$ in (27) can be evaluated exclusively by a classical sampling approach.

C. Measured Rényi relative entropy

We estimate the measured Rényi relative entropy via a similar approach using the variational representation in (8). A few minor modifications to the steps in Sec. III A are required, as delineated next. Using the parametrization of Hermitian operators given in (17), we obtain the variational lower bound

$$\begin{aligned} D_{M,\alpha}(\rho \parallel \sigma) &\geq \sup_{\mathbf{w}, \theta} \left(\frac{\alpha}{\alpha - 1} \ln \text{Tr}[e^{(\alpha-1)H(\mathbf{w}, \theta)} \rho] - \ln \text{Tr}[e^{\alpha H(\mathbf{w}, \theta)} \sigma] \right). \end{aligned} \quad (29)$$

With the same definitions of the distributions p_θ^ρ and q_θ^σ as in (19) and (20) and using (23), we rewrite the trace terms in (29) as

$$\text{Tr}[e^{(\alpha-1)H(\mathbf{w}, \theta)} \rho] = \sum_{i=1}^d p_\theta^\rho(i) e^{(\alpha-1)f_w(i)}, \quad (30)$$

$$\text{Tr}[e^{\alpha H(\mathbf{w}, \theta)} \sigma] = \sum_{i=1}^d q_\theta^\sigma(i) e^{\alpha f_w(i)}. \quad (31)$$

Following the same sampling step as in Sec. III A, we approximate the expected values in (31) and (32) by sample means and arrive at the estimator

$$\hat{D}_{M,\alpha}^n : \sup_{\mathbf{w}, \theta} \mathcal{L}_\alpha^n(\mathbf{w}, \theta), \quad (32)$$

where

$$\begin{aligned} \mathcal{L}_\alpha^n(\mathbf{w}, \theta) &:= \frac{\alpha}{\alpha - 1} \ln \frac{1}{n} \sum_{k=1}^n e^{(\alpha-1)f_w(i_k(\theta))} \\ &\quad - \ln \frac{1}{n} \sum_{k=1}^n e^{\alpha f_w(j_k(\theta))}. \end{aligned} \quad (33)$$

D. Rényi entropy

We briefly state the variational estimator for the Rényi entropy. From (11), along with our parametrization procedure, we obtain

$$\begin{aligned} H_\alpha(\rho) &\leq \ln d - \sup_{\mathbf{w}, \theta} \left(\frac{\alpha}{\alpha - 1} \ln \text{Tr}[e^{(\alpha-1)H(\mathbf{w}, \theta)} \rho] \right. \\ &\quad \left. - \ln \text{Tr}[e^{\alpha H(\mathbf{w}, \theta)} \pi_d] \right). \end{aligned} \quad (34)$$

The estimator that results from replacing expectations with sample means is thus

$$\begin{aligned} \hat{H}_\alpha^n &:= -\ln d - \sup_{\mathbf{w}, \theta} \left(\frac{\alpha}{\alpha - 1} \ln \frac{1}{n} \sum_{k=1}^n e^{(\alpha-1)f_w(i_k(\theta))} \right. \\ &\quad \left. - \ln \frac{1}{n} \sum_{k=1}^n e^{\alpha f_w(j_k(\theta))} \right), \end{aligned} \quad (35)$$

where $\tilde{j}_\ell(\theta), \dots, \tilde{j}_\ell(\theta)$ are samples from the uniform distribution. As in the case of estimating the von Neumann entropy, a quantum computer is not required to sample from the uniform distribution, and so the term $\frac{1}{n} \sum_{k=1}^n e^{\alpha f_w(\tilde{j}_k(\theta))}$ above can be evaluated using a classical sampling approach.

E. Root fidelity

Finally, we estimate the root fidelity $\sqrt{F(\rho, \sigma)}$ via a similar approach using the variational representation in (15). Employing the same parametrization of Hermitian operators given in (17), we obtain the variational upper bound

$$\sqrt{F(\rho, \sigma)} \leq \frac{1}{2} \inf_{\mathbf{w}, \theta} \{ \text{Tr}[e^{-H(\mathbf{w}, \theta)} \rho] + \text{Tr}[e^{H(\mathbf{w}, \theta)} \sigma] \} \quad (36)$$

and then rewrite the trace terms as

$$\text{Tr}[e^{-H(\mathbf{w}, \theta)} \rho] = \sum_{i=1}^d p_\theta^\rho(i) e^{-f_w(i)}, \quad (37)$$

$$\text{Tr}[e^{H(\mathbf{w}, \theta)} \sigma] = \sum_{i=1}^d q_\theta^\sigma(i) e^{f_w(i)}, \quad (38)$$

where p_θ^ρ and q_θ^σ are given in (19) and (20). Following the same sampling step as in Sec. III A, we approximate the expected values in (36) and (37) by sample means and arrive at the estimator for the root fidelity

$$\hat{F}_n := -\inf_{\mathbf{w}, \theta} \mathcal{L}_F^n(\mathbf{w}, \theta), \quad (39)$$

where

$$\mathcal{L}_F^n(\mathbf{w}, \theta) := \frac{1}{2n} \sum_{k=1}^n (e^{-f_w(i_k(\theta))} + e^{f_w(j_k(\theta))}). \quad (40)$$

Remark. Let us compare our VQA to existing variational estimators. Alternative variational methods for estimating fidelity have been proposed recently [89,90]. Some of the methods rely on Uhlmann's theorem [80], which involves a maximization. However, those approaches require purifications of the states ρ and σ in order to estimate their fidelity. Since state purifications are not easily attainable from samples, the fact that our algorithm avoids this need represents an advantage. Another variational method was proposed in Algorithm 7 in [90]. While this algorithm does not require purifications, the estimator employed there for the objective function is biased. Our approach, on the other hand, estimates the objective function from the right-hand side of (35) in an unbiased fashion, as given in (39).

IV. IMPLEMENTATION AND EXPERIMENTS

We numerically simulate our quantum neural estimation algorithm and assess its accuracy for estimating the measured relative entropy, von Neumann entropy, measured Rényi relative entropy, and Rényi entropy. For each quantity, we benchmark the performance against an estimator that avoids the use of a classical neural network and instead explicitly stores the eigenvalues. The latter approach is more expressive than the former; however, storing all eigenvalues is infeasible for systems with a large number of qubits. The quantum neural

estimation approach circumvents this overhead by parametrizing the eigenvalues with a reasonably sized classical neural network. The comparison between the two approaches aims to surface the potential performance loss that results for the neural net parametrization, but no noticeable drop in performance is seen in our simulations. For our simulations, we use PENNYLANE's default qubit as our quantum circuit simulator [91]. We next describe our methodology and then follow with the results.

A. Methodology

All our simulations follow the approach described below.

Preparing input quantum states. We consider examples for which the input states are either two-qubit or six-qubit mixed states. This is to give a sense of how the algorithms scale with the number of qubits. To prepare these input mixed states, we first prepare purifications thereof and then trace out the reference systems. Formally, let ρ_S be a mixed state of a system S and let $|\psi\rangle\langle\psi|_{RS}$ be a purification of ρ_S such that $\text{Tr}_R[|\psi\rangle\langle\psi|_{RS}] = \rho_S$, where R is a reference system. For our purposes, it is sufficient to consider the reference system to have the same dimension as the system, that is, in order to prepare a six-qubit mixed state, we prepare its 12-qubit purification. Overall, with this approach, we generate four input instances (a pair of two-qubit mixed states and a pair of six-qubit mixed states) for each quantity of interest.

Parametrized quantum circuits. Our algorithm samples from certain distributions by applying a quantum circuit $U^\dagger(\theta)$ to the input state [cf., e.g., (19)]. We use PENNYLANE's `qml.RandomLayers` subroutine to prepare a parametrized quantum circuit with a random structure and then keep this structure fixed throughout multiple runs of a specific simulation. We only change the structure when it is not sufficiently expressive, in the sense that the set of generated unitaries is not comparable to the set of all unitaries. The `qml.RandomLayers` subroutine creates a parametrized quantum circuit with multiple layers, each built by randomly selecting a subset of qubits and applying single-qubit or two-qubit parametrized quantum gates to them. For a two-qubit example, we use three layers, each with three to four parametrized quantum gates, for a total of 9 to 12 parameters. For a six-qubit example, we use five layers, each with three to four parametrized quantum gates, for a total of 15 to 20 parameters. To evaluate the gradient of a given cost function with respect to these parameters, we use the parameter-shift rule [85–87].

Classical neural network. The neural estimator uses a classical neural network to parametrize eigenvalues of Hermitian observables. For our simulations, we consider a 2-10-1 fully connected architecture with sigmoidal activations for two-qubit examples and a 6-30-1 fully connected architecture with sigmoidal activations for six-qubit examples. Here 2 and 1 are the input and output dimensions, respectively, for the two-qubit examples, and similarly, 6 and 1 are the input and output dimensions for the six-qubit examples. A sigmoid is not applied at the output since it restricts the values to $[0,1]$, while the eigenvalues that the neural networks aim to approximate may be outside this interval. Gradients are

evaluated using the PYTORCH automatic differentiation subroutine `torch.autograd`.

Number of runs. We estimate expectations of the form $\text{Tr}[\cdot]$, which appear in the objective functions of our quantities, using sample means (cf., e.g., (21) or (22)). We find that 100 samples suffice for our experiments, and so we use that sample size throughout. We plot the mean (solid or dotted line) and standard deviation (shaded area) of each experiment over ten runs.

B. Results

The simulation results for two-qubit and six-qubit examples are shown in Figs. 1 and 2, respectively.

von Neumann entropy. Figures 1(a) and 2(a) show that our estimator, both with and without the neural network, accurately retrieves the von Neumann entropy with only a small error. We note, however, that the convergence rate is quite slow in this case, as it takes around 600 and 2000 epochs to come within 1% error of the ground truth for two-qubit and six-qubit examples, respectively. To evaluate the ground truth, we use the closed-form expression of the von Neumann entropy in (4).

Measured relative entropy. Figures 1(b) and 2(b) plot the quantum neural estimation error curves for the case of measured relative entropy. To compute the ground truth, we use the fact that the measured relative entropy is equal to the quantum relative entropy when the quantum states are positive definite and commute with each other (see Proposition 5 in [20]). This perspective allows us to work around the issue that there is no closed-form expression for the measured relative entropy and use that for quantum relative entropy instead:

$$D(\rho\|\sigma) := \text{Tr}[\rho(\ln \rho - \ln \sigma)]. \quad (41)$$

Rényi entropy. Figures 1(c) and 2(c) present simulations for the Rényi entropy of orders $\alpha = 2.5$ and 1.05, respectively. The closed-form expression from (9) is used to compute the ground truth. Again, we observe that the neural estimator converges fast and accurately recovers the true Rényi entropy value.

Root fidelity. We simulate quantum neural estimation of the root fidelity [see (15)] and show the results in Figs. 1(d) and 2(d). The ground truth fidelity value is computed from (13). We see that the model that employs the neural network approaches the ground truth faster than the one that directly optimizes the eigenvalues. For the second instance shown in Fig. 1(d), the neural estimator fluctuates around the ground truth from the onset, while the estimator without the neural net approaches it from above. This effect persists in repeated runs of the simulation.

Measured Rényi relative entropy. Figures 1(e) and 2(e) plot the quantum neural estimation error curves for the case of measured Rényi relative entropy with $\alpha = 2.5$ and 1.05, respectively. To compute the ground truth, we use the fact that the measured Rényi relative entropy is equal to the sandwiched Rényi relative entropy when the quantum states are positive definite and commute with each other (see Theorem 6 in [20]). This perspective allows us to get around the fact that there is no closed-form expression for measured Rényi relative entropy and instead use that of sandwiched Rényi

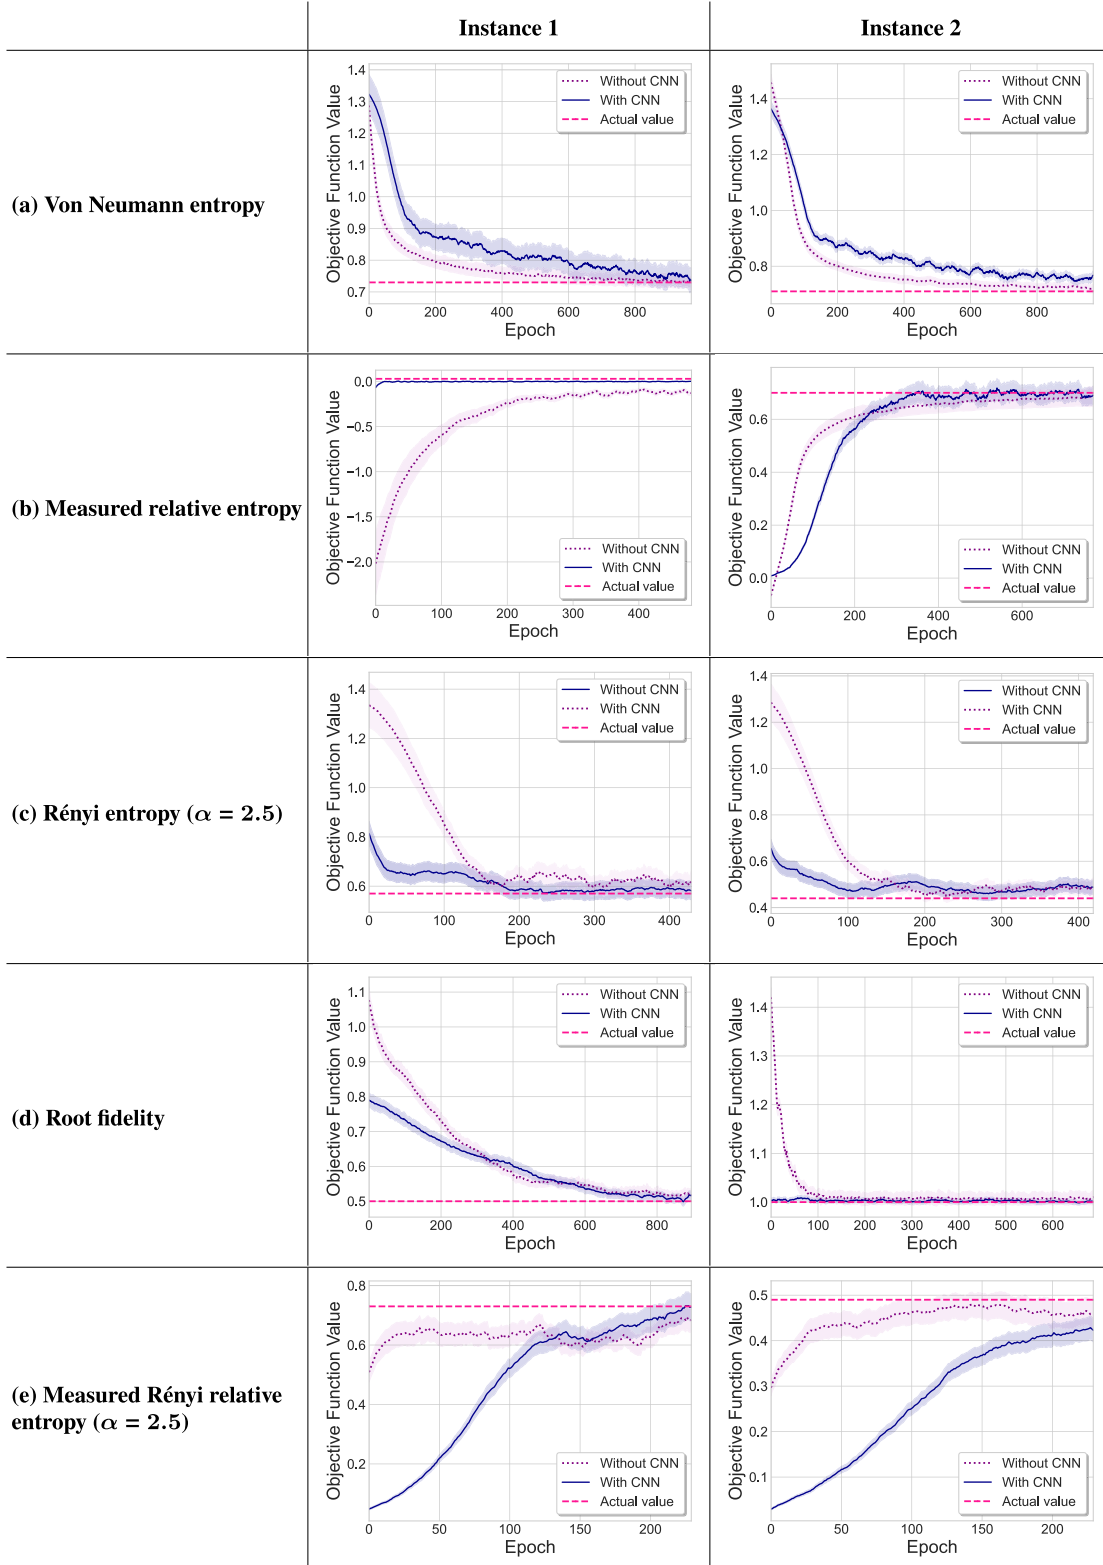


FIG. 1. Two-qubit examples. Convergence of the quantum neural estimation algorithm is shown for the (a) von Neumann entropy, (b) measured relative entropy, (c) Rényi entropy ($\alpha = 2.5$), (d) root fidelity, and (e) measured Rényi relative entropy ($\alpha = 2.5$). Two input instances are generated at random for each of these quantities. For each instance, results are presented both with (blue) and without (purple) the classical neural network approximation of eigenvalues. Each solid or dotted line shows the mean value of the estimate, the shaded region represents the standard deviation, and the dashed line shows the ground truth. PYTHON source files for reproducing these plots are available in the Supplemental Material [92].

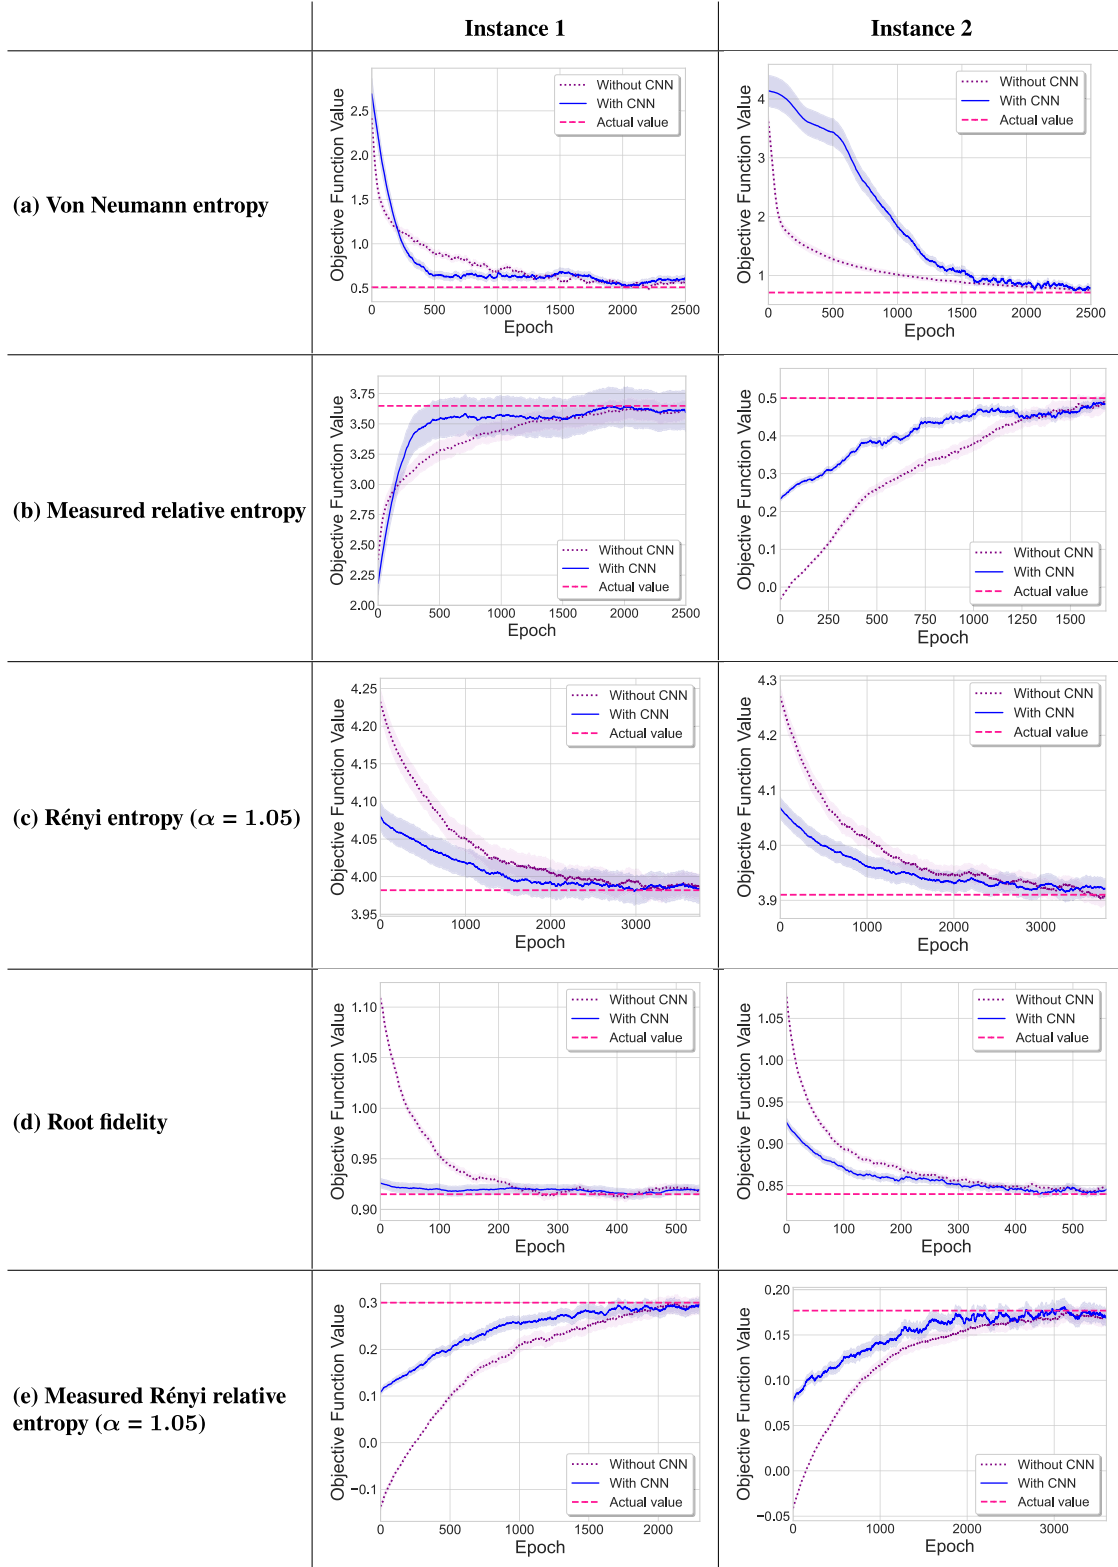


FIG. 2. Six-qubit examples. Convergence of the quantum neural estimation algorithm is shown for the (a) von Neumann entropy, (b) measured relative entropy, (c) Rényi entropy ($\alpha = 1.05$), (d) root fidelity, and (e) measured Rényi relative entropy ($\alpha = 1.05$). Two input instances are generated at random for each of these quantities. For each instance, results are presented both with (blue) and without (purple) the classical neural network approximation of eigenvalues. Each solid or dotted line shows the mean value of the estimate, the shaded region represents the standard deviation, and the dashed line shows the ground truth. PYTHON source files for reproducing these plots are available in the Supplemental Material [92].

relative entropy:

$$\tilde{D}_\alpha(\rho\|\sigma) := \frac{1}{\alpha-1} \ln \text{Tr}[(\sigma^{(1-\alpha)/2\alpha} \rho \sigma^{(1-\alpha)/2\alpha})^\alpha]. \quad (42)$$

V. CONCLUDING REMARKS AND FUTURE WORK

This work proposed a quantum neural estimation algorithm for estimating quantum information measures, spanning various entropies and measured relative entropies. Our estimator utilizes parametrized quantum circuits and classical neural networks to approximate a variational form of the measure of interest, which enables efficient estimation by sampling and optimization. Numerical experiments that validate the accuracy of the proposed approach on two-qubit and six-qubit settings were provided. Scalability of this approach to larger instances and extensions to other quantum entropies and divergences are important future avenues that we briefly discuss below. It is also of interest to obtain formal performance guarantees for our VQA in terms of its runtime and copy complexity, i.e., the number of quantum states required for the algorithm to achieve a given accuracy.

Scalability and the barren-plateau problem. To achieve the desired scalability, one would have to overcome the barren-plateau problem [93]. It refers to the phenomenon that the gradients of a cost function with respect to the quantum circuit parameters in VQAs tend to become exponentially small as the circuit depth and the number of qubits grow. This can have a significant impact on the performance and scalability of VQAs, limiting their ability to find optimal solutions efficiently. Addressing the barren-plateau problem is an active area of research, and various heuristic approaches towards mitigating its impact have been proposed [94–98]. In particular, we think that the approach of [97] is quite promising for overcoming the barren-plateau problem. The authors of [97] considered the variational quantum eigensolver problem and searched for an initial state by means of a classical tensor-network approach. They then used this pretrained state as an initialization for a variational quantum algorithm. These authors gave numerical evidence that this approach avoids the barren-plateau problem for examples consisting of up to 100 qubits. It is an important open question to determine how to apply the method of [97] to quantum neural estimation of entropies. The main difference between our setting and that of [97] is that the states ρ and σ are provided in sample form on a quantum computer, and so it is not clear how to simulate the measurement of observable expectations classically and efficiently via tensor networks.

In the present paper, we did not explicitly investigate the presence of barren plateaus in the context of our quantum neural estimation algorithm. Even though our numerical simulations show good performance, they do not involve sufficiently many qubits for observing barren plateaus. We speculate that the presence of a classical neural network may help to alleviate the barren-plateau problem. This is due to the fact that the cost functions considered in this work are dependent not only on the quantum circuit parameters but also on the neural network parameters. As a result, while gradients with respect to quantum circuit parameters may vanish exponentially with the number of qubits, this may not be the case for gradients with respect to neural network parameters

(which are independent of the number of qubits and the type of parametrized quantum circuit used). We leave for future work an in-depth study of this possibility and generally of the optimization landscapes. In addition, we plan to explore potential restrictions and formal limitations that may result from barren plateaus and design regularization methods to mitigate them. A promising method towards that end is to first use a tensor network to find a good initialization point for the quantum circuit parameters and then run the VQA. As mentioned above, this approach was recently shown to provide an advantage for the variational quantum eigensolver problem in terms of alleviating barren plateaus and improving the algorithm runtime [97].

Full-state tomography and performance guarantees. An appealing research avenue is to gauge how efficient quantum neural estimation is in comparison to an algorithm that estimates the state via tomography and then computes entropy based on the estimate. When the dimension of the states is small, as in the numerical examples presented herein, we anticipate that full-state tomography is more efficient than quantum neural estimation since the former does not require an optimization over classical or quantum circuit parameters. However, we expect that when the dimension is moderate to large, quantum neural estimation will be more efficient than tomography-based approaches. A formal comparison would require establishing the copy complexity of quantum neural estimation along with its associated runtime. A plausible approach for obtaining an upper bound on this quantity is to first quantify the various components of the overall error, such as approximation, estimation, and optimization errors, in terms of the number of the copies of state used and the circuit parameters. Then the parameters can be optimized to minimize the error of the algorithm and its runtime. We leave this aspect as an open question for future investigation.

Choice of Hamiltonian ansatz. The parametrization in (17) accounts for a general Hamiltonian ansatz. More specifically, the classical neural net and the quantum circuit parametrize the eigenvalues and eigenvectors, respectively, of the class of all Hermitian operators (Hamiltonians), without restricting the structure or complexity of these circuits. However, provided additional knowledge on the underlying quantum states, one may restrict the optimization to a smaller class of Hamiltonians without sacrificing performance, while possibly reducing computational or statistical complexity. For instance, if the eigenvalues or eigenvectors of considered states are from a bounded set or a restricted class of unitaries, then one may consider an appropriate ansatz that encodes this structure into the classical neural network and parametrized quantum circuit. Constraints may also arise organically from the considered quantum divergence. A pertinent example is the normalized trace distance $\frac{1}{2}\|\rho - \sigma\|_1 = \sup_{0 \leq M \leq I} \text{Tr}[M(\rho - \sigma)]$, in which case one may restrict the output of the classical neural net parametrizing M to $[0,1]$. This could potentially lead to computational and/or statistical gains, as is often observed in the classical setting.

Other quantum relative entropies. Another important direction is to extend the quantum neural estimation approach to the relative entropy [11] and the sandwiched Rényi relative entropy [26,27], defined in (40) and (41), respectively.

These quantities admit the following variational forms (see [20,99,100] and [20,101], respectively):

$$D(\rho\|\sigma) = \sup_H \{\text{Tr}[H\rho] - \ln \text{Tr}[\exp(H + \ln \sigma)]\} \quad (43)$$

$$= \sup_H \{\text{Tr}[H\rho] - \text{Tr}[\exp(H + \ln \sigma)]\} + 1, \quad (44)$$

$$\tilde{D}_\alpha(\rho\|\sigma) = \sup_H \left(\frac{\alpha}{\alpha-1} \ln \text{Tr}[H\rho] - \ln \text{Tr}[(H^{1/2} \sigma^{(\alpha-1)/\alpha} H^{1/2})^{\alpha/(\alpha-1)}] \right). \quad (45)$$

The difficulty in estimating these objective functions has to do with the second terms in (42)–(44), due to noncommutativity. A possible approach for evaluating the second term in (43) is to employ the Lie-Trotter product formula, which implies

$$\text{Tr}[\exp(H + \ln \sigma)] = \lim_{\ell \rightarrow \infty} \text{Tr}[(e^{H/\ell} \sigma^{1/\ell})^\ell]. \quad (46)$$

To realize (a finite approximation of) the right-hand side of (45), one could employ a variant of multivariate trace estimation [102] along with quantum singular value transformation [103] as a means to realize fractional powers of the density operator σ from block encodings of it. This exploration as well as the associated error analysis is left for future work.

An alternative approach could use the relations between the measured and unmeasured quantities [104],

$$D_M(\rho\|\sigma) \leq D(\rho\|\sigma) \leq D_M(\rho\|\sigma) + 2 \ln |\text{spec}(\sigma)|, \quad (47)$$

$$D_{M,\alpha}(\rho\|\sigma) \leq \tilde{D}_\alpha(\rho\|\sigma) \leq D_{M,\alpha}(\rho\|\sigma) + 2 \ln |\text{spec}(\sigma)|, \quad (48)$$

with the latter holding for all $\alpha \in (0, 1) \cup (1, \infty)$. In the above, $\text{spec}(\sigma)$ denotes the set of distinct eigenvalues of σ . Now, for $n \in \mathbb{N}$, using the fact that $|\text{spec}(\sigma^{\otimes n})| \leq (n+1)^{d-1}$, as well as the additivity of the unmeasured relative entropies, we obtain

$$\left| \frac{D_M(\rho^{\otimes n}\|\sigma^{\otimes n})}{n} - D(\rho\|\sigma) \right| \leq \frac{2(d-1)}{n} \ln(n+1), \quad (49)$$

$$\left| \frac{D_{M,\alpha}(\rho^{\otimes n}\|\sigma^{\otimes n})}{n} - \tilde{D}_\alpha(\rho\|\sigma) \right| \leq \frac{2(d-1)}{n} \ln(n+1). \quad (50)$$

We can then rewrite (3) and (8) as

$$D_M(\rho^{\otimes n}\|\sigma^{\otimes n}) = \sup_{\omega^{(n)} > 0} \{\text{Tr}[(\ln \omega^{(n)})\rho^{\otimes n}] - \text{Tr}[\omega^{(n)}\sigma^{\otimes n}]\} + 1, \quad (51)$$

$$D_{M,\alpha}(\rho^{\otimes n}\|\sigma^{\otimes n}) = \sup_{\omega^{(n)} > 0} \left(\frac{\alpha}{\alpha-1} \ln \text{Tr}[(\omega^{(n)})^{(\alpha-1)/\alpha} \rho^{\otimes n}] - \ln \text{Tr}[\omega^{(n)}\sigma^{\otimes n}] \right). \quad (52)$$

It follows from operator concavity of $x \mapsto \ln x$ and $x \mapsto x^{(\alpha-1)/\alpha}$ for $\alpha > 1$, the operator convexity of $x \mapsto x^{(\alpha-1)/\alpha}$ for $\alpha \in (1/2, 1)$, and the permutation invariance of the tensor-power states $\rho^{\otimes n}$ and $\sigma^{\otimes n}$ that in (50) and (51) it suffices to optimize only over permutation invariant observables $\omega^{(n)}$.

This significantly simplifies the optimization since the parameter space occupied by permutation invariant observables grows only polynomially in the number of copies n (this is related to some recent observations of geometric quantum machine learning [105,106]). However, in spite of this reduction, the error bounds in (48) and (49) scale as d/n , and thus the number of copies n must be larger than the dimension d to obtain a good approximation. As d is exponential in the number of qubits, this approach may not be feasible, and new ideas are needed to address this problem.

ACKNOWLEDGMENTS

We thank Paul Alsing, Daniel Koch, Saahil Patel, Shannon Ray, and Soorya Rethinasamy for insightful discussions. D.P. and M.M.W. acknowledge support from the National Science Foundation under Grant No. 1907615. Z.G. was partially supported by NSF Grants No. CCF-2046018, No. DMS-2210368, and No. CCF-2308446 and the IBM Academic Award. S.S. acknowledges support from the Excellence Cluster Matter and Light for Quantum Computing.

APPENDIX: CONNECTION TO THE DONSKER-VARADHAN FORMULA FOR RELATIVE ENTROPY AND NEURAL ESTIMATION OF RELATIVE ENTROPY

It is worthwhile to note the connection between the variational formulas for measured relative entropy in (2) and (3) and the Donsker-Varadhan (DV) formula for the classical relative entropy of two probability distributions p and q on $\{1, \dots, d\}$. For the latter, see Lemma 2.1 in [74] and Lemma 1 and Eq. (8) in [75]. Based on this connection, we also note here an additional interpretation of our neural quantum estimation approach for estimating measured relative entropy, namely, our approach can be understood as the parametrized quantum circuit searching for the optimal measurement to perform and the classical neural network searching for the optimal function in the DV formula.

To begin with, let us recall the definition of the classical relative entropy

$$D(p\|q) := \sum_{i=1}^d p(i) \ln \left(\frac{p(i)}{q(i)} \right). \quad (A1)$$

The DV variational formula for $D(p\|q)$ is

$$D(p\|q) = \sup_f \left(\sum_{i=1}^d p(i) f(i) - \ln \sum_{i=1}^d q(i) e^{f(i)} \right) \quad (A2)$$

$$= \sup_f \left(\sum_{i=1}^d p(i) f(i) - \sum_{i=1}^d q(i) e^{f(i)} \right) + 1, \quad (A3)$$

where the supremum is over all functions $f : \{1, \dots, d\} \rightarrow \mathbb{R}$. By including an optimization over all rank-one projective measurements, the measured relative entropy can thus be written as

$$D_M(\rho\|\sigma) = \sup_{\mathcal{P}} D(p_{\mathcal{P}}\|q_{\mathcal{P}}), \quad (A4)$$

where $\mathcal{P} := \{|\phi_i\rangle\langle\phi_i|\}_{i=1}^d$ denotes a rank-one projective measurement, i.e., such that $\langle\phi_i|\phi_j\rangle = \delta_{ij}$ and $\sum_{i=1}^d |\phi_i\rangle\langle\phi_i| = I$, and

$$p_{\mathcal{P}} := \text{Tr}[|\phi_i\rangle\langle\phi_i|\rho], \quad q_{\mathcal{P}} := \text{Tr}[|\phi_i\rangle\langle\phi_i|\sigma]. \quad (\text{A5})$$

Theorem 2 of [20] established the fact that it suffices to restrict the optimization in (A4) to rank-one projective measurements. For every orthonormal basis $\{|\phi_i\rangle\}_{i=1}^d$, there is a unitary U that relates it to the computational basis, i.e., $|\phi_i\rangle = U|i\rangle$ for all $i \in \{1, \dots, d\}$. As such, the optimization in (A4) is equivalent to

$$D_M(\rho\|\sigma) = \sup_U D(p_U\|q_U), \quad (\text{A6})$$

where

$$p_U := \text{Tr}[|i\rangle\langle i|U^\dagger \rho U], \quad q_U := \text{Tr}[|i\rangle\langle i|U^\dagger \sigma U]. \quad (\text{A7})$$

Then plugging into (A2) and (A3), we find that

$$\begin{aligned} D_M(\rho\|\sigma) &= \sup_U \sup_f \left(\sum_{i=1}^d p_U(i) f(i) - \ln \sum_{i=1}^d q_U(i) e^{f(i)} \right) \\ &= \sup_U \sup_f \left(\sum_{i=1}^d p_U(i) f(i) - \sum_{i=1}^d q_U(i) e^{f(i)} \right) + 1. \end{aligned} \quad (\text{A8})$$

Let us define

$$H = \sum_{i=1}^d f(i) U|i\rangle\langle i|U^\dagger \quad (\text{A9})$$

and observe that

$$\sum_{i=1}^d p_U(i) f(i) = \sum_{i=1}^d f(i) \text{Tr}[|i\rangle\langle i|U^\dagger \rho U] \quad (\text{A10})$$

$$= \text{Tr} \left[\sum_{i=1}^d f(i) U|i\rangle\langle i|U^\dagger \rho \right] \quad (\text{A11})$$

$$= \text{Tr}[H\rho], \quad (\text{A12})$$

$$\sum_{i=1}^d q_U(i) e^{f(i)} = \sum_{i=1}^d e^{f(i)} \text{Tr}[|i\rangle\langle i|U^\dagger \sigma U] \quad (\text{A13})$$

$$= \text{Tr} \left[\sum_{i=1}^d e^{f(i)} U|i\rangle\langle i|U^\dagger \sigma \right] \quad (\text{A14})$$

$$= \text{Tr}[e^H \sigma]. \quad (\text{A15})$$

Then by inspecting (A8)–(A15) and using the fact that every Hermitian operator can be written as in (A9), we note that the above is an alternate derivation of the formulas in (2) and (3). This alternate derivation instead uses the DV formula as a starting point. Indeed, one can also make this observation and connection by inspecting the proof of Lemma 1 in [20].

Now comparing (A8) and (18)–(22), we see that the parametrized optimization in (18) is equivalent to optimizing over parametrized unitaries in order to find the optimal measurement and optimizing over parametrized functions according to the DV formula. Thus, our development here shows that our approach for quantum neural estimation of the measured relative entropy is equivalent to using classical neural estimation according to the DV formula combined with a parametrized quantum circuit to find the optimal measurement.

-
- [1] A. Wehrl, General properties of entropy, *Rev. Mod. Phys.* **50**, 221 (1978).
 - [2] J. von Neumann, Thermodynamik quantenmechanischer gesamtheiten, *Nachr. Ges. Wiss. Göttingen Math.-Phys. Kl.* **1927**, 273 (1927).
 - [3] E. H. Lieb and J. Yngvason, A Guide to Entropy and the Second Law of Thermodynamics, in *Statistical Mechanics*, edited by B. Nachtergaele, J. P. Solovej, and J. Yngvason (Springer, Berlin, Heidelberg, 1998), pp. 353–363.
 - [4] J. J. Halliwell, J. Pérez-Mercader, and W. H. Zurek, *Physical Origins of Time Asymmetry* (Cambridge University Press, Cambridge, 1996).
 - [5] C. E. Shannon, A mathematical theory of communication, *Bell Syst. Tech. J.* **27**, 379 (1948).
 - [6] A. Rényi, On measures of entropy and information, *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, edited by J. Neyman [Berkeley Symp. Math. Stat. Probab. 4.1, 547 (1961)].
 - [7] S. Kullback and R. A. Leibler, On information and sufficiency, *Ann. Math. Stat.* **22**, 79 (1951).
 - [8] I. Csiszár, Information type measure of difference of probability distributions and indirect observations, *Stud. Sci. Math. Hungarica* **2**, 299 (1967).
 - [9] S. M. Ali and S. D. Silvey, A general class of coefficients of divergence of one distribution from another, *J. R. Stat. Soc. B* **28**, 131 (1966).
 - [10] T. Morimoto, Markov processes and the H -theorem, *J. Phys. Soc. Jpn.* **18**, 328 (1963).
 - [11] H. Umegaki, Conditional expectations in an operator algebra. IV. Entropy and information, *Kodai Math. Sem. Rep.* **14**, 59 (1962).
 - [12] D. Petz, Quasi-entropies for states of a von Neumann algebra, *Publ. Res. Inst. Math. Sci.* **21**, 787 (1985).
 - [13] D. Petz, Quasi-entropies for finite quantum systems, *Rep. Math. Phys.* **23**, 57 (1986).
 - [14] K. Matsumoto, A New Quantum Version of f -Divergence, in *Reality and Measurement in Algebraic Quantum Theory. NWW2015; Springer Proceedings in Mathematics & Statistics*, edited by M. Ozawa, J. Butterfield, H. Halvorson, M. Rédei, Y. Kitajima, and F. Buscemi (Springer, Singapore, 2018), Vol. 261.
 - [15] M. M. Wilde, Optimized quantum f -divergences and data processing, *J. Phys. A: Math. Theor.* **51**, 374002 (2018).
 - [16] C. Hirche and M. Tomamichel, Quantum Rényi and f -divergences from integral representations, [arXiv:2306.12343](https://arxiv.org/abs/2306.12343).
 - [17] A. Greven, G. Keller, and G. Warnecke, *Entropy* (Princeton University Press, Princeton, 2014), Vol. 47.

- [18] D. Petz, *Quantum Information Theory and Quantum Statistics* (Springer Science + Business Media, New York, 2007).
- [19] J. Watrous, Simpler semidefinite programs for completely bounded norms, *Chic. J. Theor. Comput.* **19**, 1 (2013).
- [20] M. Berta, O. Fawzi, and M. Tomamichel, On variational expressions for quantum relative entropies, *Lett. Math. Phys.* **107**, 2239 (2017).
- [21] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
- [22] J. Jäger and R. V. Krems, Universal expressiveness of variational quantum classifiers and quantum kernels for support vector machines, *Nat. Commun.* **14**, 576 (2023).
- [23] S. Arora, R. Ge, Y. Liang, T. Ma, and Y. Zhang, in *Proceedings of the 34th International Conference on Machine Learning*, edited by D. Precup and Y. W. Teh (PMLR, Cambridge, 2017), Vol. 70, pp. 224–232.
- [24] M. I. Belghazi, A. Baratin, S. Rajeshwar, S. Ozair, Y. Bengio, A. Courville, and D. Hjelm, in *Proceedings of the 35th International Conference on Machine Learning*, edited by J. Dy and A. Krause (PMLR, Cambridge, 2018), Vol. 80, pp. 531–540.
- [25] Y. Mroueh, I. Melnyk, P. Dognin, J. Ross, and T. Sercu, Improved mutual information estimation, *Proc. AAAI Conf. Artif. Intel.* **35**, 9009 (2021).
- [26] M. Müller-Lennert, F. Dupuis, O. Szechr, S. Fehr, and M. Tomamichel, On quantum Rényi entropies: A new generalization and some properties, *J. Math. Phys.* **54**, 122203 (2013).
- [27] M. M. Wilde, A. Winter, and D. Yang, Strong converse for the classical capacity of entanglement-breaking and Hadamard channels via a sandwiched Rényi relative entropy, *Commun. Math. Phys.* **331**, 593 (2014).
- [28] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu, *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, 2016), pp. 913–925.
- [29] J. Wright, How to learn a quantum state, Ph.D. thesis, Carnegie Mellon University, 2016, available at <http://reports-archive.adm.cs.cmu.edu/anon/2016/CMU-CS-16-108.pdf>.
- [30] R. O'Donnell and J. Wright, *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing* (Ref. [28]), pp. 899–912.
- [31] H. Yuen, An improved sample complexity lower bound for (fidelity) quantum state tomography, *Quantum* **7**, 890 (2023).
- [32] T. Li and X. Wu, Quantum query complexity of entropy estimation, *IEEE Trans. Inf. Theory* **65**, 2899 (2017).
- [33] Y. Subasi, L. Cincio, and P. J. Coles, Entanglement spectroscopy with a depth-two quantum circuit, *J. Phys. A: Math. Theor.* **52**, 044001 (2019).
- [34] J. Acharya, I. Issa, N. V. Shende, and A. B. Wagner, *Proceedings of the 2019 IEEE International Symposium on Information Theory, Paris, 2019* (IEEE, Piscataway, 2019), pp. 3012–3016.
- [35] S. Subramanian and M.-H. Hsieh, Quantum algorithm for estimating Rényi entropies of quantum states, *Phys. Rev. A* **104**, 022428 (2021).
- [36] A. Gilyén and T. Li, in *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, edited by T. Vidick, Leibniz International Proceedings in Informatics (LIPIcs) Vol. 151 (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, 2020), pp. 25:1–25:19.
- [37] T. Gur, M.-H. Hsieh, and S. Subramanian, Sublinear quantum algorithms for estimating von Neumann entropy, [arXiv:2111.11139](https://arxiv.org/abs/2111.11139).
- [38] Y. Wang, B. Zhao, and X. Wang, Quantum algorithms for estimating quantum entropies, *Phys. Rev. Appl.* **19**, 044041 (2023).
- [39] Q. Wang, J. Guan, J. Liu, Z. Zhang, and M. Ying, New quantum algorithms for computing quantum entropies and distances, [arXiv:2203.13522](https://arxiv.org/abs/2203.13522).
- [40] A. Gilyén and A. Poremba, Improved quantum algorithms for fidelity estimation, [arXiv:2203.15993](https://arxiv.org/abs/2203.15993).
- [41] Y. Wang, L. Zhang, Z. Yu, and X. Wang, Quantum phase processing and its applications in estimating phase and entropies, *Phys. Rev. A* **108**, 062413 (2023).
- [42] Q. Wang, Z. Zhang, K. Chen, J. Guan, W. Fang, J. Liu, and M. Ying, Quantum algorithm for fidelity estimation, *IEEE Trans. Inf. Theory* **69**, 273 (2023).
- [43] Q. Wang and Z. Zhang, Time-efficient quantum entropy estimator via sampler, [arXiv:2401.09947](https://arxiv.org/abs/2401.09947).
- [44] A. Barenco, A. Berthiaume, D. Deutsch, A. Ekert, R. Jozsa, and C. Macchiavello, Stabilization of quantum computations by symmetrization, *SIAM J. Comput.* **26**, 1541 (1997).
- [45] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf, Quantum fingerprinting, *Phys. Rev. Lett.* **87**, 167902 (2001).
- [46] S. J. van Enk and C. W. J. Beenakker, Measuring $\text{Tr } \rho^n$ on single copies of ρ using random measurements, *Phys. Rev. Lett.* **108**, 110503 (2012).
- [47] A. Elben, B. Vermersch, M. Dalmonte, J. I. Cirac, and P. Zoller, Rényi entropies from random quenches in atomic Hubbard and spin models, *Phys. Rev. Lett.* **120**, 050406 (2018).
- [48] T. Brydges, A. Elben, P. Jurcevic, B. Vermersch, C. Maier, B. P. Lanyon, P. Zoller, R. Blatt, and C. F. Roos, Probing Rényi entanglement entropy via randomized measurements, *Science* **364**, 260 (2019).
- [49] L. Leone, S. F. E. Oliviero, G. Esposito, and A. Hama, Phase transition in stabilizer entropy and efficient purity estimation, *Phys. Rev. A* **109**, 032403 (2024).
- [50] E. Gerjuoy, A. R. P. Rau, and L. Spruch, A unified formulation of the construction of variational principles, *Rev. Mod. Phys.* **55**, 725 (1983).
- [51] T. Cubitt and A. Montanaro, Complexity classification of local Hamiltonian problems, *SIAM J. Comput.* **45**, 268 (2016).
- [52] D. R. Hartree, The wave mechanics of an atom with a non-Coulomb central field. Part II. Some results and discussion, *Math. Proc. Cambridge Philos. Soc.* **24**, 111 (1928).
- [53] C. D. Sherrill and H. F. Schaefer, The Configuration Interaction Method: Advances in Highly Correlated Approaches, in *Advances in Quantum Chemistry*, edited by P.-O. Löwdin, J. R. Sabin, M. C. Zerner, and E. Brändas (Academic, New York, 1999), Vol. 34, pp. 143–269.
- [54] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwek, and A. Aspuru-Guzik, Noisy intermediate-scale quantum (NISQ) algorithms, *Rev. Mod. Phys.* **94**, 015004 (2022).
- [55] X. Yuan, S. Endo, Q. Zhao, Y. Li, and S. C. Benjamin, Theory of variational quantum simulation, *Quantum* **3**, 191 (2019).

- [56] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
- [57] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).
- [58] D. Patel, P. J. Coles, and M. M. Wilde, Variational quantum algorithms for semidefinite programming, [arXiv:2112.08859](https://arxiv.org/abs/2112.08859).
- [59] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature (London)* **549**, 195 (2017).
- [60] J. Devlin, M. Chang, K. Lee, and K. Toutanova, in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, 2019*, edited by J. Burstein, C. Doran, and T. Solorio (Association for Computational Linguistics, Stroudsburg, 2019), Vol. 1, pp. 4171–4186.
- [61] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, Language models are few-shot learners, in *Proceedings of the Annual Conference on Advances in Neural Information Processing Systems 33*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M.-F. Balcan, and H.-T. Lin (Curran, Red Hook, 2020).
- [62] C. Wang, A. Bochkovskiy, and H. Liao, *Proceedings of the 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition, Vancouver, 2023* (IEEE Computer Society, Los Alamitos, 2023), pp. 7464–7475.
- [63] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, Hierarchical text-conditional image generation with clip latents, [arXiv:2204.06125](https://arxiv.org/abs/2204.06125).
- [64] T. Karras, S. Laine, and T. Aila, *Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach* (IEEE, Piscataway, 2019), pp. 4401–4410.
- [65] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever, *Proceedings of the International Conference on Machine Learning* (Curran, Red Hook, 2021), pp. 8821–8831.
- [66] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. Denton, S. K. S. Ghasemipour, R. Gontijo-Lopes, B. K. Ayan, T. Salimans, J. Ho, D. J. Fleet, and M. Norouzi, in *Advances in Neural Information Processing Systems 36*, edited by A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho (Curran, Red Hook, 2022).
- [67] P. Zhang, Q. Liu, D. Zhou, T. Xu, and X. He, *Sixth International Conference on Learning Representations, Vancouver, 2018* (ICLR, La Jolla, 2018).
- [68] S. Sreekumar, Z. Zhang, and Z. Goldfeld, *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics* (PMLR, Cambridge, 2021), pp. 3322–3330.
- [69] S. Sreekumar and Z. Goldfeld, Neural estimation of statistical divergences, *J. Mach. Learn. Res.* **23**, 1 (2022).
- [70] M. Shin, J. Lee, and K. Jeong, Estimating quantum mutual information through a quantum neural network, *Quantum Info. Proc.* **23**, 1 (2024).
- [71] S. Lee, H. Kwon, and J. S. Lee, Estimating entanglement entropy via variational quantum circuits with classical neural networks, [arXiv:2307.13511](https://arxiv.org/abs/2307.13511).
- [72] M. J. Donald, On the relative entropy, *Commun. Math. Phys.* **105**, 13 (1986).
- [73] M. Piani, Relative entropy of entanglement and restricted measurements, *Phys. Rev. Lett.* **103**, 160504 (2009).
- [74] M. D. Donsker and S. R. S. Varadhan, Asymptotic evaluation of certain Markov process expectations for large time, I, *Commun. Pure Appl. Math.* **28**, 1 (1975).
- [75] X. Nguyen, M. J. Wainwright, and M. I. Jordan, Estimating divergence functionals and the likelihood ratio by convex risk minimization, *IEEE Trans. Inf. Theory* **56**, 5847 (2010).
- [76] C. A. Hein, Entropy in operational statistics and quantum logic, *Found. Phys.* **9**, 751 (1979).
- [77] M. M. Wilde, *Quantum Information Theory*, 2nd ed. (Cambridge University Press, Cambridge, 2017), Theorem 11.1.1.
- [78] C. Fuchs, Distinguishability and accessible information in quantum theory, Ph.D. thesis, University of New Mexico 1996, [arXiv:quant-ph/9601020](https://arxiv.org/abs/quant-ph/9601020).
- [79] J. Birrell, P. Dupuis, M. A. Katsoulakis, L. Rey-Bellet, and J. Wang, Variational representations and neural network estimation of Rényi divergences, *SIAM J. Math. Data Sci.* **3**, 1093 (2021).
- [80] A. Uhlmann, The “transition probability” in the state space of a $*$ -algebra, *Rep. Math. Phys.* **9**, 273 (1976).
- [81] C. A. Fuchs and C. M. Caves, Mathematical techniques for quantum communication theory, *Open Syst. Inf. Dyn.* **3**, 345 (1995).
- [82] G. Verdon, J. Marks, S. Nanda, S. Leichenauer, and J. Hidar, Quantum Hamiltonian-based models and the variational quantum thermalizer algorithm, [arXiv:1910.02071](https://arxiv.org/abs/1910.02071).
- [83] J.-G. Liu, L. Mao, P. Zhang, and L. Wang, Solving quantum statistical mechanics with variational autoregressive networks and quantum circuits, *Mach. Learn.: Sci. Technol.* **2**, 025011 (2021).
- [84] N. Ezzell, E. M. Ball, A. U. Siddiqui, M. M. Wilde, A. T. Sornborger, P. J. Coles, and Z. Holmes, Quantum mixed state compiling, *Quantum Sci. Technol.* **8**, 035001 (2023).
- [85] J. Li, X. Yang, X. Peng, and C.-P. Sun, Hybrid quantum-classical approach to quantum optimal control, *Phys. Rev. Lett.* **118**, 150503 (2017).
- [86] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, Quantum circuit learning, *Phys. Rev. A* **98**, 032309 (2018).
- [87] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, Evaluating analytic gradients on quantum hardware, *Phys. Rev. A* **99**, 032331 (2019).
- [88] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, 2016).
- [89] R. Chen, Z. Song, X. Zhao, and X. Wang, Variational quantum algorithms for trace distance and fidelity estimation, *Quantum Sci. Technol.* **7**, 015019 (2021).
- [90] S. Rethinasamy, R. Agarwal, K. Sharma, and M. M. Wilde, Estimating distinguishability measures on quantum computers, *Phys. Rev. A* **108**, 012409 (2023).
- [91] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin *et al.*, PennyLane: Automatic differentiation of hybrid quantum-classical computations, [arXiv:1811.04968](https://arxiv.org/abs/1811.04968).
- [92] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevA.109.032431> for PYTHON source files to reproduce figures.

- [93] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).
- [94] B. T. Kiani, S. Lloyd, and R. Maity, Learning unitaries by gradient descent, [arXiv:2001.11897](#).
- [95] T. L. Patti, K. Najafi, X. Gao, and S. F. Yelin, Entanglement devised barren plateau mitigation, *Phys. Rev. Res.* **3**, 033090 (2021).
- [96] H.-Y. Liu, T.-P. Sun, Y.-C. Wu, Y.-J. Han, and G.-P. Guo, Mitigating barren plateaus with transfer-learning-inspired parameter initializations, *New J. Phys.* **25**, 013039 (2023).
- [97] M. S. Rudolph, J. Miller, J. Chen, A. Acharya, and A. Perdomo-Ortiz, Synergy between quantum circuits and tensor networks: Short-cutting the race to practical quantum advantage, [arXiv:2208.13673](#).
- [98] X. Liu, G. Liu, J. Huang, H.-K. Zhang, and X. Wang, Mitigating barren plateaus of variational quantum eigensolvers, [arXiv:2205.13539](#).
- [99] D. Petz, A variational expression for the relative entropy, *Commun. Math. Phys.* **114**, 345 (1988).
- [100] H. Kosaki, Relative entropy of states: A variational expression, *J. Operat. Theory* **16**, 335 (1986).
- [101] R. L. Frank and E. H. Lieb, Monotonicity of a relative Rényi entropy, *J. Math. Phys.* **54**, 122201 (2013).
- [102] Y. Quek, E. Kaur, and M. M. Wilde, Multivariate trace estimation in constant quantum depth, *Quantum* **8**, 1220 (2024).
- [103] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, Phoenix* (ACM, New York, 2019), pp. 193–204.
- [104] M. Tomamichel, *Quantum Information Processing with Finite Resources: Mathematical Foundations* (Springer, Berlin, 2015), Vol. 5, Proposition 4.12.
- [105] Q. T. Nguyen, L. Schatzki, P. Braccia, M. Ragone, P. J. Coles, F. Sauvage, M. Larocca, and M. Cerezo, Theory for equivariant quantum neural networks, [arXiv:2210.08566](#).
- [106] J. J. Meyer, M. Mularski, E. Gil-Fuster, A. A. Mele, F. Arzani, A. Wilms, and J. Eisert, Exploiting symmetry in variational quantum machine learning, *PRX Quantum* **4**, 010328 (2023).